

THE UNIVERSITY OF CHICAGO

SEARCH FOR AN ELECTROWEAK TENSOR COUPLING WITH TRAPPED ${}^8\text{Li}$
 β -DECAY

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
DEPARTMENT OF PHYSICS

BY
LOUIS JAMES VARRIANO

CHICAGO, ILLINOIS

AUGUST 2023

Copyright © 2023 by Louis James Varriano
All Rights Reserved

For my family, friends, and Michael.

Table of Contents

LIST OF FIGURES	viii
LIST OF TABLES	xxi
LIST OF ACRONYMS	xxiii
ACKNOWLEDGMENTS	xxv
ABSTRACT	xxvi
1 INTRODUCTION	1
1.1 The Standard Model and weak interaction	1
1.2 Nuclear β -decay observables	4
1.3 Measurements of the $\beta - \nu$ angular correlation	8
1.3.1 Enhancement to sensitivity from β -delayed particle decays	10
1.3.2 Sensitivity to the Fierz interference term	13
1.4 Correlations between a and b	14
1.5 Structure and decay of ${}^8\text{Li}$ and ${}^8\text{B}$	22
1.6 $\beta - \nu$ angular correlation in trapped ${}^8\text{Li}$	25
1.6.1 Recent studies of the ${}^8\text{B}$ 2α spectrum	28
1.7 This work	32
2 THE BETA-DECAY PAUL TRAP MK IV	34
2.1 A brief history of the Beta-decay Paul Trap	34

2.2	Physics of a Paul trap	36
2.2.1	Ion cloud energy and spatial distributions	40
2.3	General BPT description	41
2.4	β scattering in the previous BPT	42
2.5	Design of the BPT Mk IV	45
2.5.1	Suitability of graphite for ultra-high vacuum	52
2.5.2	Estimated cloud size in the BPT Mk IV	54
2.5.3	RF resonator	55
2.6	Detectors	58
2.6.1	RF pickup	59
2.6.2	Notch filters	62
2.7	Data acquisition system	63
2.8	Calibration sources in the BPT Mk IV	66
3	^8Li DATA CAMPAIGN	71
3.1	^8Li production, capture, and transport	71
3.2	To the BPT Mk IV	76
3.2.1	Off-line commissioning	76
3.2.2	On-line commissioning	77
3.3	Trapping cycle of the BPT	81
3.3.1	Impact on background measurement	85
3.4	Experiment analysis	89
3.4.1	Simulations	89
3.4.2	Timing information	90
3.4.3	DAQ linearity	92
3.4.4	On-going data analysis	96
4	DSSD AND CALIBRATION SOURCE CHARACTERIZATIONS	97

4.1	Calibration α sources	98
4.2	Silicon detector response to α particles	98
4.2.1	Dead layers	102
4.2.2	Non-ionizing energy losses	103
4.2.3	Electron-hole pair statistics	107
4.2.4	Electronic noise and energy resolution	107
4.2.5	Silicon response non-linearity (pulse height defect)	108
4.3	Calibration source characterization	111
4.4	DSSD description	115
4.4.1	Resolution improvement from simultaneous e-h measurements	120
4.4.2	Interstrip gap region	122
4.4.3	Timing resolution	126
4.5	α beam characterization of a DSSD	128
4.5.1	Experiment details	130
4.6	Analysis and results	134
4.6.1	Data cuts and classification	135
4.6.2	Calibration	137
4.6.3	Primary DSSD dead layer	139
4.6.4	Additional DSSD dead layers	144
4.6.5	Back strip charge-sharing	153
4.6.6	Front strip charge-sharing	157
4.6.7	Silicon response non-linearity	162
4.6.8	Summary of results	164
5	CONCLUSIONS AND FUTURE DIRECTIONS	166
5.1	Future directions	167
5.2	Experiment upgrades and variations	168

Appendices

A	AVAILABILITY OF PRODUCTS FROM THIS WORK	172
B	EXPECTED SENSITIVITIES USING THE ASIMOV DATA SET	173
B.1	A simple example	175
C	MICROPARTICLE LINEAR QUADRUPOLE ION TRAP FOR DEMONSTRATIONS	180
D	COLLABORATION PICTURES	184
	REFERENCES	188

List of Figures

1.1	The Standard Model. From Wikipedia.	1
1.2	Four interconnected questions about fundamental interactions and the observed universe that can be explored with nuclear physics. From Ref. [3].	2
1.3	From Ref. [7]: Marginalized $\Delta\chi^2 \equiv \chi^2 - \chi_{\min}^2$ distributions for the Wilson coefficients C_V^- (left) and C_T^- (right), with (red) and without (blue) taking account the input from mirror beta decay.	7
1.4	Select experiments measuring $a_{\beta\nu}$ in nuclei and neutrons.	10
1.5	From Ref. [40]: Coincidence data for the most populated ^{20}Na β -delayed α level showing the energy shift as a function of the measured β -particle energy with the angle between the α and β indicated. The curves show the expected kinematic shift for a pure Fermi (F) and a pure Gamow-Teller (G) transition.	12
1.6	Comparison of $E_e, \theta_{\beta\nu}$, and p_r spectra for various values of C_T^+/C_A^+ assuming $C_T^- \equiv 0$; these are for β^- -decay with $Q = 13$ MeV. Vertical axis is relative difference from the SM ($C_T^+ = 0$), shown in blue, with the a contribution in orange and the b contribution in green. Note that, in the $\theta_{\beta\nu}$ and p_r spectra, as $C_T^+ \rightarrow 0$, the b contributes more compared to the a term. There are near-exact cancellations between the a and b terms for these two spectra.	16
1.7	At each value of C_T^+/C_A^+ (assuming $C_T^- \equiv 0$), the absolute value of the relative difference between the SM and BSM p_r spectra is found, then averaged across p_r . The average relative difference is highest for $Q \sim 2$ MeV, leading to the highest sensitivity. Points of near cancellation are observed, which occur when the a and b terms approximately cancel; these are dependent on the Q value.	17

1.8	Median expected sensitivities at 68% CL (dashed) and 95% CL (solid) for $Q = 13$ MeV with 1 million and 4 million events. The colors refer to measurements of the different spectra: E_e (green), p_r (pink), $\theta_{\beta\nu}$ (orange). The blue curve is a measurement of the joint spectrum, either $(E_e, \theta_{\beta\nu})$ or (E_e, p_r) , which give identical limits as expected. See text for additional details.	18
1.9	Median expected sensitivity at 68% CL to (a) C_T^+/C_A^+ (assuming $C_T^-/C_A^+ \equiv 0$) and (b) C_T^-/C_A^+ (assuming $C_T^+/C_A^+ \equiv 0$) as a function of the number of events for $Q = 13$ MeV.	19
1.10	Median expected sensitivity at 68% CL to (a) $ C_T^+/C_A^+ $ (assuming $C_T^-/C_A^+ \equiv 0$) and (b) $ C_T^-/C_A^+ $ (assuming $C_T^+/C_A^+ \equiv 0$) as a function of Q value for 100 million events.	20
1.11	Median expected sensitivity at 68% CL for two different Q values with 4 million events, demonstrating the differing results obtained with different spectra and analysis approaches. See text for details.	21
1.12	Median expected sensitivities at 68% CL for different Q values with 4 million events; the limit is (nearly) symmetric about zero and hence the absolute value is shown. (a) Expected sensitivity for $ C_T^+/C_A^+ $ assuming $C_T^-/C_A^+ \equiv 0$, and (b) expected sensitivity for $ C_T^-/C_A^+ $ assuming $C_T^+/C_A^+ \equiv 0$	22
1.13	Decay scheme of ${}^8\text{Li}$ and ${}^8\text{B}$ with only relevant levels shown.	23
1.14	From Ref. [38]: Level scheme of ${}^8\text{Be}$ with proposed intruder state. Two different R-matrix fits are shown (a) without the intruder state and (b) with the intruder state. The fit is equally good and cannot distinguish between the two cases. In light blue, the selected excitation energy range around the 3 MeV peak for the angular correlation measurement to reduce possible contamination from the intruder state.	24
1.15	From Ref. [38]: General setup of the BPT measurement of ${}^8\text{Li}$ decays.	26

1.16	From Ref. [38]: Final α energy difference spectrum which is compared to simulation (black curve) to place limits on tensor couplings. The blue curve shows the expected spectrum in the case of purely tensor couplings. n.b. A +0.0005 correction was added to the result shown in the figure to account for possible contamination from the intruder state; see Ref. [38] for details.	27
1.17	Effects involved in the final distribution of observables (the α energy difference spectrum) for the BPT measurement.	33
2.1	Illustration of a quadrupole potential with some equipotential lines shown. Solid lines are positive potential; dashed lines are negative. The straight lines along the diagonals are at zero potential.	38
2.2	Maxwell-Boltzmann distributions of energies at liquid nitrogen temperature (80 K, blue) and at room temperature (300 K, orange). (a) Probability density. (b) Cumulative distribution with the energy extents of 90%, 95%, and 99% of particles at 80 K indicated.	40
2.3	Effect on β scattering of material and thickness choices for the previous BPT RF hood from simulations. Unluckily, the 0.006" stainless steel (blue) design used in the previous BPT actually maximized the β scattering.	44
2.4	Comparison of major design elements between the previous BPT, as used in Refs. [38, 75], and the BPT Mk IV, performed in Autodesk Inventor [105]. Unlabeled components are made of stainless steel (grey) or alumina ceramic (white). (a) and (b) compare a single quadrant of the electrodes with the beam axis horizontal. The different DC voltage regions in the Mk IV are lettered in yellow. (c) and (d) compare the electrodes along the beam axis; note that the cross section in the BPT Mk IV changes along the beam axis—cf. (b)—and only the rods are present in the center trapping region. The voltage pattern applied the Mk IV is labeled in yellow. (e) and (f) compare a sectioned view of the detector and RF shielding assemblies. See text for additional details.	46

2.5	Sectioned view of the final design of the BPT Mk IV performed in Autodesk Inventor [105]. See Figure 2.4 and text for additional details.	47
2.6	Design process for the BPT Mk IV to define the region where the electrodes would not obscure line of sight from the ion cloud to the DSSDs. The angle of the new RF hood design is additionally defined by this sketch.	49
2.7	Comparison between 0.5 mm rod designs (blue) and 1.0 mm rod (orange) designs; (a) separation between pairs of electrode rods as calculated from the allowed electrode region; (b) radius at which terms higher than quadratic contribute $\geq 2\%$ to the potential; (c) required voltage relative to previous BPT to achieve same radial potential.	51
2.8	The BPT Mk IV during assembly. (a) Full view of the trap with liquid nitrogen circulation lines shown. (b) A view of the center trapping region. (c) Support structure for the glassy carbon rods, which are held by friction in a small collar. A spring, indicated by a white arrow in lower left, is inserted into each alignment hole to tension the rods.	53
2.9	Estimated ion cloud sizes from a simplified treatment of the potentials, described in the text. (a) Radial cloud σ as a function of RF voltage (V_{pp}), ion mass to charge ratio, and RF frequency. Dashed lines indicate some corresponding q values. (b) Mathieu parameter q as a function of RF voltage, ion mass to charge ratio, and RF frequency. The vertical axis of (a) and (b) has units of $M/Q \cdot f^2$ in $\text{amu} \cdot \text{MHz}^2$. (c) Axial cloud σ as a function of DC voltages applied to the inner and rod electrodes (see section 2.5.	56
2.10	Circuit diagram of the resonator for the BPT Mk IV. See text for details.	57
2.11	Waveforms from the output of a preamplifier showing the size of RF pickup relative to an α pulse. RF mitigation including notch filters were in place.	60

2.12	Illustration of the phase of RF pickup across the DSSD surface. Front strips are parallel with the electrodes and the beam axis. Electrodes shown are from the previous BPT.	60
2.13	Measured RF amplitude and phase from the output of a preamplifier on the front strips of a DSSD.	61
2.14	Circuit diagram of notch filter; notch filters are only used with the front strips. .	62
2.15	Schematic of key features of the data acquisition chain from the DSSDs. See text for details.	64
2.16	Shaping amplifier output with a shaping time of (a) 2 μ s and (b) 4 μ s, showing that the RF pickup is effectively integrated out at 4 μ s. Images are from an oscilloscope showing energy output (yellow) and trigger output (blue).	65
2.17	Calibration α source locations in the previous BPT. Reproduced from Ref. [86].	67
2.18	Example of ^{148}Gd source shadowing in the previous BPT. Two sources each illuminated half of the detector, hence the higher counts in the middle of the detector. Front strip shadowing from the RF hood, indicated by the red arrows, was unpredictable due to their hand-made construction. Back strip shadowing, indicated by the blue arrow, made calibration difficult on these strips in previous data sets.	68
2.19	Source holder for the BPT Mk IV. Source foils (one ^{148}Gd , one ^{244}Cm) are sandwiched between two 0.002" thick stainless steel pieces, held together by M1.2 screws. The 1.5 mm diameter holes in the plates allow the source foils to see the detector. See Figure 2.4f for the mounting position.	68
2.20	Distribution of counts across one of the DSSDs from its ^{148}Gd source. Front strip 24 is dead on this detector, hence no counts. It is not clear where the two small peaks near the front come from. See text for details.	70
3.1	ATLAS facility layout. The BPT is located in the Triangle Room, labeled as "Trap Area." Other production elements are located in Area II. From Ref. [130].	72

3.2	From Ref. [41]: An overview of the beamline used to create and transfer ^8Li to the BPT.	73
3.3	The BPT Mk IV installed in its vacuum chamber, with major components labeled.	77
3.4	Axial DC potential of the BPT Mk IV during the commissioning ^8Li experiment, as determined from SIMION using the applied voltages from the experiment. Applied voltages are shown in Table 3.1. The potential during trapping (blue), during a capture pulse (orange) and during ejection (green) are shown. The widths of the DC electrodes are shown as grey bands. The center region has no DC electrode, with the potential supplied directly onto the rods.	79
3.5	^8Li ion cloud cooling as a function of cycle time with (a) showing axial cooling and (b) showing radial cooling. The first few ms of data are unusable as noted in the text. The two different DSSDs pairs are shown separately in each plot, (right-left, R-L, and top-bottom, T-B). There is good agreement between the two pairs, and discrepancies may be due to non-functioning strips, which have not been accounted for.	80
3.6	Trap cycle model comparison with the ^8Li , including a larger first injection and a finite trap lifetime. Events where two α were detected are used. See text for details.	85
3.7	Untrapped ion background from current cycle and previous cycle, demonstrating the need to account for untrapped ions leftover from the previous measurement cycle.	87
3.8	Background scale factors as a function of the total number of BPT injections and the number of injections to the IS. The other parameters are from the most recent ^8Li run; the selected number of IS captures was 4, which gives a ratio slightly higher than the uniform background assumption (black dashed line) but results in a higher data rate efficiency.	88

3.9	Time of flight for α (blue) and β (orange) particles as a function of kinetic energy, assuming a 7.5 cm distance from the trap center to the DSSD surface.	91
3.10	Timing information from triple correlation events (only one run is shown). Timing is relative to the trigger of the first particle; a considerable walk of up to 5 μ s can be seen as a function of particle energy. The blank space around 700 keV comes from a particle identification cut.	92
3.11	Time structure of a pulser test. Vertical axis has arbitrary units but scales with the number of triggers (blue). The green line is used to determine the time cuts (orange lines) to separate the pulses.	94
3.12	(a) Linearity evaluation from a pulser test, showing significant deviation about ADC channel 4000. (b) Residuals show a strange behavior that comes from pulser attenuation settings used, as detailed in the text.	95
3.13	(a) Linearity evaluation from a pulser test restricted only to peaks with identical attenuation settings (colors). (b) Residuals.	95
4.1	Intermediate and final distributions for a ^{148}Gd source, assuming a 100 nm Si dead layer and 8 keV electronic noise σ . See text for details.	101
4.2	Total stopping powers (including both electronic and nuclear stopping) and projected ranges and range straggling from SRIM for α particles of materials relevant to this work.	103
4.3	Helium stopping in water vapor and solid, suggesting a correction factor of 0.95 for solid water. From Ref. [151]	104
4.4	Non-ionizing energy loss distributions for several α energies simulated by SRIM.	105
4.5	Non-ionizing energy loss distribution of a 1500 keV α from this work and Ref. [148].	106

4.6	The number of electron-hole pairs N_{eh} produced by an alpha in silicon depends non-linearly on its initial energy, according the model of Ref. [152]. (a) N_{eh} as a function of initial alpha energy according to the linear and non-linear models. (b) The relative difference in N_{eh} between these two models. (c) A hypothetical linear calibration to energy with ^{148}Gd and ^{244}Cm ignoring the pulse height defect (blue), demonstrating that this non-linear effect can have a significant impact on the BPT calibration. The calibration corrected for the pulse height defect (orange) is much improved.	110
4.7	Setup of calibration sources and PIPS detector for source characterization. The sources are mounted to a rotatable post; other hardware is present from a parallel experiment.	113
4.8	Illustration of the source distribution model adopted for the calibration sources: two parameters were needed, the thickness and length scale of the exponential distribution. The source dead layer with no active material (green) was found to be unnecessary and set to zero.	114
4.9	Decrease in the peak energy with the modeled calibration source distribution compared to a monoenergetic α , assuming a 100 nm Si detector dead layer and 8 keV electronic noise σ	115
4.10	A DSSD mounted to the detector assembly of the BPT; junction side shown. . .	116
4.11	Dimensions of the DSSD junction side surface features, not to scale. Approximate dimensions only; see Table 4.3.	118

4.12	DSSD response from a single pixel (single front strip and single back strip) in ADC channels. Two calibration alpha sources were used as indicated on the figure; a low-energy background peak from cosmic rays is also present. A: normal calibration line; B: back strip charge-sharing from electron cloud drift; C: front strip charge-sharing from interstrip gap; D: incomplete charge collection, presumably without front strip charge-sharing and with opposite polarity pulse on neighboring strip.	118
4.13	Improvement in resolution from taking a weighted average compared to a simple average, assuming two independent Gaussian noise contributions.	120
4.14	Resolution improvement from a simple average of front strip and back strip energies demonstrated with a 1550 keV α beam.	122
4.15	DSSD interstrip gap using optical interferometry. This picture was taken at the far edge of a strip, hence the rounded corners. Note that SiO ₂ is optically transparent at the wavelength of light used and therefore appears to have a lower than expected height. (Credit to Tsviki Hirsh and Tal Zaharoni, Soreq NRC, Israel.)	123
4.16	Heights and widths of the DSSD interstrip gap features using a Bruker DektakXT profilometer [164]. (Credit to Tsviki Hirsh and Tal Zaharoni, Soreq NRC, Israel.)	124
4.17	(a) Distortion of the electric field in the interstrip gap region can cause holes to be trapped at the SiO ₂ interface. This diagram of simulated data demonstrates how the different regions of the field distortion lead to different pulse heights in the neighboring strips. Taken from Ref. [165], color annotations added. (b) Demonstration of inverted polarity signals in neighboring front strips using a micro-proton beam across the interstrip gap. The lower figure shows how the back strip also suffers from a loss of electrons that are taken up by the SiO ₂ interface. Taken from Ref. [166].	127
4.18	Set-up of experiment #1905 (see text for details).	131

4.19	A defocused α beam as imaged by the micro-channel plate (MCP)-based beam monitor and viewed on an oscilloscope.	132
4.20	Model of the double-sided silicon strip detector (DSSD) vacuum chamber assembly. The differentially-pumped rotary platform is at the top. Cooling is provided by an attached liquid nitrogen dewar (not shown) and a copper rod along the rotation axis. This assembly allows in-vacuum rotation while maintaining cooling. The detector mask is shown at the front.	133
4.21	Beam spots at 1500 keV and 7500 keV. Pixels inside of the yellow box are fully illuminated by the beam, along with all charge-sharing, and were primarily used for the analysis.	135
4.22	Total energy uncertainty at 1σ from calibration for the α beam characterization of a DSSD for each strip.	140
4.23	Mean front-back energy difference for single pixel events.	140
4.24	Dependence of the peak energy as a function of angle with the nominal 1500 keV α beam.	142
4.25	Stopping power comparison between a layer composed of only silicon and one composed of silicon and water.	144
4.26	Energy spectrum for 1552.0 ± 1.3 keV α particles that deposited charge in a single front strip and single back strip (no charge-sharing) showing the contributions from each of the different dead layers. All layers were also assumed to have 85 nm of H ₂ O on top of them. This spectrum is from a single pixel.	146
4.27	Thickness determination for (a) the periphery Si layer and (b) the Al layer using the nominal 1500 keV α beam, assuming 85 nm of H ₂ O on top. Each color represents a different pixel. The points were fit independently, and colored horizontal lines show the result for a simultaneous fit across the different angles; uncertainties on these are statistical only. Black horizontal lines show the average and total uncertainty.	147

4.28	Determination of the SiO ₂ dead layer thickness from the back energies of front strip charge-sharing. Uncertainties include statistical uncertainty and the systematic uncertainty from the primary dead layer. A spline (black) is produced from the average with uncertainties estimated by bootstrapping.	149
4.29	Thickness of the periphery p+ implant from the back energies of front strip charge-sharing. Results agree well with the thickness determined from the single pixel spectra at 1500 keV nominal beam energy.	150
4.30	Fraction of events passing through the Al electrode dead layer.	151
4.31	(a) Fraction of all events passing through the SiO ₂ layer. (b) Fraction of all events passing through the SiO ₂ layer that deposit energy in a single front strip. Compare to Figure 4.40 and see text for details.	152
4.32	Fraction of events that pass through the periphery p+ layer and deposit energy in a single pixel. Compare to Figure 4.40 and see text for details.	153
4.33	Spectrum of back strip charge-sharing (f1b2) events compared to single pixel events for the nominal 1500 keV α beam. These spectra are the sum of the four pixels of interest.	154
4.34	Average charge loss in back strip charge-sharing (f1b2) events as a function of α energy, corrected for calibration disagreements between individual front and back strips (see Figure 4.23). A simple line (in black) is fit to the mean of the individual pixels, with uncertainties shown (dashed).	155
4.35	Back strip energy spectrum of a single pixel for back strip charge-sharing events with (a) nominal 1500 keV α and (b) nominal 8500 keV α . (Note that this spectrum is the summation of several pixels to increase statistics.) A model based on the beta distribution (Eq. 4.7) is shown in orange; see text for details.	156
4.36	Back strip charge sharing fraction; a smoothed spline in black is overlaid.	157

4.37	Spectrum of front strip charge-sharing (f2b1) events compared to single pixel events for (a) the nominal 1500 keV α beam and (b) the nominal 7500 keV α beam. These spectra are the sum of the four pixels of interest.	158
4.38	Front strip energy spectrum of a single pixel for front strip charge-sharing events with nominal 8500 keV α . (Note that this spectrum is the summation of several pixels to increase statistics.) A model based on the beta distribution (Eq. 4.7) is shown in orange.	159
4.39	Average charge loss in front strip charge-sharing (f2b1) events as a function of α energy, corrected for calibration disagreements between individual front and back strips (see Figure 4.23). A simple line (in black) is fit to the mean of the individual pixels, with uncertainties shown (dashed).	160
4.40	Front strip charge sharing fraction of events by dead layer. Circles show the total front strip charge sharing; squares and circles show contributions from the SiO ₂ layer and the periphery p+ implant. Lines show the adopted models for these contributions, with uncertainties estimated in dashed lines. The 3.5 MeV points were taken from January data.	161
4.41	Energy difference between nominal beam energy and energy determined from the DSSD. See text for details about uncertainties.	163
4.42	Event fractions for the different dead layers and the number of front strips in which charge is deposited.	165
B.1	(a) A typical single experiment in our Monte Carlo example. (b) The t_μ for this single experiment. Note that the $\hat{\mu}$ for this particular experiment is not at the SM value of zero.	176
B.2	Distributions of t_μ for different values of μ . Note that the distributions are identical for the positive μ (symmetric about $\mu = 0$). See text for explanation. . . .	178
B.3	Expected sensitivities at two different p -values using the Asimov treatment for estimating two independent, Gaussian-distributed parameters. See text for details.	179

C.1	The demonstration linear quadrupole ion trap and associated equipment.	181
C.2	Spores (green dots) trapped in the microparticle linear quadrupole trap.	181
C.3	Circuit diagram of the demonstration ion trap.	182
C.4	A single line of spores (green dots) trapped in the microparticle linear quadrupole trap. Smearred lines show the micromotion oscillation of the spores. The nut visible is roughly 8 mm wide.	183
D.1	From left: Jason A. Clark, Guy Savard, Aaron T. Gallant, Mary T. Burkey, Louis Varriano, and Nicholas D. Scielzo.	184
D.2	From left: Louis Varriano, Bernhard Maaß, Jason A. Clark, Kay Kolos, Guy Savard, Daniel P. Burdette, Peter Müller, Adrian A. Valverde, W. Sam Porter, Aaron T. Gallant, and Zach Purcell.	185
D.3	From left: Biying Liu, Adrian A. Valverde, Jason A. Clark, Maxime Brodeur, Brenden Longfellow, Dwaipayyan Ray, and Alec Cannon.	186
D.4	From left: Tsviki Hirsh and the author, looking extremely haggard after a month of taking data.	187

List of Tables

1.1	Experiments included in Figure 1.4; results from the BPT are combined into a single line for ${}^8\text{Li}$. Fermi (F) decays are sensitive to scalar currents; Gamow-Teller (GT) decays are sensitive to tensor currents.	11
1.2	Uncertainty budget from measurement of Ref. [38] and the expected uncertainty budget from the ${}^8\text{Li}$ data campaign performed in this work. ***Key recoil-order terms will be attempted to be measured experimentally from the data by performing a “global” fit to the data, which can reduce or eliminate this systematic uncertainty.	27
1.3	Some recent experiments studying the ${}^8\text{B}$ 2α spectrum with silicon detectors. . .	29
2.1	A brief history of changes to the BPT; the last row is the BPT Mk IV. The version with a trap electrode radius to center $r = 11$ mm was used primarily to study isotopes that undergo β -delayed neutron decay. Pictures taken from Refs. [78, 80, 81, 87]	35
3.1	Voltage settings for the BPT Mk IV during the commissioning ${}^8\text{Li}$ run. Electrodes are labeled as in Figure 2.4b. Electrode “A” is pulsed during capture; electrodes “D” and “E” are pulsed during ejection. The resulting axial DC potential is shown in Figure 3.4.	79
4.1	Relevant α energies and intensities from ${}^{148}\text{Gd}$ [144] and ${}^{244}\text{Cm}$ [143].	98
4.2	Contributions to the energy loss distribution and detector response function. Adapted from Ref. [145].	100
4.3	Dimensions of key surface features of the BPT DSSDs. Profilometer and optical interferometry measurements performed by Tsviki Hirsh and Tal Zaharoni (Soreq NRC, Israel). Details of the α beam characterization of the double-sided silicon strip detector (DSSD) are presented in section 4.5 and section 4.6.	119

4.4	Nominal beam energies and energy determined by DSSD for the March run of ATLAS experiment #1905.	136
4.5	Events are classified by the number of neighboring strips with deposited energy, as shown.	136

LIST OF ACRONYMS

- ADC** analog-to-digital converter
- ANL** Argonne National Laboratory
- ATLAS** the Argonne Tandem Linear Accelerator System
- BPT** Beta-decay Paul Trap
- BSM** beyond-Standard Model
- CAD** computer-aided design
- CL** confidence level
- DAQ** data acquisition system
- DSSD** double-sided silicon strip detector
- FWHM** full width at half maximum
- IS** isotope separator
- ISOLDE** the Isotope mass Separator On-Line facility
- MC** Monte Carlo
- MCP** micro-channel plate
- MUSIC** MUlti-Sampling Ionization Chamber
- NIEL** non-ionizing energy loss distribution
- PCB** printed circuit board
- PDF** probability density function
- PIPS** passivated implanted planar silicon

PTFE polytetrafluoroethylene

RFQ radio frequency quadrupole

SM Standard Model

SRIM the Stopping and Range of Ions in Matter program [\[1\]](#)

SSB silicon surface barrier

UHV ultra-high vacuum

ACKNOWLEDGMENTS

This work would not have been possible without the contributions and support of many people. In particular, I would like to thank my advisor, Guy Savard, for his mentorship and continuous support. Tremendous thanks also to Jason Clark, without whom much would be impossible.

Thanks also to the students and postdocs in our group of the past few years, including Daniel P. Burdette, Mary T. Burkey, Nathan Callahan, Biying Liu, Graeme Morgan, Dwai-payan Ray, Adrian A. Valverde, Gemma Wilson, and Xinliang Yan. Thanks to the advice and support of the collaborators on the BPT project, particularly Aaron T. Gallant, Tsviki Hirsh, Brenden Longfellow, Nicholas D. Scielzo, and Ralph E. Segel. Thanks especially to all those who took shifts during the ^8Li data campaign, who are listed in Appendix D.

Thanks also to the postdocs and staff scientists of the Physics Division of Argonne National Laboratory (several of whom have moved to different institutions), particularly Matt Gott, Heshani Jayatissa, Ben Kay, Bernhard Maaß, Peter Müeller, Claus Müller-Gattermann, Daniel Santiago-Gonzalez, Ryan Tang, and Ivan Tolstukhin. The support staff of the Physics Division also enabled much of the work performed here, including Edward J. Boron, III, Michael Oberling, and John Rohrer. Thanks also to the ATLAS operators and support staff who enabled the experiments described here to run smoothly.

Thank you to the committee members for their time and advice over the years, which included Edward Blucher, Marcela Carena, David DeMille, and Abigail Viereg. (John P. Schiffer served for several years on the committee before his untimely death.) Thanks also to my undergraduate research advisor, Yuri Kamyshev, who initially mentored me and inspired much of my interest in fundamental symmetries, as well as the many other mentors I had during my undergraduate career.

Finally, thanks to my family and friends for their continued support over the years.

ABSTRACT

The Beta-decay Paul Trap (BPT) measures the beta-neutrino angular correlation coefficient $a_{\beta\nu}$ in the Gamow-Teller decays of ${}^8\text{Li}$ and ${}^8\text{B}$ (decaying to ${}^8\text{Be}^* \rightarrow 2\alpha$) to search for an electroweak tensor coupling, a beyond-Standard Model possibility. The new BPT Mk IV trap reduces β -scattering by a factor of 4, a key source of systematic uncertainty, and has recently been commissioned with ${}^8\text{Li}$ with 2.7 million triple coincidence events. In addition, a complete detector characterization has been performed with an α beam to reduce systematic uncertainties associated with the silicon strip detector response. This new high-statistics ${}^8\text{Li}$ data set, coupled with the experimental advancements, will allow for improved sensitivity to a tensor contribution once data analysis is completed and unblinded. In the near future, the BPT should be able to reach an uncertainty of $\Delta|C_T/C_A|^2 \leq 10^{-3}$, though the theoretical understanding of the nuclear structure still requires improvement.

Chapter 1

INTRODUCTION

1.1 The Standard Model and weak interaction

As with any experiment looking for “new physics,” we begin with a (very) brief description of the Standard Model (SM), whose particle content is shown Figure 1.1. Skipping over the decades of theoretical and experimental work to create and verify the SM, the SM is a chiral gauge theory with $SU(3)_C \times SU(2)_L \times U(1)_Y$ symmetry, based on the strong interaction ($SU(3)_C$), weak isospin ($SU(2)_L$), and weak hypercharge ($U(1)_Y$), as well as the Higgs boson. From only a few free parameters, the SM predicts particle and interaction properties with an astonishing degree of accuracy [2].




Figure 1.1: The Standard Model. From Wikipedia.

However, we know from several different observations that the SM is incomplete, including: the fact that gravity is not included in the theory, that neutrinos are observed to have mass, that matter is dominant over antimatter, and that dark matter and dark energy make

up the vast majority of the observable universe. These questions, as well as others, open the door to tests of the SM that interrogate it to a high degree of precision to search for new physics that can improve the theory.

Low-energy tests of the SM are uniquely situated to provide broadband tests of new physics at energy scales well above their typical kinetic energies [3]. By measuring the β -decay spectrum (\sim MeV scale) very precisely, deviations from the SM can indicate the presence of additional electroweak interactions and new particles at above the TeV-scale, in a broadband, model-independent way.

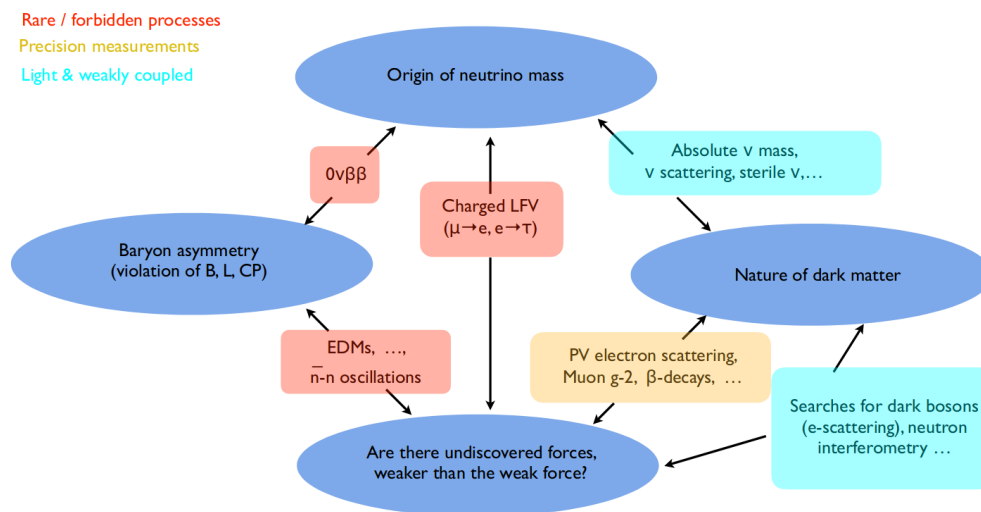


Figure 1.2: Four interconnected questions about fundamental interactions and the observed universe that can be explored with nuclear physics. From Ref. [3].

In the SM, the weak interaction is a chiral interaction that maximally violates parity conservation, as posited by Lee and Yang [4] and confirmed by Wu et al. [5]. The long history of the development of the weak interaction is available in e.g. Ref. [6]. Assuming no neutrino mass, the nucleon-level effective field theory Lagrangian describing weak interactions to leading order is given in general by [7]

$$\begin{aligned}
\mathcal{L} = & -\bar{p}\gamma^\mu n(C_V^+\bar{e}\gamma_\mu\nu_L + C_V^-\bar{e}\gamma_\mu\nu_R) + \bar{p}\gamma^\mu\gamma_5 n(C_A^+\bar{e}\gamma_\mu\nu_L + C_A^-\bar{e}\gamma_\mu\nu_R) \\
& -\bar{p}n(C_S^+\bar{e}\nu_L + C_S^-\bar{e}\nu_R) - \frac{1}{2}\bar{p}\sigma^{\mu\nu}n(C_T^+\bar{e}\sigma_{\mu\nu}\nu_L + C_T^-\bar{e}\sigma_{\mu\nu}\nu_R) \\
& -\bar{p}\gamma_5 n(C_P^+\bar{e}\nu_L - C_P^-\bar{e}\nu_R) + \text{h.c.},
\end{aligned} \tag{1.1}$$

where the C_X^\pm parameterize the strengths of the effective interactions between the nucleons and leptons at low energies. The superscripts refer to the neutrino chirality where $+$ ($-$) is left-handed (right-handed). The subscripts V, A, S, T , and P (vector, axial vector, scalar, tensor, and pseudoscalar) are the 5 possible Lorentz-invariant form factors that can contribute to the weak interaction. Although, in principle, any of these 5 terms could be present, only 2 have been found to exist in reality, leading to the vector minus axial-vector ($V - A$) structure of the weak interaction in the SM.

Historically, equation (1.1) has usually been written with the Wilson coefficients $C_X^{(\prime)}$ [8] where¹

$$C_X = \frac{C_X^+ + C_X^-}{2}, \quad C'_X = \frac{C_X^+ - C_X^-}{2}. \tag{1.2}$$

It is physically more transparent to instead write the coefficients in a way where the neutrino chirality is apparent. This form makes it easier to connect the effective nucleon-level interactions to the chiral effective field theory effective coupling strengths that describe the quark-level interactions [7, 9]. Additionally, the C_X^\pm form reduces correlations for experiment interpretations.

An imaginary part of the C_X^\pm would imply time-reversal violation, or equivalently, simultaneous charge-conjugation and parity violation; searches for time-reversal violation in β -decay require polarized nuclei [8]. New efforts at measuring this are underway (e.g. [10]).

¹There have been several different conventions for writing these equations with different numerical factors depending on the specific convention; this makes interpretation across experiments slightly tricky. This work tries to use C^\pm notation but may report some results using $C^{(\prime)}$ coefficients.

Searches for time-reversal violation can also be performed using electric dipole moments [3]. New sources of charge-parity violation are required to help explain the abundance of matter over antimatter [3]. This work does not have access to any imaginary portions of the C_X^\pm , and so they are taken to be real here. Additionally, contributions from the pseudo-scalar term are highly suppressed, and they are also neglected [7].

1.2 Nuclear β -decay observables

β -decay can be a sensitive probe of the weak interaction, on par with sensitivities achieved in high-energy experiments and from other bounds [9, 11–13]. The theory of β -decay developed over the course of the 20th century [6]; the entire history will not be recited here. Transitions in β -decay are dependent on the Fermi (F) and Gamow-Teller (GT) nuclear matrix elements. In a Fermi decay, the outgoing leptons (β, ν) have spins that are anti-aligned, requiring that the spin of the baryons are the same before and after the decay (angular momentum $\Delta J = 0$ and isospin $\Delta T = 0$). In a Gamow-Teller decay, the leptons have aligned spins, and the spin of the baryon can flip ($\Delta J = 0, \pm 1, \Delta T = 0, \pm 1$, but not $J = 0 \rightarrow 0$ or $T = 0 \rightarrow 0$). Mixed transitions contain both F and GT contributions. Transitions where the leptons carry away no angular momentum ($L = 0$) are known as “allowed” transitions. Relating these to equation (1.1), Fermi transitions are sensitive only to vector (C_V^\pm) and scalar (C_S^\pm) couplings, while Gamow-Teller transitions are sensitive to axial-vector (C_A^\pm) and tensor (C_T^\pm) couplings.

Assuming time-reversal conservation (real C_X^\pm) and summing over the polarizations of the child nucleus and β , the differential decay width for an allowed transition ($J \rightarrow J' = J, J \pm 1$) to leading order is given by [7, 8, 12, 14]

$$\frac{d\Gamma}{dE_e d\Omega_e d\Omega_\nu} = \hat{\xi} \frac{M_F^2}{64\pi^5} F(\Delta - E_e)^2 p_e E_e \left\{ 1 + b \frac{m_e}{E_e} + a \frac{\vec{p}_e \cdot \vec{p}_\nu}{E_e E_\nu} + A \frac{\vec{p}_e \cdot \langle \vec{J} \rangle}{E_e J} + B \frac{\vec{p}_\nu \cdot \langle \vec{J} \rangle}{E_\nu J} + c \left[\frac{\vec{p}_e \cdot \vec{p}_\nu}{3E_e E_\nu} - \frac{(\vec{p}_e \cdot \vec{j})(\vec{p}_\nu \cdot \vec{j})}{E_e E_\nu} \right] \left[\frac{J(J+1) - 3(\langle \vec{J} \rangle \cdot \vec{j})^2}{J(2J-1)} \right] \right\}, \quad (1.3)$$

where J refers to the parent nucleus spin. Experiments typically measure the various correlation coefficient parameters (b, a, A, B, c), where the sensitivity to particular terms depends on the experiment details (including which particles are detected and whether polarization is measured). Typically, experiments seeking to place limits on beyond-Standard Model (BSM) terms have relied upon pure F or GT transitions; mixed decays require an understanding of the relative ratio between the two types of transitions. A notable exception is the mixed decay of the neutron, which is a key system to study due to its relative simplicity. Additionally, there is recent interest in exploring forbidden transitions [15]. Full expressions for these correlation coefficients may be found in e.g. Ref. [7, 12], but we focus on two coefficients that are relevant to this work, a and b . For these correlation coefficients, we have

$$\begin{aligned} \hat{\xi} &\equiv (C_V^+)^2 + (C_S^+)^2 + (C_V^-)^2 + (C_S^-)^2 + \frac{(C_V^+)^2}{(C_A^+)^2} \tilde{\rho}^2 [(C_A^+)^2 + (C_T^+)^2 + (C_A^-)^2 + (C_T^-)^2] \\ b\hat{\xi} &\equiv \pm 2 \left\{ C_V^+ C_S^+ + C_V^- C_S^- + \frac{(C_V^+)^2}{(C_A^+)^2} \tilde{\rho}^2 [C_A^+ C_T^+ + C_A^- C_T^-] \right\} \\ a\hat{\xi} &\equiv (C_V^+)^2 - (C_S^+)^2 + (C_V^-)^2 - (C_S^-)^2 - \frac{1}{3} \frac{(C_V^+)^2}{(C_A^+)^2} \tilde{\rho}^2 [(C_A^+)^2 - (C_T^+)^2 + (C_A^-)^2 - (C_T^-)^2]. \end{aligned} \quad (1.4)$$

The b coefficient is known as the Fierz interference term as it clearly contains interference between the different C_X^\pm ; the $+(-)$ in equation (1.4) is from electron (positron) decay. The a coefficient is the β - ν angular correlation coefficient, and it is sensitive to angular correlations between these particles, as seen from equation (1.3). The $\tilde{\rho}$ is the mixing ratio for mixed decays. It can be clearly seen that, for the SM, $b = 0$ and $a = -1/3$ for GT decays, such as

in this work. A deviation from these values (aside from SM corrections) would indicate new physics.

Experiments can measure several different observables or their combination, such as the β energy spectrum or the recoiling nucleus energy spectrum which are discussed in the next section. With the advent of ion and atom trapping techniques, modern experiments are often able to measure both spectra, as well as angular correlations directly. Historically, experiments have tended to measure a as an asymmetry integrated over the energy spectrum of the decay, assuming that $b = 0$ (or, equivalently, that $C_T^+ = 0$ with SM $C_A^- = 0$). This then gives sensitivity only to C_T^- through the a term. This results in a measured correlation coefficient different than the true correlation coefficient as

$$\tilde{a} \equiv \frac{a}{1 + b\langle m_e/E_e \rangle}, \quad (1.5)$$

where $\langle m_e/E_e \rangle$ is an average over the energies that contribute in the experiment. If the restriction that $b = 0$ is lifted, equation (1.5) can then be used to reinterpret the measured \tilde{a} to provide a limit on C_T^+ as well. Note that this treatment is not wholly correct for measurements of differential spectra [16, 17]. Additionally, this treatment ignores the correlations between a and b , which are explored further in section 1.3.2.

Combining data from different isotopes and measurements of different correlation coefficients leads to better sensitivities. Using data from superallowed β -decay measurements, neutron decay measurements, and mirror β -decay measurements, a global limit on the BSM coupling coefficients has recently been placed [7]². Interestingly, the data from β -decay seem to favor a non-zero C_T^- , though this is pulled by the result of Ref. [18]. This is shown in Figure 1.3. The values of the C_X^\pm of the global fit [7] are

²Note that this global fit [7] does not include any of the results of the BPT measurements with ^8Li and ^8B .

$$v^2 \begin{pmatrix} C_V^+ \\ C_A^+ \\ C_S^+ \\ C_T^+ \end{pmatrix} = \begin{pmatrix} 0.98501_{(-114)}^{(+75)} \\ -1.2544_{(-11)}^{(+14)} \\ -0.0007_{(-14)}^{(+29)} \\ -0.0010_{(-22)}^{(+33)} \end{pmatrix}, \quad \begin{pmatrix} v^2 |C_V^-| < 0.053 \\ v^2 |C_A^-| < 0.063 \\ v^2 |C_S^-| < 0.050 \\ v^2 |C_T^-| \in [0.072, 0.099] \end{pmatrix}. \quad (1.6)$$

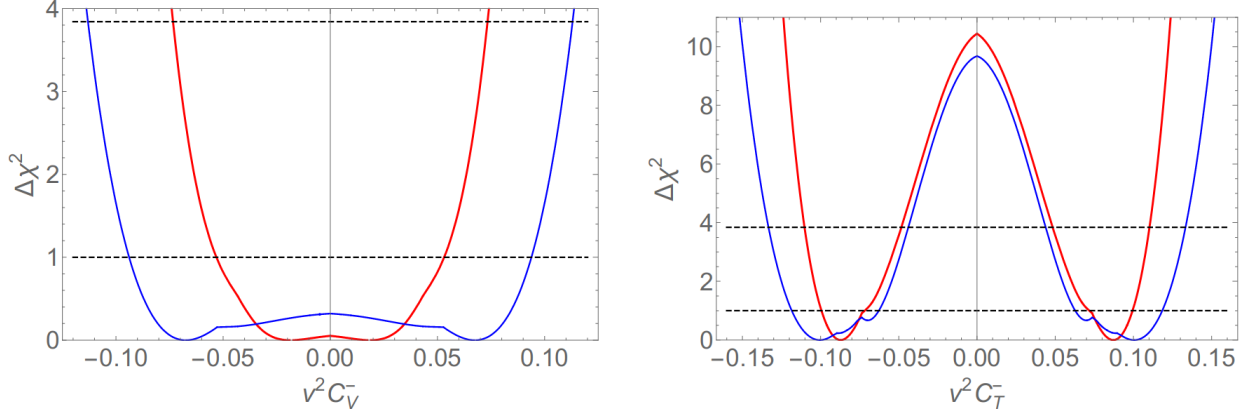


Figure 1.3: From Ref. [7]: Marginalized $\Delta\chi^2 \equiv \chi^2 - \chi_{\min}^2$ distributions for the Wilson coefficients C_V^- (left) and C_T^- (right), with (red) and without (blue) taking account the input from mirror beta decay.

Recently, it has been suggested to measure observables in unique first-forbidden β decay transitions ($\Delta J = 2$ with a change in parity, sensitive to A and T couplings) as a complement to the usual measurements of allowed transitions [15]. Specifically, it is found that the β energy spectrum for first-forbidden transitions is sensitive to C_T^- (right-handed neutrino tensor couplings). This is in contrast to the β energy spectrum for allowed transitions, which are insensitive to C_T^- and only sensitive to C_T^+ . Note that recoil energy spectrum in allowed transitions, however, is sensitive to both couplings, though some care must be taken in extracting the limits [17] (discussed in section 1.3.2).

There are several essential corrections to low-energy observables of the weak Lagrangian; these stem largely from electromagnetic corrections due to the fact that the β -decay occurs inside a nucleus. Corrections to this Lagrangian from electromagnetic contributions have been noted for decades [19–21]. Electromagnetic corrections are typically divided into “inner”

and “outer” radiative corrections [22, 23]. The outer radiative corrections are due to the fact that real and virtual photons can contribute to altering the decay spectrum; these corrections depend on the nucleus under study and its structure. Due to its impact in determining V_{ud} , the u, d quark mixing matrix element of the CKM matrix, the inner radiative correction has been a topic of considerable recent interest by a number of authors (e.g. Ref. [24]). At the level of precision of modern experiments, recoil-order corrections also contribute to the β -decay spectrum shape [21]. These corrections are essential for the study of ${}^8\text{Li}$ and ${}^8\text{B}$, the subject of this thesis.

1.3 Measurements of the $\beta - \nu$ angular correlation

As noted in the previous section, measurements of the $\beta - \nu$ angular correlation in β -decay can provide limits on S and T couplings present in the SM [13]. To consider this further, we rewrite equation (1.3) after averaging over the spin variables of the nucleus and electron [25]. The decay rate is then proportional to

$$N(E_e, \theta_{\beta\nu})dE_e d\Omega_\theta \propto p_e E_e p_\nu^2 \left[1 + b \frac{m_e}{E_e} + a \frac{p_e}{E_e} \cos(\theta_{\beta\nu}) \right] dE_e d\Omega_\theta. \quad (1.7)$$

Note that E_e and $\theta_{\beta\nu}$ are independent variables in this equation, with the assumption that the neutrino is massless.

It is not strictly necessary to measure the angle between the β and the ν (which cannot be detected) to measure the angular correlation. Instead, the recoil momentum (or, usually, the recoil energy) spectrum is measured, which can be seen through a change of variables in the above equation [17, 25]. Neglecting the energy of the recoiling nucleus, this transformation leads to

$$N(E_e, p_r)dE_e dp_r \propto \frac{1}{2} p_r E_e p_\nu \left[1 + b \frac{m_e}{E_e} + a \frac{p_r^2 - p_e^2 - p_\nu^2}{2E_e p_\nu} \right] dE_e dp_r, \quad (1.8)$$

where p_r is the momentum of the recoiling nucleus. Note that p_r is dependent upon E_e with

the restriction that $|p_e - p_\nu| \leq p_r \leq |p_e + p_\nu|$ for a fixed value of E_e . This effectively means that the recoil spectrum is, to some degree, sensitive to both E_e and $\theta_{\beta\nu}$.

Considering only GT decays, which are sensitive to A and T couplings, and taking $C_A^- \equiv 0$ (no right-handed neutrino A couplings)³, we have from equation (1.4),

$$a_{\text{GT}} = -\frac{1}{3} \frac{(C_A^+)^2 - (C_T^+)^2 - (C_T^-)^2}{(C_A^+)^2 + (C_T^+)^2 + (C_T^-)^2} = -\frac{1}{3} \frac{1 - (C_T^+/C_A^+)^2 - (C_T^-/C_A^+)^2}{1 + (C_T^+/C_A^+)^2 + (C_T^-/C_A^+)^2} \quad (1.9)$$

and

$$b_{\text{GT}} = \pm 2 \frac{C_A^+ C_T^+}{(C_A^+)^2 + (C_T^+)^2 + (C_T^-)^2} = \pm 2 \frac{C_T^+/C_A^+}{1 + (C_T^+/C_A^+)^2 + (C_T^-/C_A^+)^2}, \quad (1.10)$$

where it is more experimentally useful to consider the relative strengths of the couplings, C_T^+/C_A^+ and C_T^-/C_A^+ . Similar expressions are found for allowed F decays, with S and V couplings and a factor of $+1$ instead of $-1/3$ in a .

Note that a measurement of b is sensitive only to C_T^+ and not C_T^- from the numerator, while a measurement of a can constrain both couplings. However, note also that a depends quadratically on the BSM couplings, while b depends linearly. This suggests that a measurement of b will be more sensitive to small couplings, a feature which is discussed in more detail in the next section. Note also that the sign of b depends on whether the decay is β^- or β^+ ; the sign is taken as $+(-)$ for electron (positron) decay. This means that measurements of both types of decays can be combined for a joint limit that can be much more sensitive; this has recently been performed with ^8Li and ^8B by the BPT [26].

The $\beta - \nu$ angular correlation coefficient a has been measured in many isotopes over the years. A selection of (mostly) recent results measuring the angular correlation coefficient is presented in Figure 1.4; these agree with the SM for the most part. Many of the latest and most accurate measurements in nuclei come from ion or atom traps, as indicated in Table

³For correlation measurements such as that of the BPT, the only sensitive to C_A^- comes from the Fierz interference term, as the correlation a cannot distinguish between C_A^- and C_A^+ . The global fit [7] results shown in equation (1.6) show that the limits on C_A^- and C_T^- are sufficiently small that their combination $C_A^- C_T^-$ in b is vanishingly small compared to the $C_A^+ C_T^+$ term.

1.1. With traps, it is possible to detect the energies of the β and the recoil (and/or delayed particle emission) as well as their positions; this allows for a full kinematic reconstruction to determine the ν energy and momentum.

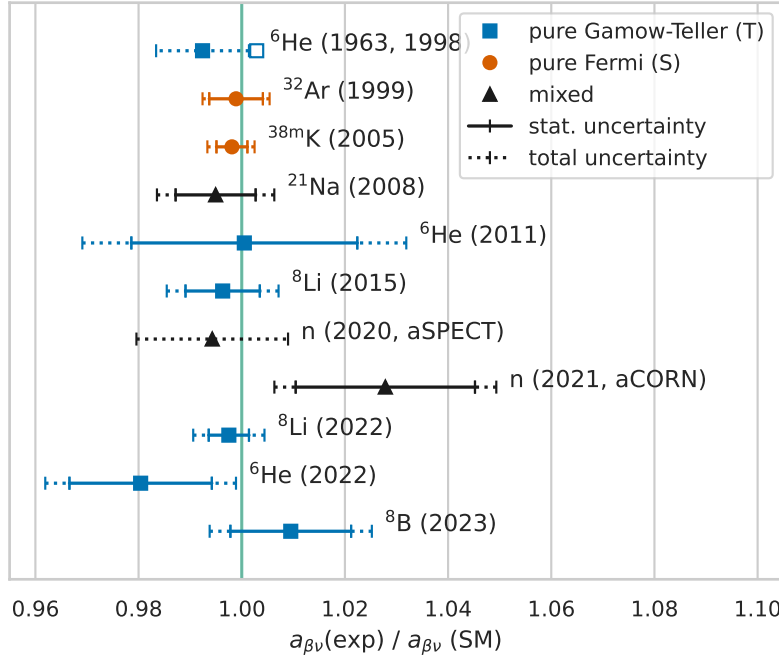


Figure 1.4: Select experiments measuring $a_{\beta\nu}$ in nuclei and neutrons.

1.3.1 Enhancement to sensitivity from β -delayed particle decays

It was noticed by Holstein in Ref. [21] that β -delayed particle decays have an enhancement to their sensitivity to recoil-order contributions. This comes from the fact that the child nucleus is oriented after the initial β -decay, and the subsequent emission of a particle is sensitive to the recoil of the child nucleus. In particular, delayed particles emitted in the direction of the recoil momentum will have their energies Doppler shifted in what has been called a kinematic shift. For events in which the delayed particle is emitted in the opposite direction as the β , the kinetic shift is maximum. The size of the enhancement depends on the spin sequence in the decay with a derivation demonstrating this enhanced sensitivity provided in Refs. [40, 41].

Isotope (F/GT)	Notes	Year(s)	Relevant references
${}^6\text{He}$ (GT)	recoil energy	1963, 1998	[27, 28]
${}^{32}\text{Ar}$ (F)	delayed p energy	1999	[29]
${}^{38\text{m}}\text{K}$ (F)	full kinematics, atom trap	2005	[30]
${}^{21}\text{Na}$ (mixed)	recoil energy, atom trap	2008	[31, 32]
${}^6\text{He}$ (GT)	full kinematics, ion trap	2011	[33]
n , aSPECT (mixed)	recoil energy	2020	[18]
n , aCORN (mixed)	β and recoil energies	2021	[34, 35]
${}^8\text{Li}$ (GT)	full kinematics, ion trap (BPT)	2013, 2015, 2022	[36–38]
${}^6\text{He}$ (GT)	full kinematics, atom trap	2022	[39]
${}^8\text{B}$ (GT)	full kinematics, ion trap (BPT)	2023	[26]

Table 1.1: Experiments included in Figure 1.4; results from the BPT are combined into a single line for ${}^8\text{Li}$. Fermi (F) decays are sensitive to scalar currents; Gamow-Teller (GT) decays are sensitive to tensor currents.

Ref. [40] measured the kinematic shift in ^{20}Na β -delayed α decay. ^{20}Na decays to a number of levels in ^{20}Ne before subsequently emitting an α ; there are pure GT, pure F, and mixed transitions. The spin sequence in this decay is $2^+ \rightarrow 2^+ \rightarrow 0^+$. By measuring the energy shift of the α when the α and β are anti-parallel, a factor of 3 enhancement in sensitivity is gained. The results of Ref. [40] for the most populated GT level is shown in Figure 1.5. The measurement was used to classify the different transitions and to measure the mixing ratio; the results are not sensitive enough to place limits on exotic currents today.




Figure 1.5: From Ref. [40]: Coincidence data for the most populated ^{20}Na β -delayed α level showing the energy shift as a function of the measured β -particle energy with the angle between the α and β indicated. The curves show the expected kinematic shift for a pure Fermi (F) and a pure Gamow-Teller (G) transition.

It is not strictly necessary to measure the β energy to perform a kinematic shift measurement, provided that the angle between the α and the β is known. If the β energies are

integrated over, there will be still be a non-zero kinematic shift, as can be seen from 1.5.

The kinematic shift technique was also successfully employed in a delayed p measurement from ^{32}Ar decay in Ref. [29], which placed strict limits on scalar currents. Using a delayed γ , a Doppler shift from the β -decay was used also to measure the angular correlation in ^{18}Ne and ^{14}O , but these have rather low precision of only a few percent [42, 43]. By detecting the delayed α from the decay of ^8Li , limits on tensor currents have been placed by the Beta-decay Paul Trap (BPT) [36–38]. Future experiments are planned with other delayed proton emitters [44, 45] that are sensitive to both scalar and tensor currents.

1.3.2 Sensitivity to the Fierz interference term

There has been considerable discussion in recent years about measuring the Fierz interference term b in addition to the angular correlation a [12, 17]. In allowed transitions, both the β energy spectrum and the recoiling energy spectrum (or angular correlation) are sensitive to the Fierz interference term, though this may not be immediately obvious in the latter case. To see this qualitatively, note that the recoil energy (from equation (1.8)) is dependent upon both E_e and $\theta_{\beta\nu}$. Additionally, if only the recoil energy spectrum is measured, this integrates out the β energy and leads to a normalization factor that includes b . This results in the expression given in equation (1.5), which can be seen to include a dependence on b . Historically, measurements of the recoil energy spectrum have taken $b \equiv 0$ initially (equivalent to an assumption of no left-handed neutrino couplings), and fit the measured spectrum with a (now called \tilde{a}) as a free parameter in equation (1.8). Then, by the prescription of equation (1.5), the $b = 0$ assumption can be lifted to report uncertainties on both a and b (or on the BSM couplings directly). This is effectively the treatment that the BPT has used in its recent measurements [26, 38] to report joint limits on C_T^+ and C_T^- .

The \tilde{a} prescription is problematic when applied to the recoil energy spectrum, as pointed out in Ref. [17]. This is due to the fact that the recoil energy depends on the β energy, and therefore the b term cannot be cleanly integrated. This has led to future experiments

that are poised to be sensitive to b (and C_T^+) only, such as a measurement of the β energy spectrum in ${}^6\text{He}$ using the cyclotron radiation emission spectroscopy technique [46].

Though Ref. [17] suggests fitting both a and b simultaneously in the analysis of an experiment, this treatment leaves out the correlation between these two parameters, namely that they contain different combinations of the same BSM coupling strengths (equations (1.9) and (1.10)), which are the real parameters of interest. Since it is really the BSM coupling constants that these experiments are interested in, it seems of limited utility to discuss measuring a and b , though it can be easier to fit these constants in the analysis of an experiment. For the level of precision that experiments are aiming towards these days, it is likely more instructive to more fully consider the correlations between these parameters. In addition, all modern experiments rely on Monte Carlo simulations of decays, including recoil order terms, to evaluate their system response and the effect of various analysis cuts. It is not clear how more simple treatments of a and b are affected when these cuts and other possible experimental correlations are taken into account.

1.4 Correlations between a and b

We now more properly evaluate the sensitivity and biases of the various methods to analyze the results of a β energy or recoil energy spectrum measurement. In particular, we want to look more closely at the correlation between the angular correlation coefficient a and the Fierz interference term b and their effects on the various spectra that can be measured. For now, we consider only allowed GT decays, whose spectra are given in equations (1.7) and (1.8). Correlation coefficients a and b are given in equations (1.9) and (1.10), respectively. The \tilde{a} prescription is given in equation (1.5). Compared to the \tilde{a} prescription of equation (1.5), which treats a and b independently, this method keeps the correlations between the parameters of interest.⁴ Note that we do not yet consider the case of a delayed particle emission; these have a larger phase space (4D instead of 2D) and are more complicated.

⁴At large numbers of events ($N \gtrsim 10$ million events, depending on Q), these correlations are less important due to the dominance of the b term relative to the a term.

These will be analyzed in the future.

First, we demonstrate that the $\theta_{\beta\nu}$ and p_r spectra are more sensitive to b than to a as $C_T^+ \rightarrow 0$ (assuming $C_T^- \equiv 0$). This is expected behavior since b depends linearly on C_T^+ while a has a quadratic dependence. This can be seen explicitly by comparing the spectra (integrated over the β energy E_e) at various values of C_T^+/C_A^+ , as shown in Figure 1.6 for a β^- -decay with a Q value of 13 MeV (the ${}^8\text{Li}$ β -decay Q value). The E_e spectrum is generally more sensitive than the other spectra to C_T^+ (though insensitive to C_T^-), but there are many experimental difficulties in an accurate measurement of the β energy spectrum. Interestingly, note that there is a near-exact cancellation for a particular value of C_T^+/C_A^+ in the $\theta_{\beta\nu}$ and p_r spectra; in the 2D parameter space $(C_T^+/C_A^+, C_T^-/C_A^+)$, the combinations of parameters creating a near-exact cancellation forms a circle.

The sensitivity for different Q values and the locations of these approximate cancellations in the p_r spectrum can be seen in Figure 1.7. At each value of C_T^+/C_A^+ (assuming $C_T^- \equiv 0$), the absolute value of the relative difference between the SM and BSM p_r spectra is found, then averaged across p_r . The average relative difference is highest for $Q \sim 2$ MeV, leading to the highest sensitivity. Points of near cancellation are observed, which occur when the a and b terms approximately cancel; these are dependent on the Q value.

To emphasize, as $C_T^+ \rightarrow 0$, the $\theta_{\beta\nu}$ and p_r spectra become more sensitive to b , which is linearly sensitive to BSM contributions. We therefore have the extremely interesting situation in which, as the number of events N increases in an experiment measuring $\theta_{\beta\nu}$ or p_r , *the expected sensitivity changes from being proportional to $1/\sqrt[4]{N}$ to $1/\sqrt{N}$* . Due to the dependence of p_r on E_e , the recoil spectrum has better than $1/\sqrt[4]{N}$ scaling in its sensitivity to begin with, but this becomes more like $1/\sqrt{N}$ scaling as $C_T^+ \rightarrow 0$.

For the case of $C_T^+ \equiv 0, C_T^- \neq 0$, the b term is absent. This means that p_r and $\theta_{\beta\nu}$ will both have $1/\sqrt[4]{N}$ scaling in their sensitivities, given that E_e is not sensitive to C_T^- .

Keeping this in mind, we now explore experimental sensitivities in the 2D parameter space for the various spectra at different Q values and for different numbers of events. We

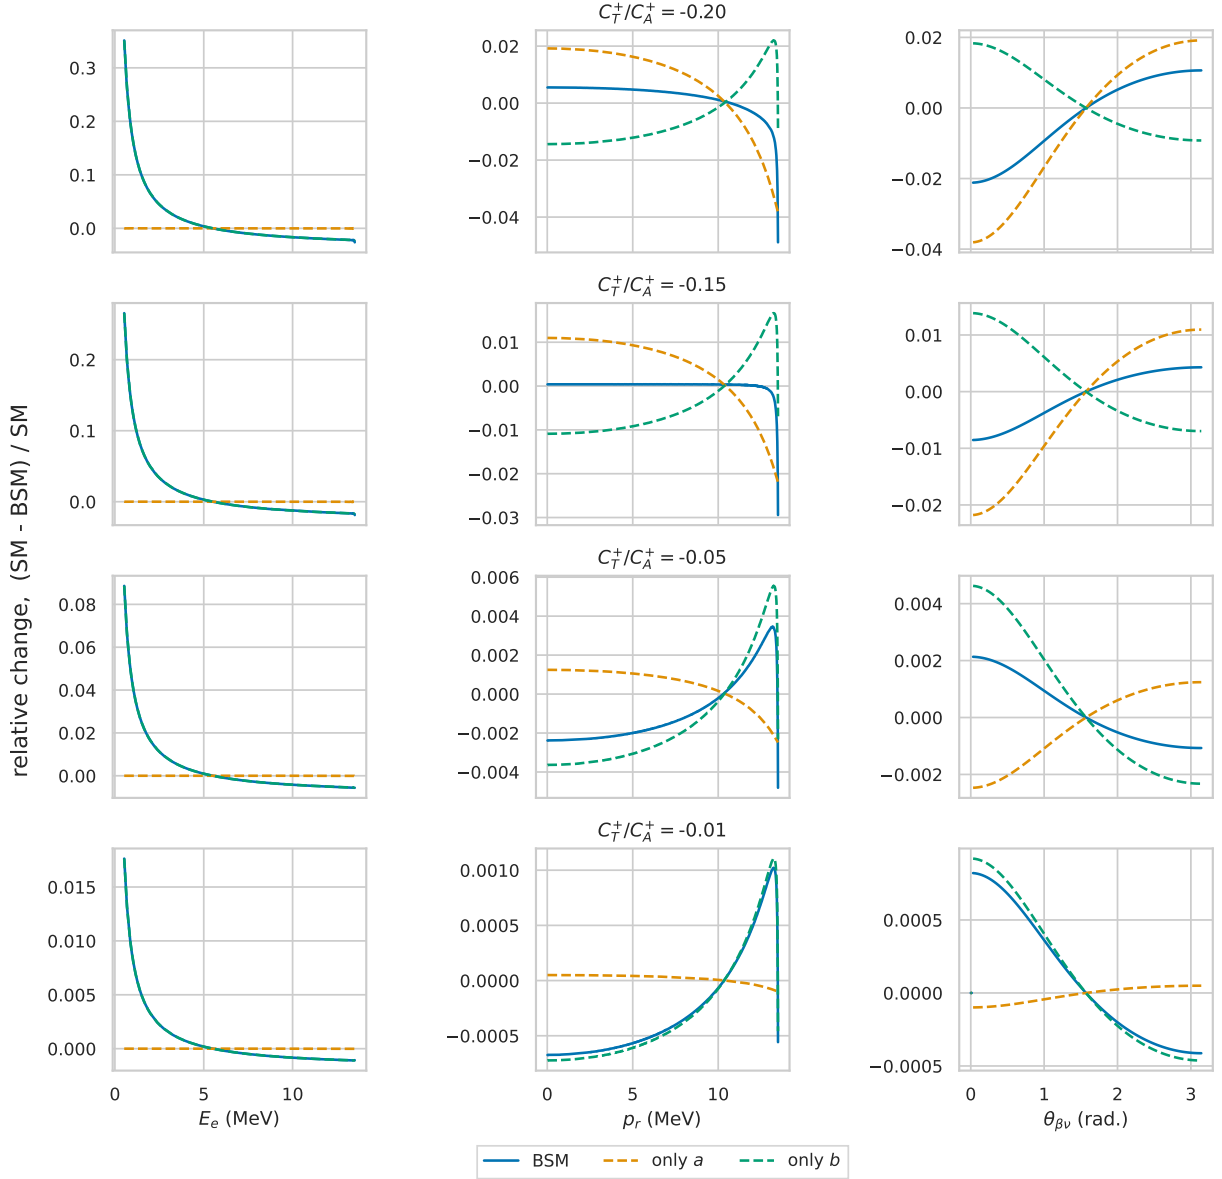


Figure 1.6: Comparison of E_e , $\theta_{\beta\nu}$, and p_r spectra for various values of C_T^+/C_A^+ assuming $C_T^- \equiv 0$; these are for β^- -decay with $Q = 13$ MeV. Vertical axis is relative difference from the SM ($C_T^+ = 0$), shown in blue, with the a contribution in orange and the b contribution in green. Note that, in the $\theta_{\beta\nu}$ and p_r spectra, as $C_T^+ \rightarrow 0$, the b contributes more compared to the a term. There are near-exact cancellations between the a and b terms for these two spectra.




Figure 1.7: At each value of C_T^+ / C_A^+ (assuming $C_T^- \equiv 0$), the absolute value of the relative difference between the SM and BSM p_r spectra is found, then averaged across p_r . The average relative difference is highest for $Q \sim 2$ MeV, leading to the highest sensitivity. Points of near cancellation are observed, which occur when the a and b terms approximately cancel; these are dependent on the Q value.

examine only β^- decays for simplicity; results will be the same for β^+ decays but mirrored about the $C_T^- = 0$ line. We compute the median expected statistical sensitivities at the 68% and 95% confidence level (CL) using the asymptotic limit from the Asimov data set [47]. This method is well-known in high-energy physics but is not as well known in low-energy nuclear physics, to my knowledge.⁵ Using this method eliminates the need for Monte Carlo simulations; these would be extremely time-consuming to perform across the large parameter space, at different Q values, and for large numbers of events. Details about this technique, including a simple example of its application, are provided in Appendix B. At each point in $(C_T^+/C_A^+, C_T^-/C_A^+)$ space on some grid, we use the equations mentioned above to compute the p -value of the Asimov data set compared to the SM data set null hypothesis; this provides the median expected sensitivities with all of the correlations between a and b included. The whole decay spectrum is used, though in an actual experiment, it is likely that the very lowest and very highest energies may not be detected or kept in the analysis.

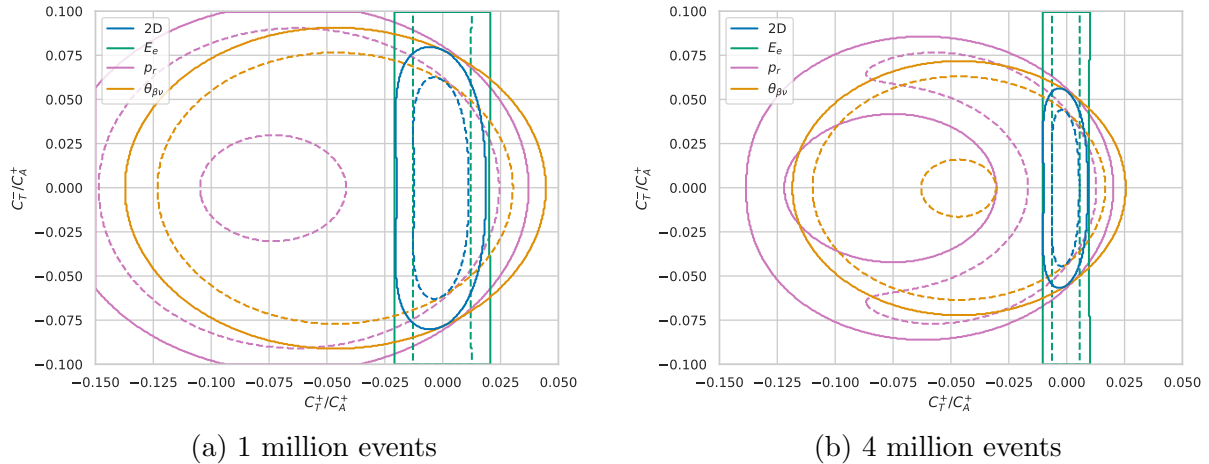


Figure 1.8: Median expected sensitivities at 68% CL (dashed) and 95% CL (solid) for $Q = 13$ MeV with 1 million and 4 million events. The colors refer to measurements of the different spectra: E_e (green), p_r (pink), $\theta_{\beta\nu}$ (orange). The blue curve is a measurement of the joint spectrum, either $(E_e, \theta_{\beta\nu})$ or (E_e, p_r) , which give identical limits as expected. See text for additional details.

Figure 1.8 shows the median expected sensitivities for 1 million and 4 million events at $Q = 13$ MeV. The colors refer to measurements of the different spectra: E_e (green), p_r

⁵Thank you to Michael D. Hank for pointing me to this method.

(pink), $\theta_{\beta\nu}$ (orange). The blue curve is a measurement of the joint spectrum, either $(E_e, \theta_{\beta\nu})$ or (E_e, p_r) , which give identical limits as expected. There is an symmetry in the sensitivity of $\theta_{\beta\nu}$ (approximate for p_r) about $C_T^+/C_A^+ = -\frac{1}{2}\langle m_e/E_e \rangle, \approx 0.048$ for this Q value.

The E_e distribution has no sensitivity to C_T^- as noted above. Note the p_r and $\theta_{\beta\nu}$ distributions explore slightly different parts of the parameter space; this comes from the fact that p_r has dependence on E_e , which also contributes to its increased sensitivity. These plots also speak to the power of experiments that can measure two spectra simultaneously to then perform a 2D analysis. These types of experiments (which can be performed using traps) can place limits on both C_T^+ and C_T^- and use all of the information from the decays.

We turn to look at the dependence on N , as shown in Figure 1.9. Note that there is a symmetry in $\theta_{\beta\nu}$ (and part of p_r) about $C_T^+/C_A^+ \approx -0.048$, and the lower branch is not shown (in part because it can be ruled out by other experiments). The $\theta_{\beta\nu}$ distribution initially shows a $1/\sqrt[4]{N}$ dependence, which then becomes a $1/\sqrt{N}$ dependence as predicted. The p_r distribution does not show this behavior for these N , but this is because the distribution is already dominated by b . This can be seen by noting that, in Figure 1.6, the a, b near cancellation value of $C_T^+/C_A^+ \approx -0.15$, which is already largely ruled out at these N .




Figure 1.9: Median expected sensitivity at 68% CL to (a) C_T^+/C_A^+ (assuming $C_T^-/C_A^+ \equiv 0$) and (b) C_T^-/C_A^+ (assuming $C_T^+/C_A^+ \equiv 0$) as a function of the number of events for $Q = 13$ MeV.

We can also make a comparison of the sensitivities of different isotopes for these various spectra. This is performed with 100 million events, the level of statistics that the next generation of β -decay experiments is aiming towards. This is shown in Figure 1.10, where the absolute value of the parameter is shown since the limits are the same above and below zero. As can be seen, the sensitivity for C_T^+/C_A^+ is best at lower Q values (~ 1 -4 MeV) with measurements of E_e . Measurements of the 2D spectra do only slightly better but are able to place limits on C_T^-/C_A^+ . The sensitivity to C_T^-/C_A^+ is much less dependent on the Q value and is more than an order of magnitude worse than for C_T^+/C_A^+ .




Figure 1.10: Median expected sensitivity at 68% CL to (a) $|C_T^+/C_A^+|$ (assuming $C_T^-/C_A^+ \equiv 0$) and (b) $|C_T^-/C_A^+|$ (assuming $C_T^+/C_A^+ \equiv 0$) as a function of Q value for 100 million events.

Finally, we make a comparison to the \tilde{a} treatment. As a reminder, this treatment sets $b \equiv 0$ and fits a as a free parameter; lifting this restriction via equation (1.5) allows for a limit on b to be set. As demonstrated in Figures 1.6 and 1.7, the BSM contribution from the b term begins to dominate over that of a as $C_T^+ \rightarrow 0$. This implies that the \tilde{a} prescription becomes invalid for experiments sensitive to small C_T^+/C_A^+ contributions (most modern experiments). Note that the \tilde{a} prescription is perfectly valid under the assumption $C_T^+ \equiv 0$ when searching for only a C_T^- contribution, as this case contains no b contribution. This is noted in Ref. [17], which suggests as a remedy to 1) fix a to the SM value and fit b or 2) fit both a and b as free parameters, noting the correlation between them.

A comparison between these several treatments is shown in Figure 1.11 for two different Q values with 4 million events. In addition to the direct comparison of the spectra, two other analyses of the p_r spectrum are shown: the \tilde{a} treatment (light pink) and a treatment in which a is fixed to the SM value and b is fit (light blue).⁶ Only the median expected 68% CL are shown for simplicity.

There are a few items of note. First, the \tilde{a} treatment results in an allowed parameter space very similar to that of the $\theta_{\beta\nu}$ spectrum; this makes sense given that the \tilde{a} treatment integrates over E_e . The \tilde{a} treatment also perfectly coincides with the proper treatment on the $C_T^+/C_A^+ = 0$ axis, where $b = 0$, as expected. Correspondingly, the treatment where b is fit agrees well with the proper treatment on the $C_T^-/C_A^+ = 0$ axis, where the contribution from b dominates over the BSM contribution from a . This agreement is shown in Figure 1.12 for different Q values with 4 million events.

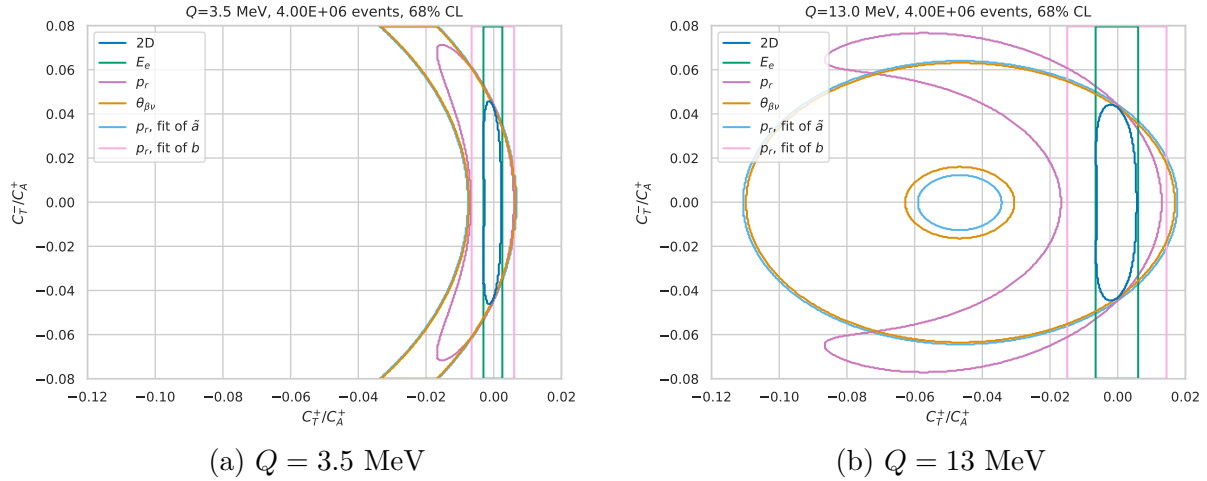


Figure 1.11: Median expected sensitivity at 68% CL for two different Q values with 4 million events, demonstrating the differing results obtained with different spectra and analysis approaches. See text for details.

Most important, though, is the fact that the \tilde{a} treatment gives an allowed parameter space that is much larger than the proper treatment. Values of $(C_T^+/C_A^+, C_T^-/C_A^+)$ that should be

⁶Fitting both a and b as free parameters and incorporating the correlation between them is difficult to interpret in the $(C_T^+/C_A^+, C_T^-/C_A^+)$ parameter space and is thought to give result largely similar to the proper treatment that already includes the correlation in the definitions of a and b , i.e. equations (1.9) and (1.10).

excluded by the experiment are not; for Q values smaller than about 3 MeV, the \tilde{a} treatment actually excludes *too much* of the parameter space compared to the proper treatment, as seen in Figure 1.12. This observation has been noted in Ref. [17] previously, though without a comparison to the proper treatment that includes the correlation between a and b .




Figure 1.12: Median expected sensitivities at 68% CL for different Q values with 4 million events; the limit is (nearly) symmetric about zero and hence the absolute value is shown. (a) Expected sensitivity for $|C_T^+/C_A^+|$ assuming $C_T^-/C_A^+ \equiv 0$, and (b) expected sensitivity for $|C_T^-/C_A^+|$ assuming $C_T^+/C_A^+ \equiv 0$.

We have not considered the enhancement to sensitivity from delayed particle decays. As mentioned above, this is more complicated to perform general sensitivities of due to the more complex expressions and higher dimensionality. This will be tackled in the future, though. For $2^+ \rightarrow 2^+ \rightarrow 0^+$ delayed α decays (the case for ${}^8\text{Li}$, ${}^8\text{B}$, and ${}^{20}\text{Na}$), the enhancement is a factor of 3 for events where the α momentum is anti-parallel to the β momentum. With 100 million of such events, this should improve the sensitivity of p_r to C_T^- in Figure 1.10(b) by a factor of 3.

1.5 Structure and decay of ${}^8\text{Li}$ and ${}^8\text{B}$

Both ${}^8\text{Li}$ ($J^\pi = 2^+$, isospin $T = 1$, $Q_\beta = 16.00413(6)$ MeV [48]) and ${}^8\text{B}$ (2^+ , 1, 16.9579(10) MeV) predominantly decay to a broad 3 MeV excited state ($\Gamma^\circ \sim 1.4$ MeV) in ${}^8\text{Be}$ (2^+ , 0), with decays to a higher 16 MeV doublet also contributing. Only the 2^+ states are populated

by the decay due to selection rules. The excited ${}^8\text{Be}^*$ subsequently α -decays within $\sim 10^{-22}$ seconds. The α particles therefore have energies up to ~ 8.5 MeV; due to the recoil momentum of the ${}^8\text{Be}^*$, the difference in α energy can be up to ~ 400 keV, with up to a few degrees of angular separation in the lab frame. Note also that the Q value of the decay is distribution and varies according to the excitation energy in ${}^8\text{Be}$ to which the decay occurs.

The known decay scheme of ${}^8\text{Li}$ and ${}^8\text{B}$ is illustrated in Figure 1.13 with only relevant levels shown. Note that the 2^+ states at 16.6 and 16.9 MeV are mixtures of pure $T=0$ and $T=1$ states [49]. Both ${}^8\text{Li}$ and ${}^8\text{B}$ are nearly pure Gamow-Teller decays; the amount of a Fermi component due to this high-lying $T=1$ state has been calculated to be below $5.0(1.5) \times 10^{-4}$ [50]. The spin sequence for the whole decay is $2^+ \rightarrow 2^+ \rightarrow 0^+$, with the α particles having orbital angular momentum $\vec{L} = 2$. As both ${}^8\text{Li}$ and ${}^8\text{B}$ are pure GT decays, they are sensitive to exotic tensor couplings, as discussed in previous sections.




Figure 1.13: Decay scheme of ${}^8\text{Li}$ and ${}^8\text{B}$ with only relevant levels shown.

In addition to the known 2^+ levels, it has been argued, first by Barker [51, 52], that an “intruder” 2^+ state exists between the 3 MeV and 16 MeV states. Such an intruder state has not been unambiguously confirmed experimentally by other authors.⁷ However, a recent calculation of the recoil corrections for the ${}^8\text{Li}$ decay suggested that such an intruder state

⁷A succinct history of the controversy over intruder states in ${}^8\text{Be}$ is provided in Ref. [53].

may exist [54]. Notably, in Ref. [54], it also was identified that the different 2^+ states in the decay have vastly different recoil-order corrections associated with them. This means that it is necessary to know the relative state contributions as a function of excitation energy in ^8Be . In principle, this can be determined by $\alpha - \alpha$ scattering phase shift data or by an R-matrix fit to the ^8Be excitation energy spectrum, but neither technique has been sensitive enough to conclusively rule out an intruder state.

In Ref. [38], due to uncertainty over the presence of an intruder state, the range of excitation energies used in the $\beta - \nu$ angular correlation measurement was confined to energies close to the 3 MeV level to reduce uncertainties associated with the presence of the intruder state and the resulting impact on recoil-order corrections. This is shown in Figure 1.14, which demonstrates that R-matrix fits to the excitation energy spectrum cannot distinguish between the two cases.




Figure 1.14: From Ref. [38]: Level scheme of ^8Be with proposed intruder state. Two different R-matrix fits are shown (a) without the intruder state and (b) with the intruder state. The fit is equally good and cannot distinguish between the two cases. In light blue, the selected excitation energy range around the 3 MeV peak for the angular correlation measurement to reduce possible contamination from the intruder state.

It is absolutely essential to include recoil-order terms in describing the decay of ^8Li and

^8B . The full, recoil-order expression for the decay may be found in Ref. [21] with a concise explanation of the various terms and estimate of their size for ^8Li provided in Ref. [41]. These terms have been measured experimentally by Refs. [55–58], but these experiments averaged over the excitation energy and did not consider the state-dependent nature of the form factors. In Ref. [54], the uncertainties on these form factors was reduced by 50% from theoretical calculations, and the state-dependent nature of the form factors was also identified.

1.6 $\beta - \nu$ angular correlation in trapped ^8Li

To measure the recoil energy spectrum of ^8Li , both α need to be detected. In addition, to make use of the enhancement to sensitivity from the delayed α (section 1.3.1), the direction of the β must also be detected. In the BPT experiment, this is performed by trapping ^8Li ions at almost zero kinetic energy in an ion cloud of a few mm^3 size. A schematic showing the general experiment setup is provided in Figure 1.15. By surrounding the ion cloud with position-sensitive silicon detectors (chapter 4), the energies and positions of the two α and the position of the β can be detected, allowing for a full kinematic reconstruction of the event. Events where two α and the β are detected are referred to as “triples” while events where only two α are detected are called “doubles.” Comparisons between these types of events allows for checks on systematic uncertainties. Events where the two α particles are perpendicular to the β can also be used for systematic checks.

Notably, only events where the β is (approximately) parallel to an α are used for the final search for tensor couplings, resulting in a factor of 3 increase in sensitivity. These events have the largest kinematic shift of the α particle energies. In addition, because the α particles are nearly back-to-back, the energy difference between the α particles contains information about the kinematic shift. Using this technique, the BPT has published several limits on tensor couplings in ^8Li [36–38] and in ^8B [26].

The α energy difference spectrum from the latest BPT measurement [38] is shown in




Figure 1.15: From Ref. [38]: General setup of the BPT measurement of ^8Li decays.

Figure 1.16. Under the assumption of $C_T \equiv -C'_T$ ($b \equiv 0$), the experiment found $|C_T/C_A|^2 = 0.0012 \pm 0.0019_{\text{stat}} \pm 0.0028_{\text{sys}}$, a result consistent with the SM. The systematic uncertainty budget from this measurement is shown in Table 1.2.

By taking the α energy difference for parallel β events, sensitivity to exotic tensor couplings is maximized, while systematic uncertainties associated with the detector response are minimized. This comes from the fact that the energy of the α anti-parallel to the β is always subtracted from that of the α parallel with the β . This effectively forms a kind of asymmetry measurement where both common and individual detector uncertainties are subtracted out, though uncertainties that are α energy-dependent can still play an important role (such as detector dead layers).

Scattering of the β is an important systematic uncertainty, as well. In the experiment, β may scatter off of detectors or trap surface and into a detector coincident with an α ; this would then look like an event where the β was parallel with the α and be included in the measurement spectrum. This effect is modeled in simulations but it is still an important




Figure 1.16: From Ref. [38]: Final α energy difference spectrum which is compared to simulation (black curve) to place limits on tensor couplings. The blue curve shows the expected spectrum in the case of purely tensor couplings. n.b. A +0.0005 correction was added to the result shown in the figure to account for possible contamination from the intruder state; see Ref. [38] for details.

Systematic uncertainty	$\Delta C_T/C_A ^2$ (2016 data)	$\Delta C_T/C_A ^2$ (2022 data)	
recoil-order terms	1.3×10^{-3}	***	} theory improvements
bremsstrahlung / radiative corrections	8×10^{-4}	8×10^{-4}	
fractional state distribution	5×10^{-4}	5×10^{-4}	
cuts to data	9×10^{-4}	$\sim 5 \times 10^{-4}$	} detector characterization
α energy corrections / calibration	7×10^{-4}	$\sim 4 \times 10^{-4}$	
α Si detector lineshape	9×10^{-4}	$\sim 5 \times 10^{-4}$	} new trap design
β scattering	1×10^{-3}	$\sim 3 \times 10^{-4}$	
Total	2.8×10^{-3}	$\sim 1.3 \times 10^{-3}$	

Table 1.2: Uncertainty budget from measurement of Ref. [38] and the expected uncertainty budget from the ^8Li data campaign performed in this work. ***Key recoil-order terms will be attempted to be measured experimentally from the data by performing a “global” fit to the data, which can reduce or eliminate this systematic uncertainty.

systematic uncertainty.

The largest uncertainties come from theory corrections, particularly the recoil-order form factors, despite the recent reduction to their uncertainty by 50% [54]. Radiative correction uncertainties come in part from the correlation with the delayed α and can likely be improved by theory. The final state distribution in ${}^8\text{Be}$ will require additional experimental and theoretical work.

1.6.1 Recent studies of the ${}^8\text{B}$ 2α spectrum

No BPT student thesis has yet discussed the ${}^8\text{B}$ 2α spectrum so a brief discussion is provided here. Notably, some possibly unidentified problems with previous measurements are pointed out.

Historically, experiments to study the decay of ${}^8\text{B}$ have been conducted for at least 70 years [59], with experiments studying ${}^8\text{Li}$ going back at least to the 1930s.⁸ More recent experiments (post-2000) have greatly improved energy resolutions and more rigorous systematic uncertainty quantification, and these experiments are generally used for comparison. Currently, the most promising technique to measure the ${}^8\text{B}$ neutrino spectrum is by using silicon detectors to detect the energies of the 2 back-to-back α particles from the decay, either in the same detector or in separate detectors. This excitation energy spectrum can then be transformed into the un-oscillated neutrino spectrum, which is an important input to the analysis of solar neutrino observatories [61]. The shape uncertainty associated with the solar neutrino spectrum is particularly important, though no quantification of uncertainty requirements for future observatories has been recently performed.

Several recent experiments from about the last 20 years are listed in Table 1.3; most of these were interested in obtaining the ${}^8\text{B}$ neutrino spectrum by measuring the 2α spectrum. Ortiz et al. [62] is particularly notable in that this was the first experiment to detect the 2 α particles in separate detectors, though the measurement suffered from unreported systematic

⁸See Ref. [60] for a brief historical overview.

Authors	Notes	Publication year(s)	Relevant references
Ortiz et al.*	foil implantation of ^8B with a PIN detector	2000	[62, 63] (see also [61])
Winter et al.	direct implantation of ^8B into an SSB	2003, 2006	[60, 64, 65] (see also [66])
Bhattacharya et al.	foil implantation of ^8B with SSBs	2006	[67]
Kirsebom et al.	foil implantation of ^8B with DSSDs	2011	[53, 68, 69]
Roger et al.	direct implantation of ^8B into a DSSD	2012	[70] (see also [68, 71, 72])
Viñals et al.	foil implantation of ^8B with DSSDs	2020	[73, 74]
Longfellow et al. (BPT)	trapped ^8B with DSSDs	2023	[75]

*Ortiz et al. suffered from additional experimental uncertainties not reported anywhere and should be disregarded: see Ref. [61], p. 225.

Table 1.3: Some recent experiments studying the ^8B 2α spectrum with silicon detectors.

uncertainties that mean it should be disregarded [61]. Winter et al. [64] and Bhattacharya et al. [67] used silicon surface barrier (SSB) detectors, while Kirsebom et al. [68] and Roger et al. [70] used double-sided silicon strip detector (DSSD)s. See Table 1.3 for additional references for these experiments. Viñals et al. [74] used DSSDs and was primarily interested in studying the 16 MeV doublet and possible electron-capture decay branch; the experiment is included in Table 1.3 as it is referenced later in this work. Commentary back and forth between these experiments (especially Refs. [66, 71, 72]) focuses primarily on detector responses and calibration, highlighting the importance that detector characterization plays in determining the energy spectrum accurately.

A key distinction between experiments over the last 20 years is whether the ${}^8\text{B}$ is implanted onto a thin catcher foil, implanted directly into a silicon detector, or held in an ion trap as with the BPT. Each technique has advantages and disadvantages. In particular, using a catcher foil means that α particles from the decay will necessarily lose energy in the foil and detector dead layers, affecting the final deposited energy of the particle. Contributions from dead layers are avoided with direct implantation into a silicon detector, but the outgoing β then contributes some energy to the signal. Ion traps avoid foil losses since the ions are held in a cloud but are still susceptible to detector dead layer losses.

At first glance, then, direct detector implantation might seem to be the most promising approach since it avoids both of these losses. However, direct implantation into a silicon detector should ideally be deep enough to ensure that α particles cannot escape the detector without depositing their entire energy. The implantation energy for ${}^8\text{B}$ should therefore be ~ 35 MeV. Unfortunately, both experiments using the direct implantation technique implanted ${}^8\text{B}$ beams at lower energies, resulting in some α particles escaping without full energy deposition. In addition, direct implantation suffers from the fact that the β from the decay will deposit energy in the detector, around 25 keV, though this can be corrected for with knowledge of the angle of the β . The effect of β -summing can also be greatly suppressed either by including a second detector to select events where the β exits the silicon detector

close to the normal of the detector surface, or by using a pixelated detector to veto events where energy is deposited outside of the implanted pixel [76].

The Winter et al. experiment implanted ^8B directly into a silicon surface barrier detector; the ^8B energy was 27.3 ± 0.2 MeV resulting in an implantation depth of 42.2 ± 2.0 μm with a distribution full width at half maximum (FWHM) of 0.7 μm [65]. The detector thickness was 91 μm to minimize β energy deposition. However, this implantation depth is not deep enough to collect the full energy from α particles $\gtrsim 7.5$ MeV exiting at angles close to normal to the detector surface.⁹ This energy loss would distort the 2α spectrum above about 14 MeV, with about 15% of these events depositing less energy than expected (depending on α energy, emission position, angle, and path straggling). Low statistics in this part of the spectrum means that the impact of this distortion is likely small. In addition, for solar neutrino experiments, the high energy neutrino spectrum is most important, corresponding to the lower energies of the 2α spectrum, so the impact here should also be very small. To the best of my knowledge, this source of distortion has not been pointed out before and could be responsible for some of the spectrum disagreements at higher energies with the other experiments.

The Roger et al. experiment also suffers from an implantation depth that is too shallow to fully capture the energies, with an even more dramatic effect. Here, the DSSD was 78 μm thick and the implantation depth was 26 ± 5 μm , with the ^8B beam having energy 18.69 ± 0.27 MeV. This means that α particles with energies as low as 3-4 MeV will escape the detector without depositing the full amount of their energies, severely distorting the spectrum. The correction factor varies from a few percent at the 3 MeV peak to up to a factor of 5 near the 16 MeV doublet. From simulations, a correction function is obtained and applied to the measured spectrum in order to recover the true energy spectrum. However, the shape uncertainty associated with the correction function is not reported nor does it seem to have been included in the recovered spectrum.

⁹Expected α ranges are provided in Figure 4.2b.

Given that the BPT experiment is primarily interested in the β - ν angular correlation in ${}^8\text{Li}$ and ${}^8\text{B}$, it is not possible to use detector implantation method since the momentum of each α cannot be measured. This means that the BPT must account for detector dead layers. Since ions are held in free space instead of being implanted in a foil, though, there are no source losses. As with the other experiments mentioned here, accurate knowledge of the BPT detector response is extremely important and the complete characterization of the DSSDs is presented in Chapter 4.

Winter et al., Kirsebom et al., and Roger et al. all used the β -delayed α decay of ${}^{20}\text{Na}$ for calibration points as low as 2.15 MeV α energy. Calibration at these lower energies is obviously very desirable given that the α energy peaks at 1.5 MeV for the break-up of ${}^8\text{Be}^*$. Unfortunately, the BPT has not been able to use ${}^{20}\text{Na}$ in its calibration; in previous attempts, the isotope has not been produced and cooled at high enough rates to merit trapping for calibration. Further beam development for ${}^{20}\text{Na}$ production would be of value for the experiment.

1.7 This work

This work reports on several experimental improvements that will reduce systematic uncertainties associated with β scattering and the detector response. The various response functions of the BPT experiment are shown in Figure 1.17. Each of these distributions contributes (to varying degrees) to the systematic uncertainty of the final result, the limit on tensor couplings. In this thesis, experimental improvements are made to reduce the systematic uncertainties associated with these distributions.

In chapter 2, the design and construction of a new trap, the BPT Mk IV, is described. This trap features a number of innovations to simultaneously reduce β scattering and detector noise from the nearby electrodes. In chapter 3, the BPT Mk IV is used for a new ${}^8\text{Li}$ data run, which collected about 30% higher statistics than previous BPT measurements

with reduced systematic uncertainties.¹⁰ In chapter 4, the silicon detectors used with the BPT are directly and fully characterized with an α beam, resulting in the measurement of dead layer thicknesses and the identification of an energy-dependent charge-sharing effect. Finally, in chapter 5, conclusions from this work are drawn and new possibilities for future experiments are suggested. In addition to the experimental improvements, a new exploration of the correlations between a and b are provided in section 1.4, with an eye towards next generation experiments.

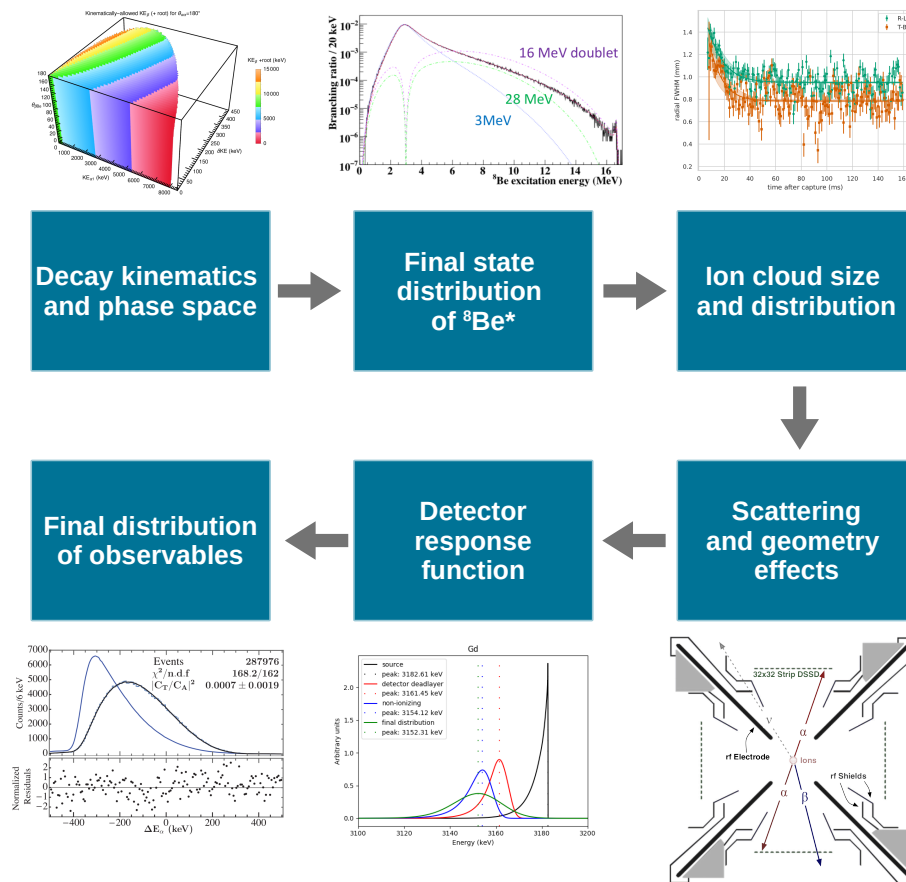


Figure 1.17: Effects involved in the final distribution of observables (the α energy difference spectrum) for the BPT measurement.

¹⁰In addition to the 2022 ^8Li data campaign, I also participated in a 2019 ^8B data campaign, as well as many experiments using the Canadian Penning Trap, which are not detailed here.

Chapter 2

THE BETA-DECAY PAUL TRAP MK IV

Portions of this chapter have been adapted from Ref. [77].

2.1 A brief history of the Beta-decay Paul Trap

The BPT has performed many successful experiments since its first publication in 2012 [78] and has been the subject of at least eight previous doctoral theses [36, 41, 79–84], with the contributions of many additional masters students, doctoral students, postdocs, and staff. Over the course of this experimental effort, reducing RF pickup has been one of the major challenges; indeed, one of the motivations for the trap redesign in this work was further reduction of RF pickup. The trap has undergone a number of redesigns with the primary aim of reducing RF pickup on the DSSDs in its history. A brief summary of the design changes in the BPT is provided in Table 2.1; this is likely not an exhaustive list of publications. Note the BPT version with a trap electrode radius to center $r = 11$ mm; this version was used primarily to study isotopes that undergo β -delayed neutron decay and required a smaller RF voltage so as to minimize the energy disruption of the recoiling ion. A dedicated trap for β -delayed neutron spectroscopy (BEARtrap) is currently under development [85].

In addition to the physical RF reduction strategies used, all of the versions of the BPT also rely upon long shaping times to integrate out the high frequency RF pick up (section 2.6.1). Some earlier versions also incorporated low pass filters as part of their shaping amplifiers. The previous version of the BPT [86] introduced tunable notch filters on each DSSD front strip that effectively remove most of the RF amplitude at the trap frequency (section 2.6.2). The previous version of the BPT also added “RF outer shields” which act as Faraday cages to suppress RF pick up; unfortunately, these shields were found to dramatically increase β scattering.

In this work, the design of the BPT is re-visited once again in an attempt to reduce RF


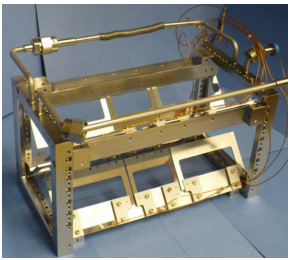
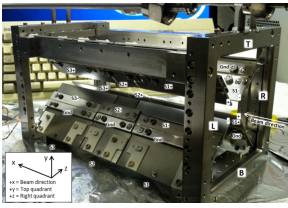


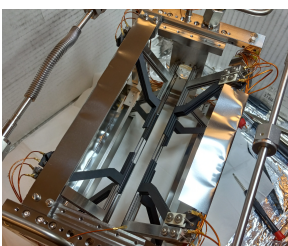
BPT picture	RF reduction strategies	First publication year	Relevant references
	RF shielding	2012	[78]
	hollow electrodes, RF shielding, nickel mesh shielding	2012	[36, 87]
	split electrodes ($r = 11$ mm), RF shielding	2013	[79, 81, 82, 84, 88–92]
	split electrodes ($r = 17.4$ mm), RF shielding	2013	[26, 37, 80]
	split electrodes ($r = 17.4$ mm), add'l. RF shielding, notch filters	2018	[38, 41, 75, 83, 86]
	rod electrodes, RF shielding, notch filters	2023	this work

Table 2.1: A brief history of changes to the BPT; the last row is the BPT Mk IV. The version with a trap electrode radius to center $r = 11$ mm was used primarily to study isotopes that undergo β -delayed neutron decay. Pictures taken from Refs. [78, 80, 81, 87]

pick up as well as to reduce the incidence of β scattering. As a result of this work, the BPT Mk IV has been designed, constructed, and commissioned and is the subject of the rest of this chapter.

2.2 Physics of a Paul trap

From Maxwell's equations for electromagnetism, it is clear that a charge cannot be stably trapped at rest by a purely electrostatic field, a result which is called Earnshaw's theorem.¹ In a simple way, this is because a stable trapping point requires a local minimum, where the divergence of the field is negative. However, this point then must act as a sink, which is not possible in free space from Gauss's law. Note, though, that this argument only holds for charges at rest in static fields.

A Paul trap uses oscillating inhomogeneous electric fields (without the use of magnetic fields) to trap ions almost at rest in free space. By selecting the correct oscillation frequency for the ion mass and charge of interest, a harmonic pseudo-potential can be formed that stably traps ions. The development of the Paul trap has had a significant scientific and industrial impact, resulting in Paul's Nobel Prize in 1989 [93]. Paul traps in various forms are now common throughout many laboratories, and their development for numerous applications continues [94]. For nuclear physics, ion traps (both Paul traps and other types of ion traps) have had many recent developments and uses, as summarized in Ref. [95]. Paul's original trap was formed by a three-dimensional hyperbolic trap. For the BPT, a linear Paul trap is instead used, where trapping along the axial direction is formed by a DC potential only.

The Mathieu equations² describe the ion motion in a Paul trap. The derivation of these equations from the potential has been shown many times over the past 70 years, including in many previous student theses (e.g. [41, 80, 87]). This work will not re-derive them but will

¹Due to the fascinating and convoluted historical development of the theory of electromagnetism in the 18th and 19th centuries, Earnshaw's result came about decades before Maxwell wrote his equations that unified the work of other prominent names. This, again, was decades before Heaviside's transcription of Maxwell's equations into the vector calculus form familiar to us today.

²Also first introduced in the 19th century.

present the most important aspects, particularly those relevant to the design of the BPT Mk IV (section 2.5). For a detailed derivation and discussion, see Ref. [94] and references therein.

The ideal electric potential for a Paul trap has a hyperbolic shape, which, in two dimensions, can be formed by four electrodes arranged with quadrupole symmetry. Considering only two dimensions, the potential at a particular point in time is written

$$\Phi(x, y) = \frac{\Phi_0}{2r_0^2}(x^2 - y^2), \quad (2.1)$$

where r_0 is the minimum distance from the center to the surface of the hyperbolic electrodes that have an applied potential of $\pm\Phi_0/2$. This hyperbolic potential is illustrated in Figure 2.1. The static field will not trap an ion since the center is clearly a saddle point. However, trapping can occur in one direction (dependent upon the sign of the ion charge), and, if the potential were inverted, then trapping could occur along the other direction. This, then, is the basic idea of a Paul trap: to oscillate a quadrupole potential so that the average potential traps an ion in both (or all three) dimensions.

In the case of the BPT, the electrodes are decidedly not hyperbolic. However, from Laplace's equation and by arguments of symmetry, it is required that the potential near the center form a hyperbolic potential. Away from the center, higher-order anharmonicities may be present which reduce the performance of the trap by introducing instabilities in the ion motion [96]. As shown in section 2.5, the goal in designing a linear Paul trap is to maximize the extent of the hyperbolic region, while maintaining optical access to the ion cloud. For the BPT Mk IV, an equivalent hyperbolic trap would have electrodes at a radius $r_0 = 16.5$ mm, compared to a front-face distance of 15.3 mm in the actual trap.

Continuing with the physics, the potential is allowed to oscillate at frequency $\Omega = 2\pi f$ with peak-to-peak amplitude V (the opposite phase electrodes have maximal amplitudes of $\pm V/2$), giving $\Phi_0 = V \cos(\Omega t)$. A DC offset $\pm U$ on the opposite phase electrodes can be added to provide mass filtering. The equations of motion for an ion of mass M and charge Q




Figure 2.1: Illustration of a quadrupole potential with some equipotential lines shown. Solid lines are positive potential; dashed lines are negative. The straight lines along the diagonals are at zero potential.

in this potential are then given by the Mathieu equations. For the mass-filtering case, they are

$$\begin{aligned} \ddot{x} + (a - 2q \cos 2\tau)x &= 0, \\ \ddot{y} - (a - 2q \cos 2\tau)y &= 0, \end{aligned} \tag{2.2}$$

where

$$a \equiv \frac{4QU}{M\Omega^2 r_0^2}, \quad q \equiv \frac{2QV}{M\Omega^2 r_0^2}, \quad \text{and} \quad \tau \equiv \frac{1}{2}\Omega t, \tag{2.3}$$

and the derivatives are with respect to τ . The BPT operates without mass-filtering ($a = 0$), and the reader is directed to Ref. [94] to further consider mass-filtering applications. Ions can be stably trapped in several regions in the (a, q) phase space (see e.g. Refs. [41, 80, 87] for examples of the stable region). Most non-mass-filtering ion traps, including the BPT, operate in the first stability region with $q < 0.908$, with $q \approx 0.4$ conventionally taken as the

ideal practical case [97].

Ions in this oscillating potential undergo both a “macromotion” and a “micromotion.” The macromotion is a simple harmonic motion in the time-averaged trap potential, modulated by small-amplitude micromotion from the driving RF field. (Again, see e.g. Refs. [41, 80, 87] for examples of the ion motion.) Since we are usually interested in trapping ions on a timescale much longer than several RF phases, this motion can be averaged over many RF periods. This leads to an approximation of the time-averaged Paul trap as a harmonic pseudo-potential (see Ref. [94] for a complete derivation). The radial pseudo-potential at a radius r is given by

$$\Phi_{\text{eff}}(r) = \frac{QV^2}{4M\Omega^2r_0^4}r^2 = \frac{M\Omega^2}{16Q}q^2r^2. \quad (2.4)$$

In a linear Paul trap, the radial pseudo-potential is modified by axial defocusing, which can usually be neglected when the central trapping segment has a length longer than the radial spacing of the electrodes. From Laplace’s equation, the DC potential applied along the axis will diverge in the radial (RF) plane. The effect on the pseudo-potential has been considered in Ref. [98]. Practically, tuning of the trapping amplitudes is required so that the RF force overcomes the defocusing. For the BPT Mk IV, this effect is observed to be small. Because the axial potential is formed by DC voltages, it is a simple harmonic potential close to the trap center that does not depend on the ion species or RF frequency, only on the field gradient along the axis.

Of critical importance to the operation of the BPT and most other Paul traps is the inclusion of a cooling mechanism, typically a helium buffer gas. The applications of buffer gas are explored in detail in Ref. [99]. In an ion trap, the buffer gas collides with the trapped ions and removes energy from them, thermalizing them to the temperature of the gas. Collisions with the buffer gas also effectively couple the radial and axial modes of motion. Buffer gas also reduces the emittance of the injected ions, allowing for more efficient trapping and transport. The effect of buffer gas on trapped ions can be modeled in the

Mathieu equations (Eq. 2.2) through the inclusion of a drag term [41, 87], which increases the size of the stability region.

2.2.1 Ion cloud energy and spatial distributions

Conveniently for the BPT, the back-to-back α particles from the decay allow for direct imaging of the ion cloud spatial distribution. The equations of motion for individual ions are not important for the BPT since the experiment is only sensitive to the spatial distribution of the decays. During the design process of the BPT Mk IV (section 2.5), an estimate of the ion cloud spatial distribution was however required.

The energy and spatial distribution of the ion cloud is of considerable interest for many experiments but is also difficult to calculate. This difficulty comes from the various effects of ion-ion interactions (space charge), RF heating, and buffer gas collisions. It has been observed and computed that a Gaussian spatial distribution describes well the ion cloud extent for helium buffer gas for most values of q when $a = 0$ [100]. In this case, it was observed that the ion velocity distribution matches well the expectation from the Maxwell-Boltzmann distribution.

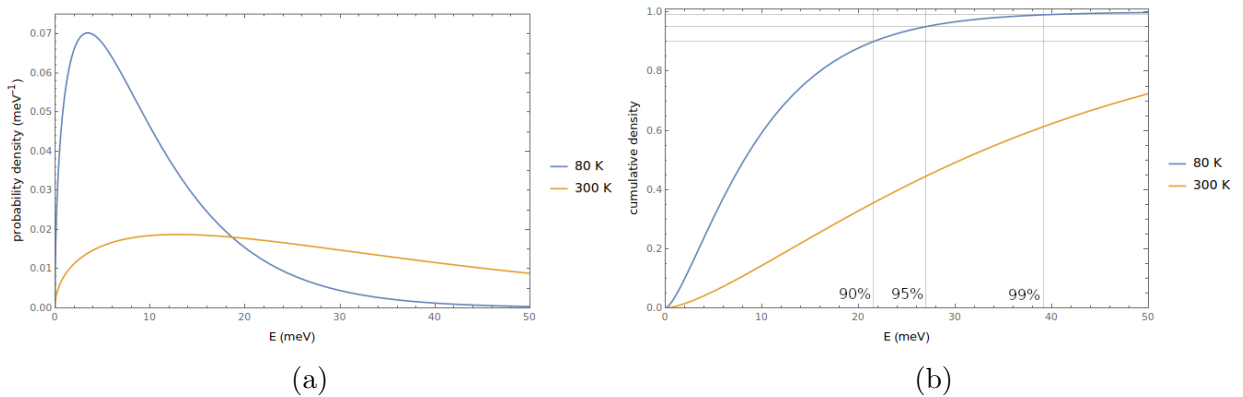


Figure 2.2: Maxwell-Boltzmann distributions of energies at liquid nitrogen temperature (80 K, blue) and at room temperature (300 K, orange). (a) Probability density. (b) Cumulative distribution with the energy extents of 90%, 95%, and 99% of particles at 80 K indicated.

In the BPT Mk IV design process, the Maxwell-Boltzmann distribution was naively taken to describe the energy distribution of the ions after thermalization with the helium buffer

gas. Using the quadratic coefficient c that describes the field gradient of the radial pseudo-potential or the axial DC potential, the energy of ions can be matched to the height of the potential to find the radius at which the ions are contained. It is also easy to solve for the required field gradient to achieve a particular ion cloud size. Note that this treats the two potentials in (r, z) independently which is not a fully accurate treatment.

$$r = \sqrt{\frac{E}{c}} \quad (2.5)$$

Two Maxwell-Boltzmann energy distributions are shown in Figure 2.2, at 80 K and at 300 K. This figure shows the significant advantage that using a cryogenic ion trap brings. At 80 K, 99% of ions have an energy less than ~ 40 meV, compared to ~ 150 meV at 300 K. The voltage required to contain the ions in the same volume is then almost a factor 2 smaller at this lower temperature.

2.3 General BPT description³

Before describing the design of the new BPT Mk IV, a general description of the trap is provided; these details are the same between trap versions.

The BPT is a linear Paul trap: it uses a quadrupole arrangement of electrodes with voltages oscillating at radio-frequencies to confine ions in the radial plane and uses segmented, static voltage electrodes to confine ions in the axial direction. An ideal Paul trap uses hyperbolic electrodes to maximize the physical extent of the trapping potential while minimizing its anharmonicity. Any arrangement of electrodes with quadrupole symmetry will have a hyperbolic potential near its center, enabling the electrode shape to have a more open geometry to allow for optical access to the ion cloud. Different electrode geometries, however, will have an impact on the size of this hyperbolic region, as discussed in section 2.5. Previous changes to the trap electrodes and experimental components have been noted in section 2.1.

³This section is adapted from Ref. [77].

The entire BPT and vacuum chamber meet ultra-high vacuum (UHV) standards to maintain an environment free from out-gassing that might affect the trapping lifetime. The typical pressure achieved prior to cryogenic cooling is $\lesssim 1 \times 10^{-8}$ mbar. The BPT uses ultra-pure 99.999% helium gas at a $\sim 10^{-5}$ mbar pressure to cool ions through collisions with the gas. The frame of the trap is hollow and liquid nitrogen is circulated to bring the trap and gas to approximately 80 K [101]: in this way, the ions of interest are thermalized to this lower temperature, reducing the overall size of the ion cloud. As shown in section 2.2.1, a low thermalization temperature is critical for the experiment. From the Maxwell-Boltzmann distribution at 80 K, 99% of ions have an energy less than ~ 40 meV, compared to ~ 150 meV at room temperature. Therefore, the voltage required to contain the ions in the same volume is almost a factor 2 smaller at this lower temperature, reducing RF pickup on the detectors. In addition, the cryogenic temperature greatly reduces out-gassing of trap construction materials, leading to a longer trap lifetime of at least tens of seconds [78]. The BPT rests on four quartz spheres to provide thermal isolation from the vacuum chamber.

The ion capture, transport, and delivery system are described in section 3.1. Details of the BPT trap cycling are provided in section 3.3.

2.4 β scattering in the previous BPT⁴

As seen from the uncertainty budget in Table 1.2, β scattering is an important experimental systematic uncertainty. In designing the BPT Mk IV, the current sources of scattering must be understood in the previous trap.

The BPT has an open geometry, but material and design choices still strongly impact β scattering. A modified version of the BPT Geant4 (release 10.5.1) simulation code was used to study the effect of different design choices on β scattering. The Geant4 simulation used the “option3” standard electromagnetic physics list [102–104]. A CAD model of the BPT was designed in Autodesk Inventor [105] then exported to STP files. Using a modified version of

⁴This section is adapted from Ref. [77].

an existing script [106], these files are then converted to GDML files readable by Geant4. An ion cloud of realistic dimensions (obtained from previous experiments) and decay spectrum was simulated. Simulated events where both α particles hit opposite DSSDs and the β hits a coincident detector were used to determine the scattered ratio (triple correlation events prior to final analysis cuts). If a detector hit of a β did not match its true original momentum, it was counted as scattered. Statistical uncertainty on the scattered fractions reported below is $\sim 0.1\%$ with $\sim 150,000$ triple coincidences from each simulation.

The previous BPT was found to have 21.1% of triple coincidence events from scattered β particles. It is not entirely possible to cleanly separate physical sources of scattering, but by adding and removing trap elements in the Geant4 model, the major contributions can be understood. Of the total 21.1% scattering, 5.5% is attributed to the 0.01” thick stainless steel “RF outer shields” (Figure 2.4e) used in an attempt to reduce to RF pickup on the DSSDs [80, 86]; these outer shields were later found to have no measurable impact on RF pickup reduction and therefore were removed for the BPT Mk IV design. Additional metal elements surrounding the DSSDs (0.006” thick stainless steel “RF hoods” and 1/16” thick gold-plated aluminum detector covers, Figure 2.4e) were found to have some RF shielding efficacy but contributed an additional 5.5% to β scattering. These elements were redesigned. A further 5.6% is attributed to the 2 mm thick stainless steel flat electrodes⁵ (Figure 2.4c), primarily the center electrode that forms the axial potential well (Figure 2.4a). The redesign of the electrodes is the focus of the next section. The remaining 4.5% comes from back-scattering on the DSSDs and scattering on the electrode support structure. This last portion is considered largely irreducible in the BPT Mk IV design, as this would require an extensive re-design of the vacuum vessel and supporting plastic scintillator detector locations.

Note that both material choice and material thickness affect β scattering. Figure 2.3 shows this impact by changing the RF hood thickness and material in the Geant4 simulation and looking at the effect on the scattered fraction of events. Note that, for stainless steel

⁵Reported incorrectly as 1.6 mm in [86].

(blue), the scattering rises to a maximum and then falls off. It is expected that a similar behavior for aluminum would occur, but likely outside of the tested (and practical) thickness range. It is thought that two competing effects are contributing: 1) the typical path length between scatters of a β in the material, and 2) the stopping power of the material. For thicknesses below or close to the typical path length between scatters in that material, multiple scatters increase the chance for a β to hit a detector. For large enough thicknesses, though, the β particles can lose most or all of their energy inside the material, leading to a reduction in the scattered fraction. In practice, it is impractical to have a very thick RF hood design (and the scattered fraction is still considerably higher than at thinner thicknesses).

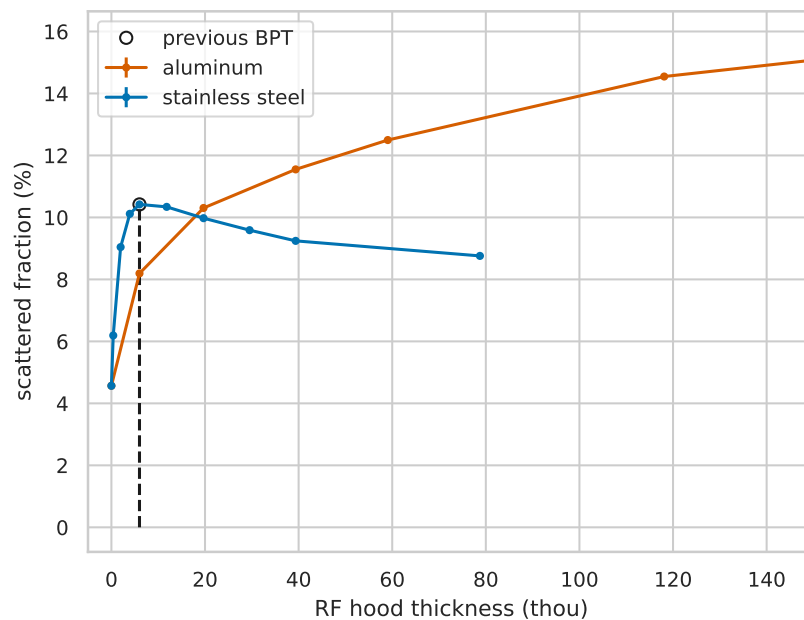


Figure 2.3: Effect on β scattering of material and thickness choices for the previous BPT RF hood from simulations. Unluckily, the 0.006” stainless steel (blue) design used in the previous BPT actually maximized the β scattering.


2.5 Design of the BPT Mk IV⁶

From the uncertainty budget in Table 1.2, it can be seen that β scattering is an important experimental uncertainty. A reduction in β -scattering is the major motivation for the redesign of the BPT. In addition, the new design reduces the required RF, leading to lower pickup on the DSSDs.


The design of the BPT Mk IV is shown in Figure 2.5. A detailed comparison of major design elements between the previous BPT and the BPT Mk IV is shown in Figure 2.4. The design goals for the BPT Mk IV were the reduction of both β scattering and RF pickup on the DSSDs [86]. These two goals compete with each other since a reduction in scattering necessitates removing material close to the ion cloud, while reducing the necessary RF amplitude requires moving the electrodes closer to the ion clouds. The new design also includes five electrode segments along the trap axis (“DC electrode regions”) in an effort to reduce disturbing the ion cloud during capture pulses. The previous BPT trap design used three DC electrode regions, leading to some heating of the ion cloud during a capture cycle as the potential at the trap center was perturbed during the capture of additional ions.

The other major redesigned elements are the detector covers and RF hoods (Figure 2.4f), which provide RF shielding of the DSSDs. These are made of 0.002” full-hard stainless steel, which is the thinnest sheet metal found to be sufficiently self-supporting. Compared to the previous electrodes and detector assembly (Figure 2.4), this design has much less material near the trapped ion cloud, reducing the chance of β scattering. The RF hood also has a much shallower angle that minimizes the effective thickness of metal that β particles see, demonstrated in Figure 2.6. These parts were photo-etched then folded by the ANL central machine shop to the final design. Though not perfect, the accuracy of the part is much better than the hand-cut and folded parts of the previous BPT. As part of the overall effort to minimize the amount of material in the vicinity of the ion trap, the stainless steel frame rails that support the electrodes and the detector mounts were also made thinner. The


⁶This section is adapted from Ref. [77].




(a) previous BPT electrodes (beam axis horizontal)




(b) BPT Mk IV electrodes (beam axis horizontal)




(c) previous BPT electrodes (along beam axis)



(d) BPT Mk IV electrodes (along beam axis)



(e) previous BPT detector assembly



(f) BPT Mk IV detector assembly

Figure 2.4: Comparison of major design elements between the previous BPT, as used in Refs. [38, 75], and the BPT Mk IV, performed in Autodesk Inventor [105]. Unlabeled components are made of stainless steel (grey) or alumina ceramic (white). (a) and (b) compare a single quadrant of the electrodes with the beam axis horizontal. The different DC voltage regions in the Mk IV are lettered in yellow. (c) and (d) compare the electrodes along the beam axis; note that the cross section in the BPT Mk IV changes along the beam axis—cf. (b)—and only the rods are present in the center trapping region. The voltage pattern applied the Mk IV is labeled in yellow. (e) and (f) compare a sectioned view of the detector and RF shielding assemblies. See text for additional details.




Figure 2.5: Sectioned view of the final design of the BPT Mk IV performed in Autodesk Inventor [105]. See Figure 2.4 and text for additional details.

endplates have a cylindrical hole and are held at a low DC potential to allow the ions to make the transition from the drift tube of the beamline to the trap (Figure 2.5).

A calibration α source illuminated an entire DSSD at a better-known position. In the previous BPT, the calibration α sources faced detectors across the trapping volume, and two sources each illuminated half of the DSSD with significant overlap in the middle of the detector. Additionally, a quarter of the detector was not illuminated, making calibration difficult on these strips. The calibration sources in the BPT Mk IV are described in more detail in section 2.8.

The RF electrodes were designed so that they did not block the line of sight between the ion cloud and the DSSDs. The ion cloud was assumed to have maximal dimensions of $1.5\text{mm} \times 1.5\text{mm} \times 1.5\text{mm}$, a realistic size from previous traps. This is illustrated in Figure 2.6. Tolerance to allow for a slightly larger size was included in the final design. The optimal angles of the RF hood were also determined from this process, essentially determined to be the angles where the β particles could not see the interior of the hood.

SIMION v8.1 was used to model the electric potential for different electrode designs [107]. A grid spacing of up to 0.05 mm/grid unit was required to accurately determine the potential differences between rods of $\lesssim 1$ mm diameter, though most design work was conducted at 0.2 mm/grid unit. The basic electrode design premise was that the center electrode region, where the ion cloud is held, labeled region “C” in Figure 2.4b, should be formed by wires or rods supported at the ends of the trap. This arrangement entirely removes the support structure near the critical scattering region nearest to the DSSDs (Figure 2.5). As shown in Figure 2.4d, the rods supply the RF potential radially; the DC electrodes provide trapping along the axis. It was found that the DC electrode structure needed to “reinforce” the RF radial potential, hence the pattern of applied voltages shown in Figure 2.4d.

To better approximate the curvature of an ideal hyperbolic electrode shape, tests were conducted with the 2 rods spaced as widely as possible in the allowable electrode region (Figure 2.7a). A 3 rod design spaced in a more hyperbolic shape was also tested. No




Figure 2.6: Design process for the BPT Mk IV to define the region where the electrodes would not obscure line of sight from the ion cloud to the DSSDs. The angle of the new RF hood design is additionally defined by this sketch.

significant difference in the electric potential near the trap center was found between the 2 and 3 rod designs, and scattering favored as few structures as possible. Thin metal wires, such as gold-plated tungsten wires common in time projection chambers, were considered, but due to the engineering challenges involved, a simpler design using larger diameter (~ 1 mm), more rigid rods was favored. Following a similar approach as in Ref. [108], the potential calculated along the radial direction from the center to the electrodes was fit to even polynomial terms of order r^0 to r^6 . The radius at which terms higher than quadratic order contribute more than 2% to the total potential is then shown in Figure 2.7b. Anharmonicities from these higher order contributions lead to greater instability in the ion motion, leading to fewer trapped ions [96]. This treatment provides a metric for determining the stable trapping region for capture before the ions cool to a much smaller radius.

In the energy range of interest (< 17 MeV), the cross section for β scattering increases with target material proton number as Z^2 , so a low Z material is intrinsically desirable [109]. Glassy carbon rods are commercially available with diameters as small as 1.0 mm and were found to contribute a relatively low amount of scattering due to both their low Z and relatively low density of 1.42 g/cm³ [110]. Glassy carbon is non-porous, chemically inert, and electrically conductive, making it an ideal material. The final design selected rods of 1.0 mm diameter at a distance of 15.5 mm from the center with a separation of 5.2 mm (Figure 2.4d), which achieves a similar ion acceptance region as the previous BPT while requiring only 68% of the previous BPT RF amplitude (Figure 2.7c). A key design challenge was the support structure for the glassy carbon rods, as these rods are flexible yet easy to shatter. In addition, the thermal contraction of the stainless steel frame amounts to about 0.8 mm reduction in length across the entire trap frame [111]. However, glassy carbon contracts about an order of magnitude less [110]. To limit the amount of sag in the glassy carbon rod, a stainless steel spring providing ~ 2 pounds of force is used to mount the rods to keep them under tension (Figure 2.8c).

To design the DC electrodes that provide axial confinement, SIMION was used to model

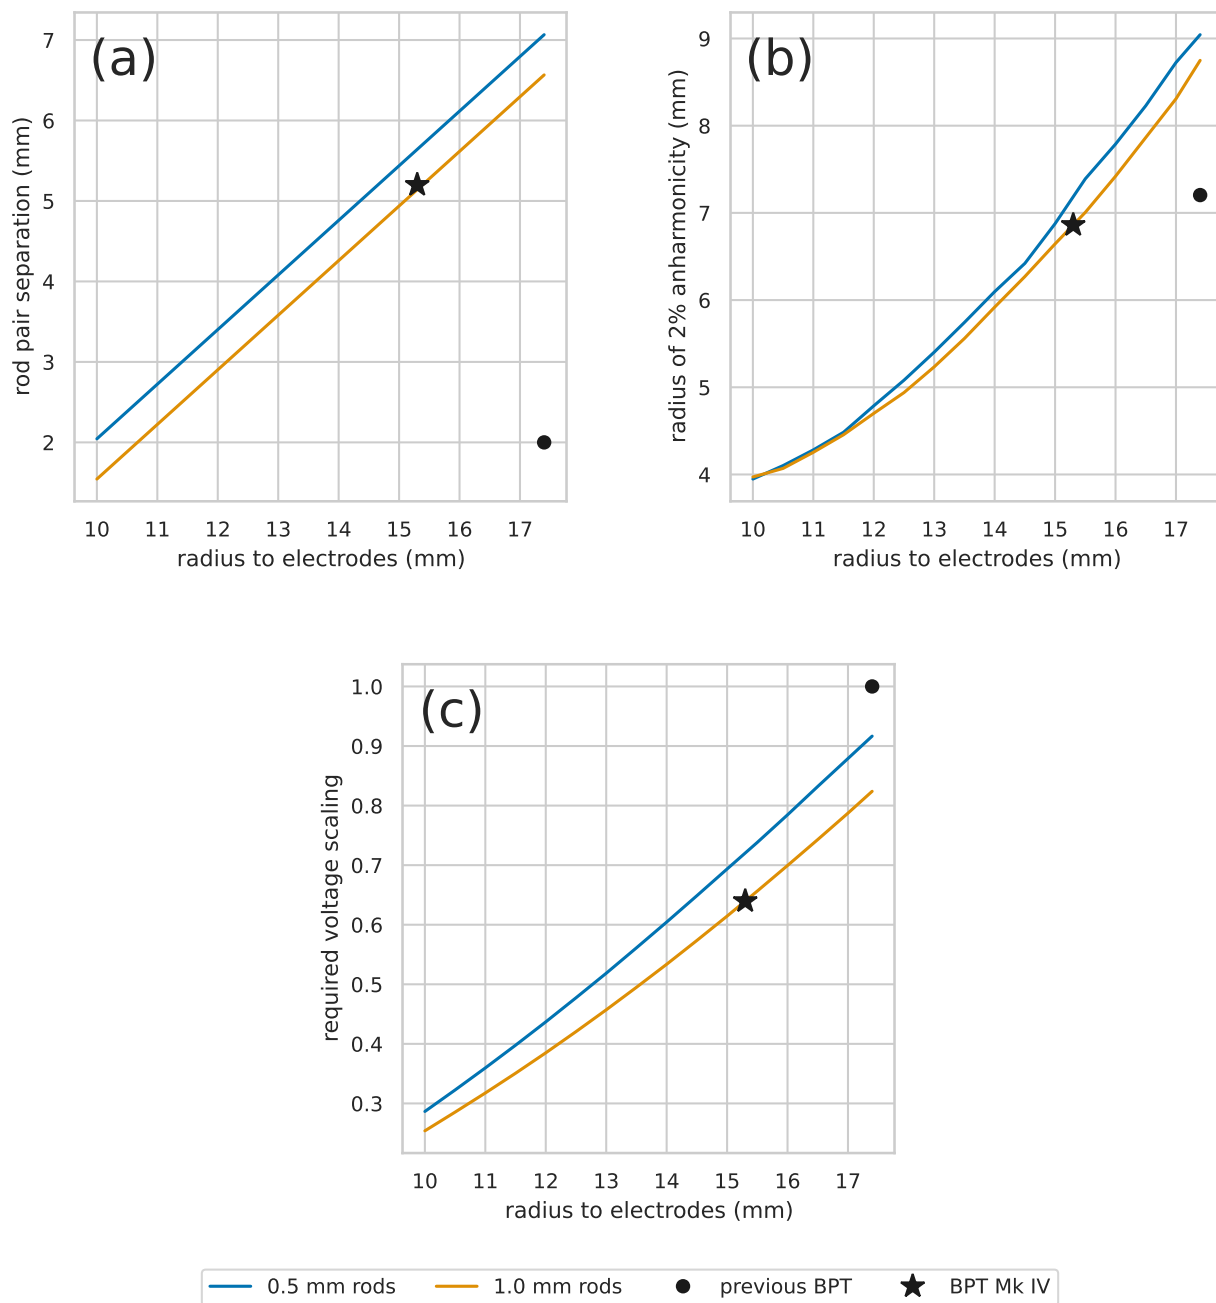


Figure 2.7: Comparison between 0.5 mm rod designs (blue) and 1.0 mm rod (orange) designs; (a) separation between pairs of electrode rods as calculated from the allowed electrode region; (b) radius at which terms higher than quadratic contribute $\geq 2\%$ to the potential; (c) required voltage relative to previous BPT to achieve same radial potential.


thin planar electrodes at a larger radius than the rods. It was found that a stack of three electrodes, with RF+DC on the outer two and DC on the middle, was needed, as shown in Figure 2.4d. This design was then optimized for the length of each electrode region (labeled as shown in Figure 2.4b) to ensure a sufficiently large axial potential gradient in the center of the trap. The minimum length of the center region (“C”) was limited by scattering and the allowable electrode region (previously described). The optimal center region length was determined to be 35 mm with an interior DC electrode (“B” and “D”) length of 40 mm. The interior electrodes “B” and “D” are made of graphite, a much less expensive option than glassy carbon, while the exterior electrodes “A” and “E,” which are further away from the ion cloud, are made of stainless steel, as they have negligible impact on scattering. Some basic tests were performed to evaluate the suitability of graphite for UHV, as described in section 2.5.1.

The BPT Mk IV during assembly is shown in Figure 2.8. All elements are constructed from UHV compatible materials and were thoroughly cleaned before assembly. The final design consisted of more than 500 individual parts. From Geant4 simulations of the final design, 5.3% of triple coincidence events came from scattered β , a roughly $4\times$ reduction from the previous trap.

Thank you to Guy Savard for his help with the installation of the glassy carbon rods; we had a 50% breakage rate during installation but fortunately had just enough spares to complete the trap.

2.5.1 Suitability of graphite for ultra-high vacuum


Graphite is a porous material and has been investigated for use in UHV environments over the decades but its use has been limited due to concerns over its purity, porosity, and contribution to out-gassing [112, 113]. However, graphite also comes in a wide variety of grades with various porosities and purities and is an inexpensive material, even at high purities and relatively low porosities. Of significant advantage to us is that graphite is a



(a) BPT Mk IV during assembly



(b) graphite and glassy carbon electrodes



(c) glassy carbon rod support structure

Figure 2.8: The BPT Mk IV during assembly. (a) Full view of the trap with liquid nitrogen circulation lines shown. (b) A view of the center trapping region. (c) Support structure for the glassy carbon rods, which are held by friction in a small collar. A spring, indicated by a white arrow in lower left, is inserted into each alignment hole to tension the rods.

conductor and that it can be easily machined to very high precision with wire electrical discharge machining.

To test the suitability of graphite for UHV, a sample of grade G535 [114] with surface area of 12 sq. in. (85 sq. in. was required in the final trap design) was tested in a small chamber at both room temperature and at ~ 100 K. The graphite was rinsed with acetone and ethanol before being installed in the test chamber; no particular preparation of the chamber was performed. The graphite sample was attached to the cold finger of a liquid nitrogen dewar that extended into the chamber. It was found that the pressure in the chamber with graphite at room temperature was the same as that of the empty chamber after 64 hours of pumping with a turbomolecular pump, with an ion gauge reading of 1.2×10^{-7} mbar. At cryogenic temperature after 6 hours, the pressure was $\lesssim 2 \times 10^{-10}$ mbar, around the limit of the ion gauge in use. This simple test is not a rigorous measurement of the out-gassing rate of this grade of graphite but provided strong indication that it would be an acceptable material for use in the BPT Mk IV.

Graphite is of interest for other ion traps measuring angular correlations that are sensitive to β scattering (e.g. [115] by private communication). For traps that do not operate at cryogenic temperatures, a more rigorous test of the suitability of graphite of UHV may be required. Note that glassy carbon is well-suited for UHV since it is non-porous, but the material is expensive both in large quantities and to machine. From limited testing, it was found that glassy carbon had a tendency to fracture like a glass but machining with diamond tooling produced acceptable results.

2.5.2 Estimated cloud size in the BPT Mk IV

The final trapping parameters of the BPT Mk IV are used to compute the estimated ion cloud size and compare it to the experimentally observed size using the model presented in section 2.2. From the final trap settings (Table 3.1) and equation 2.4, the radial pseudo-potential has an effective field gradient of 0.042 V/mm² near the trap center and $q = 0.32$.

From the Maxwell-Boltzmann distribution at 80 K, 68% of ions have an energy less than 11.6 meV. Converting this into a maximum distance from the trap center, this corresponds to 0.53 mm, a 1σ radius of the ion cloud (corresponding to a FWHM of 1.25 mm). During the actual experiment (section 3.2), the radial cloud was measured to have $\sigma = 0.4$ mm, about 20% smaller than predicted. In the axial direction, the field gradient is much shallower at 0.007 V/mm². This corresponds to an axial cloud with $\sigma = 1.31$ mm. The axial cloud during the experiment was measured to have $\sigma = 1.48$ mm, about 13% higher than predicted from the simplified model. This simple model is not accurate enough to rely on for the BPT simulations, but it is quite reasonable for design purposes and gives an indication of cloud sizes during tuning.

To facilitate tuning of the BPT Mk IV for future experiments, several plots are provided in Figure 2.9 that indicate the estimated cloud size as a function of different combinations of voltages and frequencies. These are produced by the same treatment, which takes the axial and radial potentials to be independent, which is not a wholly accurate treatment (off by up to $\sim 20\%$ as seen above). However, during tuning, it is extremely useful to have an idea of the relative change in the trap size, which these plots show. In particular, these plots show regions of the parameter space where operation of the trap is likely not possible (in yellow and black), due to a cloud size that is too large or where $q > 0.9$.

2.5.3 RF resonator⁷

A new RF resonator circuit was required for the BPT Mk IV due to the increased number of electrodes, additional electrode regions, and the different electrical properties of the trap. The circuit design is shown in Figure 2.10. A primary *LC* resonator circuit is shown on the left of the diagram, which uses a homemade air core transformer and an air variable capacitor. Five of the modules, indicated by dashed lines, are attached in parallel to the primary resonator circuit to create the independent circuits for each electrode pair. An

⁷This section is adapted from Ref. [77].

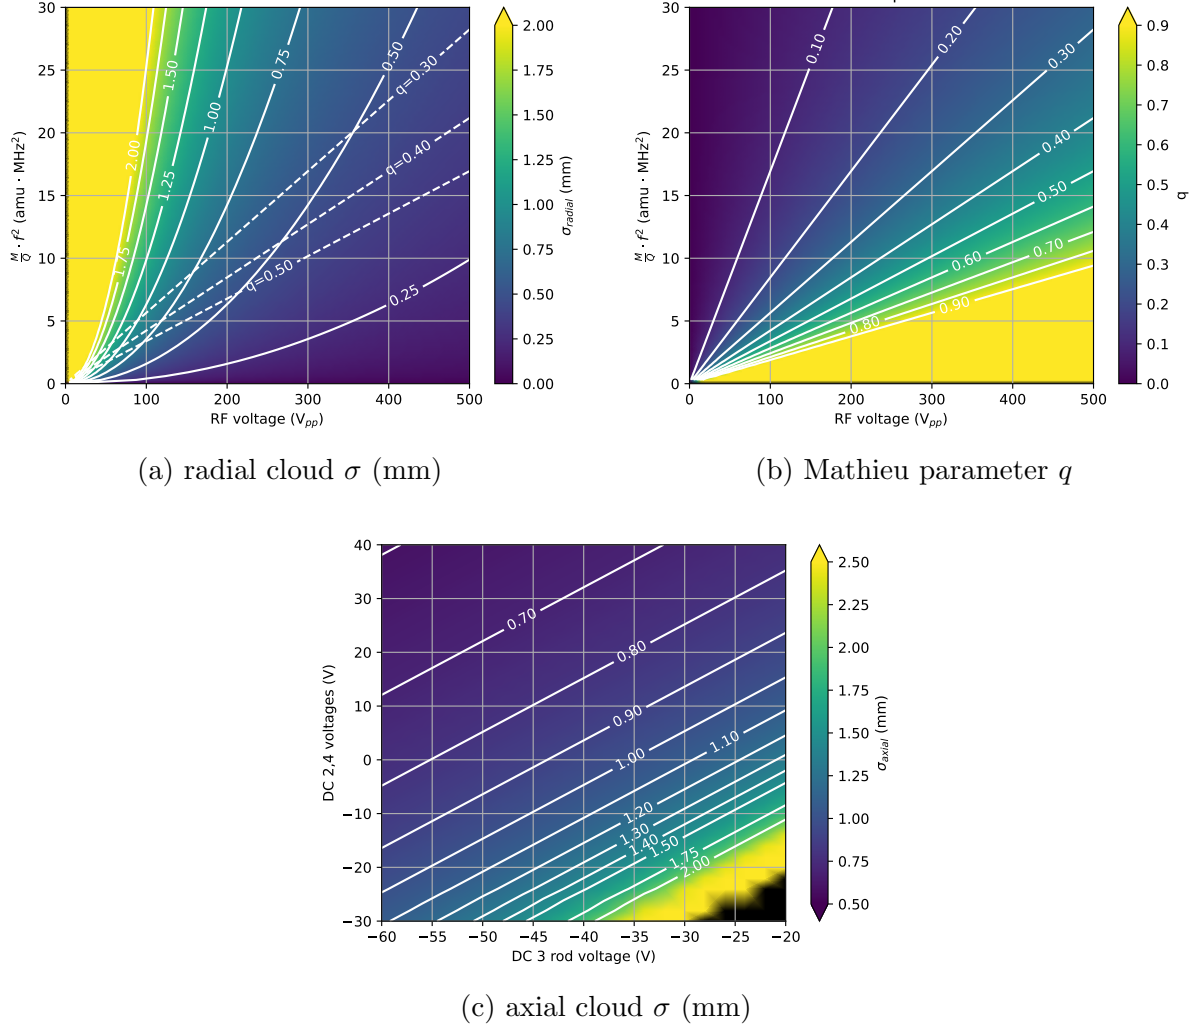


Figure 2.9: Estimated ion cloud sizes from a simplified treatment of the potentials, described in the text. (a) Radial cloud σ as a function of RF voltage (V_{pp}), ion mass to charge ratio, and RF frequency. Dashed lines indicate some corresponding q values. (b) Mathieu parameter q as a function of RF voltage, ion mass to charge ratio, and RF frequency. The vertical axis of (a) and (b) has units of $M/Q \cdot f^2$ in amu \cdot MHz 2 . (c) Axial cloud σ as a function of DC voltages applied to the inner and rod electrodes (see section 2.5).

in-house constructed common mode choke acting as a 1:1 transformer is used to provide a different DC offset to the electrode pairs while allowing each pair to have identical RF amplitudes and phases. The chokes are made using an N49 ferrite [116], which has relatively low losses around the 1 MHz frequency range of interest. Voltages for each electrode pair are taken from the “+” and “-” labeled sources, which have opposite RF phases but identical DC offsets. The DC input has a simple filter to attenuate any RF feedback to the power supply. A pulser is capacitively coupled to provide a fast voltage switch of up to a few hundred volts with ~ 10 ns rise time; these pulses are used to lower the electrode voltages during capture and ejection of the ions.




Figure 2.10: Circuit diagram of the resonator for the BPT Mk IV. See text for details.

The inductor on the primary resonator circuit is relatively low-valued due to a high trap capacitance of roughly ~ 600 pF. This is likely due to the fact that thin Kapton wires with a coaxial ground shield were used to deliver the voltages to the trap in an attempt to reduce pickup inside of the chamber; these have a stated capacitance of 137 pF/m [117]. To allow the circuit to achieve higher frequencies, an inductor can be added in parallel with the primary resonator circuit with a switch (not shown in Figure 2.10). The resonator circuit is tuned to the appropriate frequency with the variable capacitor. The resonator is driven by an RF amplifier (model T&C ULTRA 2020), in turn driven by an SRS DS345 function generator [118] providing a sine-wave of the desired frequency.

The quality factor of the resonator circuit is fairly low, with $Q \sim 10$. Measurements of

the RF voltage and phase on each electrode pair are difficult to perform accurately since attaching e.g. an oscilloscope to the electrodes slightly changes the properties of the circuit. With this in mind, the electrodes have amplitudes within $\lesssim 1\%$ of each other and no phase mismatch was noticeable.

Thank you to Michael Oberling and Daniel P. Burdette who provided significant help in the construction of the RF resonator.

2.6 Detectors

The BPT uses 4 DSSDs to detect the position and energy of the two back-to-back α particles and the position of the β . The DSSDs are 1 mm thick Micron Semiconductor Ltd. BB7(DS) 7P/2M detectors with 100 nm junction side windows and 2 mm strip pitch [119]. Each detector has 32 front (junction) strips and 32 back (ohmic) strips, for a total of 256 channels. The DSSDs are held at ~ 80 K during the experiment and fully depleted with a bias voltage of +180 V applied to the back strips. The cryogenic cooling provides for lower leakage currents, giving improved resolution. Notch filters are attached to the front strips (section 2.6.2) to reduce the RF pickup on each strip. The DSSDs have a typical α -particle resolution of about 20 keV FWHM. The minimum-ionizing β particles deposit a few hundred keV in the 1 mm thickness of the DSSD. An extensive description and full characterization of these detectors is given in Chapter 4.

A picture of a DSSD can be seen in Figure 4.10. These DSSDs are mounted to the frame of the BPT at a distance of 6.5 cm from the trap center. This arrangement provides about 25% solid angle coverage with a $\sim 2^\circ$ angular resolution per pixel. Note that the front strips run horizontally along the beam axis, and the back strips are perpendicular to the beam axis. The front strips are therefore sensitive to the radial size of the ion cloud, while the back strips are sensitive to the axial extent of the ion cloud.

A plastic scintillator with a photomultiplier tube is mounted behind each DSSD (four scintillators total) to detect the remaining β energy. These organic plastic scintillators are

made of polyvinyltoluene by Eljen Technology, model EJ-200 [120]. The scintillators are 6.2” in diameter and 6” in length. The rise time is 0.9 ns with a decay time of 2.1 ns. The scintillators are coupled to Hamamatsu R877-100 photomultiplier tubes [121] operated at -1300 V. These scintillators have been studied in detail in Ref. [83]. Due to the high outgassing rate of the plastic scintillators, they are separated from the trap vacuum chamber by a thin window. The windows are made of Kapton polyamide film (0.008 mm thickness) evaporatively coated on both sides with a 100 $\mu\text{g}/\text{cm}^2$ layer of aluminum [122]. The plastic detectors are then held at $\sim 1 \times 10^{-3}$ mbar while the trap volume is maintained at a UHV level of $\sim 1 \times 10^{-8}$ mbar, prior to the addition of helium buffer gas.

2.6.1 RF pickup

A major challenge for the BPT has been RF pickup on the DSSDs: placing charge-sensitive detectors next to oscillating high voltage is, in general, not a good idea. An oscillating baseline is induced on the DSSDs due to their close proximity to the RF electrodes. This RF pickup can be much larger than the signal induced by an α particle and can swamp the preamplifiers, destroying the resolution of the α energy. Examples of spectra with high RF pick up can be seen in Refs. [41, 86]. Figure 2.11 shows the effect of RF on a pulse from an α at the preamp output; RF mitigation including notch filters (section 2.6.2) was in place, demonstrating the considerable pickup still present. This waveform was taken from the previous BPT [86] at 470 kHz, the trapping frequency for the 2019 ^8B data set [75].

In the course of this work, the RF pickup was directly measured across a DSSD; I appreciate the aid of Daniel P. Burdette with these measurements. A phase and amplitude pattern was observed, shown in Figure 2.13, that indicated that the majority of the effect comes from the electrodes directly inducing a voltage onto the DSSDs. The amplitude and phase dependence comes from the fact that the front strips are parallel to the length of the electrodes (both parallel to the beam axis). Therefore, front strips closer to one phase of the electrodes will have a higher contribution from that phase. An illustration of this effect




Figure 2.11: Waveforms from the output of a preamplifier showing the size of RF pickup relative to an α pulse. RF mitigation including notch filters were in place.




Figure 2.12: Illustration of the phase of RF pickup across the DSSD surface. Front strips are parallel with the electrodes and the beam axis. Electrodes shown are from the previous BPT.

is shown in Figure 2.12. The middle front strip of the DSSD should therefore have very little amplitude, since the contributions of each phase of the electrodes will cancel: this is observed nearly perfectly. The off-set from the middle is likely due to the fact that, in the split electrode design, there is a metal cover on the ceramic that is also at RF voltage (shown in Figure 2.12) which causes an asymmetry in the field. The effect of the RF shielding can also be directly seen in Figure 2.13. Front strips near the edges that were shielded from the electrodes had much lower pickup.




Figure 2.13: Measured RF amplitude and phase from the output of a preamplifier on the front strips of a DSSD.

Back strips were observed to have very little, if any, RF pickup. This is likely due both to the fact that they are perpendicular to the electrodes, resulting in a cancellation of the pickup, and that they are shielded to some degree by the front strips. Given this striking pattern of RF pickup, a logical approach would be to rotate the DSSDs by 90° in order to make the front strips perpendicular to the electrodes instead. Unfortunately, the geometry of the ceramic detector mounts (Figure 4.10) and other trap components prevents this. To incorporate a rotation of the DSSDs into the BPT Mk IV design likely would have required a new vacuum chamber and/or extensive overhaul of the mounting system of the plastic scintillator detectors (section 2.6), which was beyond the scope of this work.

It is possible that there are additional contributions to the RF pickup from wires carrying RF voltages to the trap. In an effort to reduce this impact, the BPT Mk IV uses coaxial wires with grounded shields to deliver RF voltages to the trap [117].

2.6.2 Notch filters

In order to avoid saturating the input of the Mesytec MPR-32 preamplifiers [123] due to RF pickup, notch filters are added between the output of each DSSD front strip and the input of the preamplifiers. The circuit diagram of the notch filter is shown in Figure 2.14. The filters have a trimmer capacitor for fine tuning of the filtered frequency; the trimmer capacitor is quite large compared to the fixed capacitor to allow for a broad range of frequencies during the BPT trap tuning. This notch filter design is slightly different than the previous BPT [41, 86]. Here, the inductor and capacitors are attached prior to the resistor; in the previous version, they were attached after (cf. Figure 6 of Ref. [86]). This subtle change in design actually results in better filtering because the DSSD strip acts as a current source and the preamplifier has low impedance.




Figure 2.14: Circuit diagram of notch filter; notch filters are only used with the front strips.

The notch filters are in sets of 8 on a separate printed circuit board (PCB) that attaches to feed out PCBs on the vacuum chamber. A total of 16 filter sets are required to cover the 128 front strip channels.

At the beginning of the ^8Li run (chapter 3), a range of trapping frequencies were tested for their trapping efficiency and the amount of RF pickup induced on the DSSDs. During this time, the trimmer capacitors were tuned to match the trapping frequency using a ceramic

screwdriver and with the detectors at their full bias voltage. Because the capacitance of the DSSD changes when depleted, it is important to tune the notch filters with full bias. Once the final trapping frequency of 1400 kHz was selected, the trimmer capacitors were tuned a final time for the rest of the data collection.

During the course of the commissioning runs (chapter 3), a number of different filter configurations were tested with various values of resistors and inductors. Michael Oberling and Edward J. Boron, III, were indispensable in assisting with the configuration and re-configuration of the filters on very short timescales. This experiment would not have been possible without them, and I sincerely appreciate their help.

2.7 Data acquisition system

The BPT uses analog electronics for its data acquisition, with both NIM and CAMAC modules. With 256 DSSD channels, upgrading to a digital data acquisition system (DAQ) system is not an inexpensive upgrade and has not yet been necessary, though it may offer some benefits, as described in section 5.2. A schematic of key features of the data acquisition chain related to the DSSDs is shown in Figure 2.15. This figure does not show all of the elements but provides an overview. The data acquisition system used is called SCARLET [124], written and maintained by Ken Teh. SCARLET is an in-house, Linux-based data acquisition system that supports multiple CAMAC crates. Events are written to a hex format and can be later sorted into a ROOT format for further processing and analysis [125].

For the DSSDs, Mesytec MPR-32 differential output preamplifiers are used [123]. A total of 8 preamplifiers are needed for the 128 front strips and 128 back strips. An example of the output of the preamplifiers with RF can be seen in Figure 2.11. RF is therefore still a serious problem even with the tuned notch filters (section 2.6.2).

The differential output of the preamplifiers are fed to the Mesytec MSCF-16 F shaping amplifiers (“shapers”) [126] with twisted pair ribbon cables. Sixteen shapers (model number

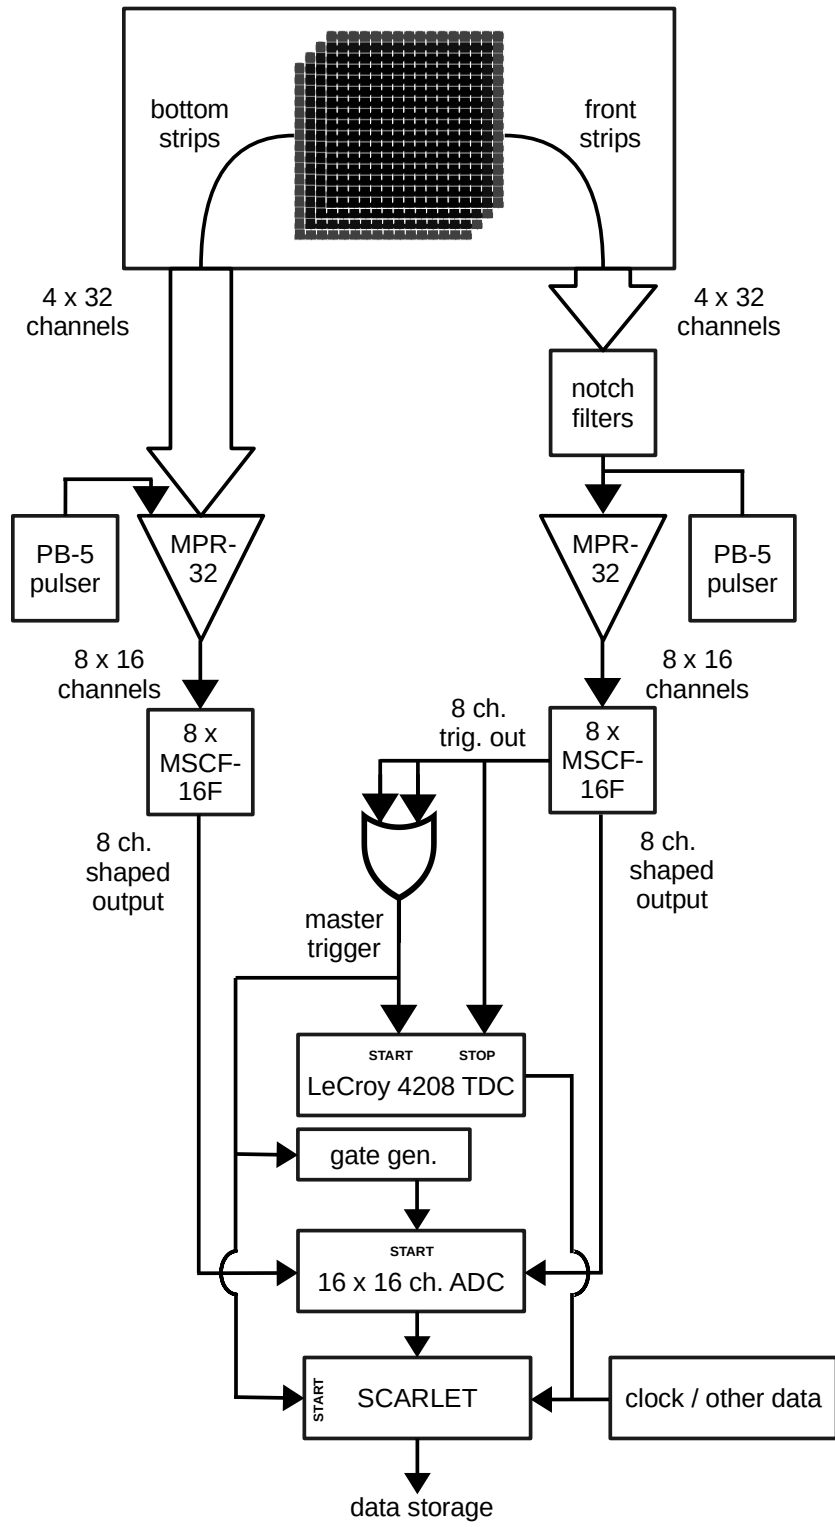


Figure 2.15: Schematic of key features of the data acquisition chain from the DSSDs. See text for details.

MSCF-16-F-Sh8-70-V-D) are used. These were ordered in two different sets of 8 during the experiment development around 2012. It was learned during the 2022 ^8Li experiment that these different sets have two different firmware versions, but this does not seem to have an impact.

A relatively long shaping time of 4 μs is used to integrate out the effect of the RF on the shaped signal. This can be seen in Figure 2.16. Typically, shaping amplifiers split an incoming signal between a fast trigger and a slower shaped energy signal. Instead, these shaping amplifiers have a custom chip to reduce the impact of RF on the trigger. This chip allows the shaping amplifiers to use a leading edge discriminator to trigger on the shaped signal, effectively removing the false triggers on the RF [127]. In a typical BPT experiment, trigger rates are well below 1 kHz total across all 128 front strips, allowing for a long shaping time without any event pile-ups.

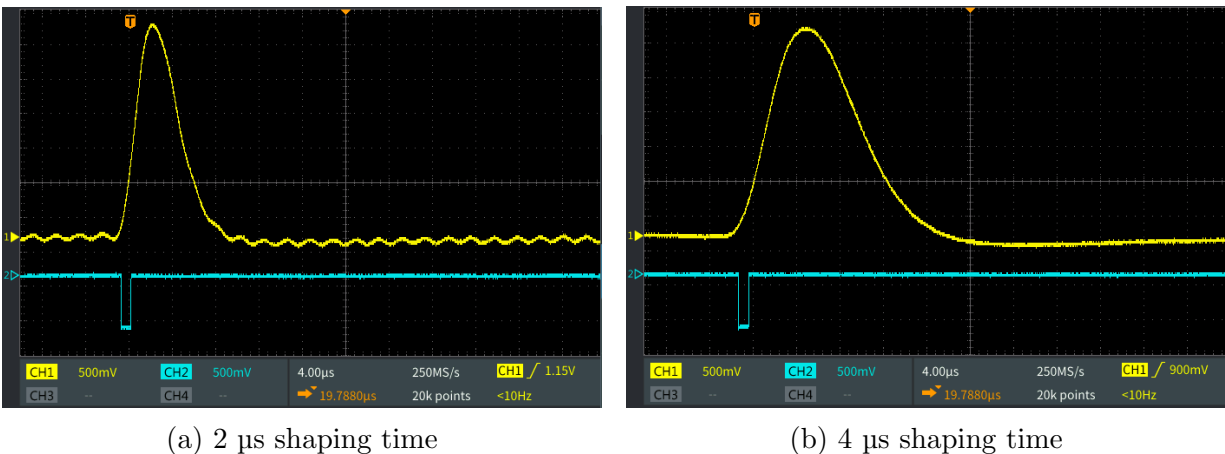


Figure 2.16: Shaping amplifier output with a shaping time of (a) 2 μs and (b) 4 μs , showing that the RF pickup is effectively integrated out at 4 μs . Images are from an oscilloscope showing energy output (yellow) and trigger output (blue).

Triggers from the front strips are combined into a master trigger (Figure 2.15). Back strip triggers are not used. The master trigger is used to generate a gate for the peak-sensing ADCs (model Phillips 7164H [128]) which measure the pulse height of the shaped output. The master trigger also starts an event clock to provide timing information about each front strip trigger, which can be used for event analysis later. For additional details about the

timing information of the events, see section 3.4.2.

Finally, the master trigger is used to start recording a SCARLET event. Data recorded in each event includes the front and back strip energies, the relative front strip trigger times, and some other timing and information about the current trap cycle, as well as data from the plastic scintillators.

One difference in the DAQ for the ^8Li data campaign (chapter 3) is the use of a baseline restoration setting on the shaping amplifiers. During initial setup and testing, it was found that the RF pickup was still quite large on the detectors. Baseline restoration can help to remove low-frequency signals (the RF is at a lower frequency than the rise time of a pulse), and this setting was found to help reduce the noise in the system. It was observed during testing that this setting also seemed to cause a non-linearity in the DAQ response, however, and so a relatively low baseline restoration strength was used (a setting of “5”). As described in section 3.4.3, a precision pulser was used to test the linearity of the system with this setting enabled in order to correct for any apparent non-linearities.

2.8 Calibration sources in the BPT Mk IV

The BPT uses in situ α sources for calibration; these are ^{148}Gd and ^{244}Cm sources and are described and characterized in sections 4.1 and 4.3. The α particles from these sources are easily removed from the final spectrum of an experiment by timing coincidence cuts and by cuts on energy and position of 2α coincidences. This section describes some specific of positioning and mounting these sources in the BPT Mk IV.

The location of these sources in the BPT Mk IV is quite different than in the previous iteration of the BPT. In the previous trap, two of each source (^{148}Gd and ^{244}Cm) were located on the RF shield opposite of the illuminated DSSD, as shown in Figure 2.17. This positioning allowed the α particles to enter the DSSD almost normal to the detector surface, with up to an 18% increase in the effective dead layer at the largest angles. Each set of sources covered about half of the detector surface. However, several downsides were present. A major issue

was that 8 back strips on each DSSD (1/4 of the total surface area) were not illuminated at all due to shadowing from the RF hoods, shown in Figure 2.18. This made calibration of these back strips difficult. In addition, the fact that two sources illuminated the middle of the detector made measuring the precise position of the sources nearly impossible. This was further complicated by the fact that the RF shields and hoods were cut and folded by hand and not machined with any precision. The exact position of the shadowing of each source was unpredictable. A great deal of time and effort was spent trying to measure the source positions from the data without success.




Figure 2.17: Calibration α source locations in the previous BPT. Reproduced from Ref. [86].

For the BPT Mk IV, the location of the calibration sources was changed. As shown in Figure 2.4f, the sources are instead mounted directly into the redesigned RF hood of the detector they illuminate. This allows the entire DSSD to be illuminated by a single calibration source, dramatically reducing the complexity of determining the source position. The sources were also re-mounted in a more precise source holder, shown in Figure 2.19. The 0.002" full-hard stainless steel plates that sandwich the source foils were made by photo-etching, reducing the uncertainty on source positions from the previous BPT mounts, which were handmade.

One disadvantage is that the α particles from the sources now enter the DSSD at much




Figure 2.18: Example of ^{148}Gd source shadowing in the previous BPT. Two sources each illuminated half of the detector, hence the higher counts in the middle of the detector. Front strip shadowing from the RF hood, indicated by the red arrows, was unpredictable due to their hand-made construction. Back strip shadowing, indicated by the blue arrow, made calibration difficult on these strips in previous data sets.




Figure 2.19: Source holder for the BPT Mk IV. Source foils (one ^{148}Gd , one ^{244}Cm) are sandwiched between two 0.002" thick stainless steel pieces, held together by M1.2 screws. The 1.5 mm diameter holes in the plates allow the source foils to see the detector. See Figure 2.4f for the mounting position.

shallower angles, leading to an increase in the effective dead layer by up to a factor of $\times 2.5$, resulting in up to about 20 keV additional energy loss. Additionally, some of the α particles exit the source at relatively shallow angles, which contributes to energy loss on the order of a few keV. From the characterization of the source dead layer (section 4.3) and the detector response (chapter 4), these effects can be modeled to produce an accurate calibration.

The count rate distribution across the DSSD can be used to fit the source position; this distribution is shown in Figure 2.20. This distribution clearly shows the expected $1/r^2$ dependence. It is not clear where the two spurious peaks at the front come from, but these are present on almost every detector/source combination and in almost the exact same position. It was not possible to remove these peaks by noise cuts or energy cuts, suggesting they are not from RF pickup or noise. It is considered unlikely that these are from the calibration sources of other DSSDs from geometry constraints. It is possible that, at these very shallow angles, some additional parts of the source foil are viewable by these pixels underneath the layer of source holder, leading to a higher rate. This is very difficult to model and is not attempted.

Using only the count rate, it is possible to fit the position of the sources (assuming a point source for simplicity). This results in fitted source positions that agree in two dimensions (front strip direction, back strip direction) within about 2-3 mm of the nominal source positions. The height of the source (distance between plane of DSSD and the source) disagrees by about 1 cm, which is not realistic, and so the nominal height is taken. In performing these fits, the parts of the DSSD with the spurious peaks are removed from the fits.




Figure 2.20: Distribution of counts across one of the DSSDs from its ^{148}Gd source. Front strip 24 is dead on this detector, hence no counts. It is not clear where the two small peaks near the front come from. See text for details.

Chapter 3

^8Li DATA CAMPAIGN

The ^8Li data campaign took place at Argonne National Laboratory (ANL) using the Argonne Tandem Linear Accelerator System (ATLAS) under experiment #1957 [129] from May 10, 2022, to June 7, 2022.

Two ^8Li commissioning runs were attempted in March and April, 2022. However, these runs ultimately did not result in any ^8Li being delivered to the BPT Mk IV. It was found after these runs that a degrader used in ^8B production had slipped into the path of the beam, resulting in the wrong energy being delivered to the gas catcher so that most of the beam was stopped near to the front of the device and not efficiently extracted.

These experiments, especially the month-long data campaign, were absolutely exhausting and would not have been possible without the aid of a team of people. In particular, I would like to thank Daniel P. Burdette, Jason A. Clark, and Guy Savard for their significant help over several months in setting up and tuning the entire transport system. Special thanks also to Aaron T. Gallant and Brenden Longfellow for their help with setting up, especially with the notch filters and the DAQ. Finally, thanks to all of my collaborators who helped with shifts and other responsibilities; see the acknowledgements for a complete list.

3.1 ^8Li production, capture, and transport

A facility map of ATLAS is shown in Figure 3.1. The BPT is located in the Triangle Room, labeled as “Trap Area.” The gas target, focusing magnet, gas catcher, and other elements are located in Area II. Due to its location, this experiment does not use the full accelerator and thus cannot reach the maximum beam energy of the facility. However, the first few accelerating segments of the linac provide the necessary energy to produce ^8Li (and ^8B). The ^8Li production, capture, and transport systems are unchanged with respect to the previous trap and experiments, with details found in Refs. [41, 78, 80, 87]. Only key features

are highlighted here.¹ Along the beamline are several different silicon detectors and micro-channel plate (MCP) detectors used for diagnostics and tuning; these are not described. Additional details of the cycling timing for the various elements is provided in section 3.3. For most of the run, the beam current as measured on the beamstop was 20-30 enA.

ATLAS ARGONNE TANDEM LINAC ACCELERATOR SYSTEM

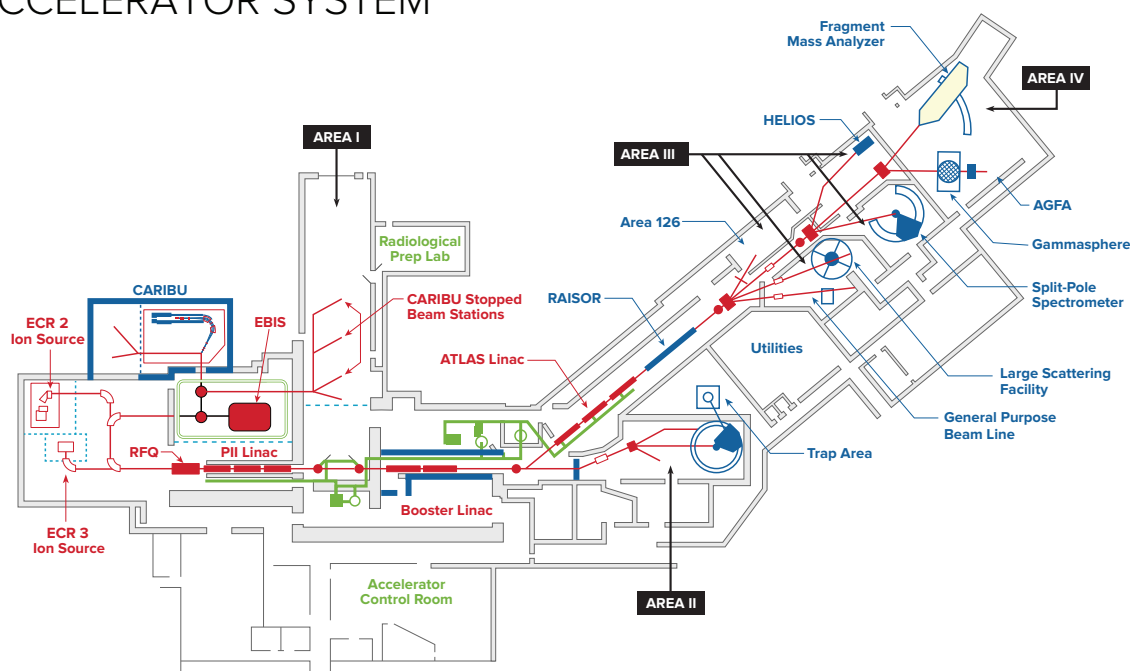


Figure 3.1: ATLAS facility layout. The BPT is located in the Triangle Room, labeled as “Trap Area.” Other production elements are located in Area II. From Ref. [130].

The experiment used the ${}^7\text{Li}(d,p){}^8\text{Li}$ reaction with a 22 MeV ${}^7\text{Li}^{2+}$ primary beam on a cryogenic deuterium gas (D_2) target. The ${}^7\text{Li}$ ions were produced via heating a stable solid source in a commercial oven.² The beam of ${}^7\text{Li}$ is then accelerated to 22 MeV towards the

¹There are too many tuning parameters for the system for the specific voltages and frequencies to be reported here, and these were changed throughout the run to continuously improve the extraction efficiency.

²On the third day of the ${}^8\text{Li}$ data campaign, it was learned that the ion source might have possible actinium contamination from a previous experiment. After the experiment was completed, it was determined that no contamination was present. However, due to safety restrictions imposed in the interim, it was not possible to

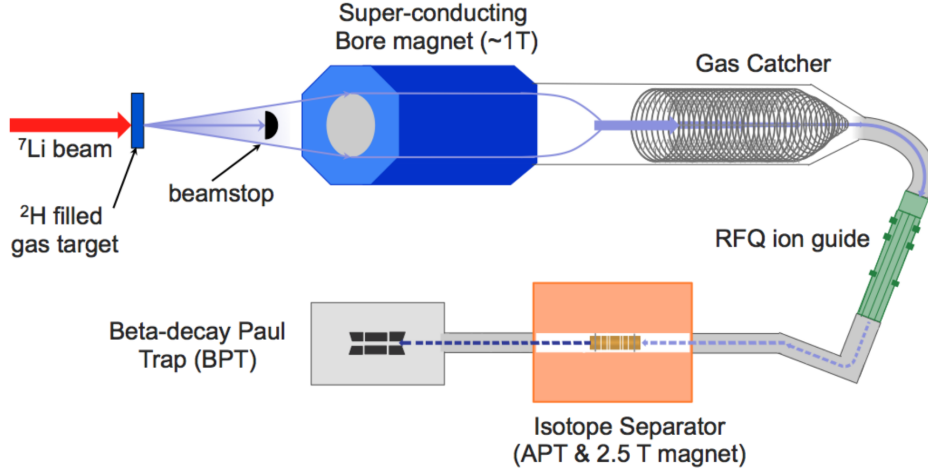


Figure 3.2: From Ref. [41]: An overview of the beamline used to create and transfer ^8Li to the BPT.

gas target. The gas target is 4.5 cm long and holds about 550 torr of deuterium at ~ 80 K. Liquid nitrogen is circulated through the gas target frame using a custom pump [101]. The entrance window of the gas target is a 1.3 mg/cm^2 titanium window; the exit window is 1.8 mg/cm^2 titanium.

The $^7\text{Li}(d,p)^8\text{Li}$ reaction has a cross section of around 75 mb [131]. Most of the unreacted primary beam that scatters less than $\sim 5^\circ$ is stopped by a tantalum beamstop downstream of the gas target. Scattered primary beam and ^8Li reaction products are focused and separated via a large bore, 1 T magnet (Figure 3.2). Backwards center-of-mass angles are selected, which reduces the cross section by an order of magnitude [132], but allows for better separation between the ^8Li and the scattered primary beam.

The resultant beam then impinges onto the gas catcher [133]. The gas catcher (Figure 3.2) uses a 3.3 mg/cm^2 aluminum window to contain about 50 torr of ultra-high purity 99.999% helium gas. This helium gas stops the ^8Li products inside of the gas catcher, producing He ions and various other molecules. Removing as much primary beam as possible is required

open the oven to replenish the lithium supply. Therefore, the experiment ran only with the limited supply of lithium initially installed, making it very uncertain when the ion supply would run out. This led us to tune the system as quickly as possible and establish the final trapping parameters of the BPT quickly. This resulted in less tuning of the BPT than may have been optimal, though the resulting ion cloud was still of a reasonable size. Though it was thought that the initial supply of lithium might last as little as a week, it ended up lasting for the entire month of the experiment.

to avoid saturating the gas catcher with charge, which leads to space charge effects that reduce the extraction efficiency of the ^8Li products. The helium gas thermalizes the ions, while RF voltage (frequency around 5 MHz) inside the gas catcher pushes ions away from the walls and focuses the ions towards the exit nozzle into a continuous, low-energy beam. A DC gradient accelerates ions towards the exit. The helium gas flow itself is key to the efficient extraction of ions. Helium is used for several reasons, an important one being that it has the highest ionization potential of any gas, which limits charge exchange with ions entering the gas catcher.

Due to the low mass of the ^8Li , the extraction efficiency is only on the order of 0.1% - 1%. The efficiency also depends highly on the amount of primary beam entering the gas catcher. In the data campaign, as in previous experiments, it is clear that the gas catcher saturates even at relatively low beam currents.³ Additionally, most of the ^8Li nuclei that exit the gas catcher are incorporated into various molecules that form inside the gas catcher. These molecules have been found to be mostly around mass 25 u (likely $^8\text{LiOH}^+$) and mass ~ 40 u (perhaps $^8\text{LiOH}\cdot\text{H}_2\text{O}^+$). Gas cell chemistry is not well-understood and is obviously complex due to the amount of energy and charge in the cell (plus the atypical species often used in other experiments). Both of these molecules originate from water contamination inside of the gas catcher; some water molecules are required to extract ^8Li nuclei, but too much contamination quickly results in inefficient extraction. To limit the amount of water in the system, the ultra-pure helium gas passes through two charcoal-filtration cold traps held in liquid nitrogen baths to freeze out water and other contaminants from the gas flow. A commercial getter further purifies the gas flow. Finally, the gas catcher can be chilled with a refrigerant fluid system to around -20 °C to further freeze out water. The temperature of the gas catcher was adjusted during the data campaign to add and remove water contamination as necessary. The gas catcher is also contained in its own vacuum chamber, held at around 1 mbar, which prevents any leaks to atmosphere from the gas catcher itself.

³This is one reason that the lithium source was able to last so long in this experiment.

In the extraction of ^8B , it is found that a high level of water contamination is essential for efficient extraction of molecules. In ^8B experiments, therefore, several of these water reduction strategies are not used. The gas catcher itself is typically vented to atmosphere and pumped down immediately before a run and can be heated above room temperature to drive water off into the helium gas. Future experiments are considering adding water directly to the helium gas flow through the use of a nafion gas-exchange membrane [134], which can selectively humidify a gas stream.

The continuous beam of low-energy ^8Li products exits the gas catcher and is directed towards a radio frequency quadrupole (RFQ), illustrated in Figure 3.2. The RFQ is yet another helium-filled, RF and DC voltage device and therefore a type of Paul trap. The RFQ is held at room temperature and the ions thermalize to this temperature through collisions with the ~ 1 mbar of helium gas. Split into several different segments, the RFQ focuses the ions exiting the gas catcher and traps them via a DC gradient into a “bunch” in the final segment.⁴ In this experiment, the RFQ was operated without a mass-selectivity DC offset (section 2.2), but some mass selectivity does occur from the applied RF frequency, about 2 MHz.

Every 40 ms (25 Hz), the cooled and bunched ions are ejected from the RFQ downstream at a voltage of 1490 V. Directly after the RFQ are two electrostatic deflectors that can be used to perform mass selection of the ion bunch. These deflectors are separated by roughly a meter of beamline, allowing for ions of different masses to achieve a few μs of time separation. Pulsing the deflectors at the right time allows only masses in a small range to enter the next beamline element, the Advanced Penning Trap, operating as an isotope separator (IS). In the data campaign, ions around 25 u were selected ($^8\text{LiOH}^+$).

A Penning trap is another type of ion trap that uses a magnetic field to confine ions instead of purely RF and DC voltages. The physics of a Penning trap is not described here but can be found in a number of publications and previous theses, e.g. [41, 80, 87, 135,

⁴Since the ions are cooled and bunched, the device is often called a RFQ cooler-buncher.

[136]. The IS operates under a magnetic field of 2.5 T and contains about 4×10^{-6} mbar of helium gas. By applying a dipole excitation pulse at ω_- to the ions, energy is added to the ${}^8\text{LiOH}^+$, which then disassociates by collision with the helium gas, forming pure ${}^8\text{Li}^+$. The disassociation energy is about 1 eV [137]. Then, by applying a quadrupole excitation at the ω_c of ${}^8\text{Li}^+$, these ions can be driven towards the center of the Penning trap, while the remaining masses are driven away. This converts the ${}^8\text{LiOH}^+$ molecules to ${}^8\text{Li}^+$ with an efficiency of close to 100%.⁵ To allow for a measurement cycle in the BPT, 4 ion bunches from the RFQ are collected in the IS at a time, before ejection to the BPT every 160 ms.

3.2 To the BPT Mk IV

The collected ion bunches are ejected from the IS and captured by the BPT, their final destination and where the measurement occurs. This ${}^8\text{Li}$ data campaign was the first experiment to use the BPT Mk IV, which performed successfully and without any major surprises (though the blinded data analysis is on-going). Figure 3.3 shows the BPT installed in its vacuum chamber, with key components labeled.

3.2.1 Off-line commissioning⁶

Offline testing of the BPT Mk IV was performed using argon and nitrogen gases in spring 2022, which have masses of 40 u and 28 u respectively. Ions were produced by introducing these gasses close to the RFQ cooler-buncher and then ionizing them with a cold cathode gauge. This ionization and trapping process has a low efficiency, but plenty of ions were produced for commissioning. Trapping voltages and frequencies for the BPT Mk IV were found to be comparable to simulations, but only a limited amount of trap tuning is possible without a radioactive beam.

⁵Unfortunately, this technique does not work for the break up of ${}^8\text{B}$ molecules, which seem to have a much larger disassociation energy. These ions are typically trapped in the BPT at mass 42 u, probably ${}^8\text{B}(\text{OH})_2^+$. Due to the extremely large Q value for the β^+ decay of 16.9 MeV and the subsequent 2α break up, there should be essentially no impact on the decay products from the molecule.

⁶This section is adapted from Ref. [77].




Figure 3.3: The BPT Mk IV installed in its vacuum chamber, with major components labeled.

Tuning the BPT Mk IV requires a slightly different technique than tuning the previous BPT or other typical beamline elements. Since the glassy carbon rods are closer to the ions than the DC electrodes, there is a voltage-screening effect, necessitating a higher applied DC voltage on the electrodes to form a trapping potential on the axis. This also means that the capture and ejection pulses applied to the DC electrodes must be relatively high (roughly -100 V). Therefore, it tends to be important to consider the relative difference between the rod DC voltage and the other voltages. Using the estimates shown in Figure 2.9 can make tuning much easier.

3.2.2 On-line commissioning⁷

As mentioned at the beginning of this chapter, two on-line commissioning runs with ^8Li were planned but were not successful in delivering beam to the BPT Mk IV. Therefore, the first time that the BPT Mk IV was tuned for ^8Li was during the data campaign. After

⁷This section is adapted from Ref. [77].

tuning the entire system for roughly a week, the BPT voltage settings were finalized and held constant during the rest of the data collection. As indicated at the beginning of the chapter, it was unclear how long the data campaign might actually last, so this length of tuning time was a compromise between a small ion cloud and the prospects of getting any useful data. Data was collected for a total of 377 hours, resulting in around 2.6×10^6 triple coincidence events.

A summary of the final trap settings is given in Table 3.1, with the resulting axial potential shown in Figure 3.4, as determined from SIMION using the applied voltages from the experiment. The applied RF voltage is known only to about 10%, due to the fact that probing the circuit changes its properties and that the probes used may not have been impedance matched to the oscilloscope properly. There was no measurable phase offset between the two phases of the electrodes; the two phases also had identical amplitudes to within $\sim 1\%$.

The helium buffer gas pressure was $\sim 1.2 \times 10^{-5}$ mbar as read by a cold cathode gauge in the trap vacuum chamber. Note that the pressure inside the BPT itself may have been up to an order of magnitude higher due to a lower pumping speed in the confined space of the trap. The capture and trapping efficiency was estimated to be roughly $\sim 25\%$; the efficiency likely could have been improved through additional tuning. The typical trap population was estimated to be a few hundred ^8Li ions. The trapping lifetime was not measured directly but was many tens of seconds, as shown in section 3.3, long enough that it was not a concern for the experiment.

During capture, only electrode “A” is pulsed down to allow ions into the trap, but this pulse causes a large amount of noise pick-up on the DSSDs. Data from the first 6 – 8 ms of each trapping cycle are unusable. It is not well understood why this time period of unusable data is so long compared to the pulse length of a few microseconds. This is not a problem for the final data analysis, as this is during the initial ion cooling period of roughly 30 ms and therefore would be discarded anyway.

electrode	DC (V)	pulse (V)
“A”	+65	-91
“B”	-15	
“C” (rods)	-35	
“D”	-15	-91
“E”	+66	-100
RF	1400 kHz, 300 V _{pp}	

Table 3.1: Voltage settings for the BPT Mk IV during the commissioning ^8Li run. Electrodes are labeled as in Figure 2.4b. Electrode “A” is pulsed during capture; electrodes “D” and “E” are pulsed during ejection. The resulting axial DC potential is shown in Figure 3.4.




Figure 3.4: Axial DC potential of the BPT Mk IV during the commissioning ^8Li experiment, as determined from SIMION using the applied voltages from the experiment. Applied voltages are shown in Table 3.1. The potential during trapping (blue), during a capture pulse (orange) and during ejection (green) are shown. The widths of the DC electrodes are shown as grey bands. The center region has no DC electrode, with the potential supplied directly onto the rods.

Due to the back-to-back α particles emitted from the ^8Li decay, it is possible to image the ion cloud directly [78]. Using the double α coincidences, the ion cloud cooling may be observed during the trapping cycle, as shown in Figure 3.5. Details of the trapping cycle are provided in section 3.3. The cooling time of about 30 ms is similar to that observed in previous BPT experiments [78]. The spatial extent of the ion cloud is determined through a comparison to simulations; in the figures shown, this was performed assuming that all DSSD strips were functional, and therefore should be taken as an estimate only. A precise determination of the ion cloud size will be performed in the course of subsequent data analysis to determine $a_{\beta\nu}$. From this initial estimate, however, the axial FWHM is close to 3.5 mm and the radial FWHM is about 0.9 mm. The axial cloud size is roughly 15% larger than that calculated from the curvature of the axial trapping potential, but this is fairly good agreement. This final ion cloud size (accounting for the extent of 99.7% of the ions) is larger than the design assumption, but a built-in tolerance means that this ion cloud size still has a full view of the detectors and the electrodes do not block decay products. (Details of the BPT Mk IV design are given in section 2.5.)




Figure 3.5: ^8Li ion cloud cooling as a function of cycle time with (a) showing axial cooling and (b) showing radial cooling. The first few ms of data are unusable as noted in the text. The two different DSSDs pairs are shown separately in each plot, (right-left, R-L, and top-bottom, T-B). There is good agreement between the two pairs, and discrepancies may be due to non-functioning strips, which have not been accounted for.

3.3 Trapping cycle of the BPT

The BPT operates by accumulating ion bunches to build up an equilibrium population during data-taking before ejecting all of the ions to perform a background measurement. The cycle for this ^8Li run was an ion bunch injection every 160 ms for a total of 70 bunches, then an 800 ms background measurement, for a total of a 12 second measurement cycle. Data is ignored during the first ~ 30 ms following injection, during thermalization with the 80 K buffer gas, to avoid systematic effects associated with a larger ion cloud size. The ion cloud cooling can be measured directly due to the back-to-back α particles in the decay. Figure 3.5 shows the cooling as a function of time after capture, demonstrating that this occurs in about 30 ms. The background measurement is used to search for a contribution to the measurement spectrum, thought to be primarily from untrapped ^8Li since random triple coincidences from cosmic rays or other sources are unlikely.

The trapping cycle of the BPT can be modeled to evaluate the trapping lifetime of the ions and to identify the optimal trapping cycle to maximize count rates. Practically, this is rather intuitive and does not require extensive modelling: ideally, enough ion bunches should be injected into the BPT to maintain an equilibrium population for as long as possible before a background measurement is performed. However, there can be some gains in efficiency through better modeling and, as a particular motivation, the specific trapping cycle can have an impact on the background measurement and scaling. Since the background measurement occurs at the end of the 12 s trapping cycle, the number of events observed during this time must be scaled to account for the ^8Li lifetime and the trapping cycle length. The model presented here is specific to the BPT due to its transport and delivery system (section 3.3) but can be easily adapted for different ion traps.

In the BPT beamline, there are four main features that play a role in the trapping cycle: 1) the continuous beam from the gas catcher, 2) the RFQ buncher that produces a bunched beam, 3) the isobar separator Penning trap (IS) that accumulates bunches before ejecting

them as a group, and 4) the BPT which accumulates these secondary bunches. The ion travel time between these elements is ignored since this is very small (tens of μs) compared to the other timescales. The cycling time for all elements is set by the buncher time t_b because this is the first element that converts ions from a continuous beam to a bunched beam. Throwing away bunches needlessly will reduce the overall data rate.

The trapping model then depends on the following variables, with values from the latest ^8Li experiment given:

- τ - the lifetime of the species, $840 \text{ ms} / \ln 2 = 1211 \text{ ms}$ for ^8Li
- t_b - the buncher time, 40 ms
- c - the number of IS injections from the buncher, 4
- d - the number of BPT injections from the IS, 70
- v - the veto time after BPT injection, 35 ms
- l - background measurement time, 800 ms

The continuous beam from the gas catcher is bunched and cooled in the buncher. We take the beginning of the bunching cycle as the start of our clock, such that the population of ions in the buncher is 0 and their lifetime clock is reset. The build-up of the ion population N in the buncher is given by

$$\frac{dN_b}{dt} = r - \frac{1}{\tau}N_b, \quad (3.1)$$

where we assume the rate of capture r is constant and τ is the lifetime of the radioactive ion (but could also include the trapping lifetime). The solution to this differential equation is

$$N_b(t_b) = \tau(1 - e^{-t_b/\tau}), \quad (3.2)$$

where we have taken $r \equiv 1$ for simplicity. At the end of every buncher cycle t_b , we have a population N_b that gets injected into the IS, ignoring any effect on the capture efficiency from shorter or longer storage times. Several bunches are accumulated in the IS for a time. We again ignore any efficiency losses that depend on e.g. the number of ions in the IS. For a single injection to the IS, the population t time after injection is given by

$$N_{IS}(t) = N_b e^{-(t)/\tau}.$$

After c number of IS injections, the population is

$$\begin{aligned} N_{IS}(t + (c - 1)t_b) &= N_b e^{-t/\tau} \sum_{m=0}^{c-1} e^{-mt_b/\tau} \\ &= N_b e^{-t/\tau} \frac{e^{t_b/\tau} - e^{-(c-1)t_b/\tau}}{e^{t_b/\tau} - 1}. \end{aligned} \tag{3.3}$$

At the end of every IS cycle, we then have a population N_{IS} that gets injected into the BPT after an additional t_b (just before the next injection when the IS cycle starts over).

$$\begin{aligned} N_{IS}(ct_b) &= N_b e^{-t_b/\tau} \frac{e^{t_b/\tau} - e^{-(c-1)t_b/\tau}}{e^{t_b/\tau} - 1} \\ &= N_b \frac{1 - e^{-ct_b/\tau}}{e^{t_b/\tau} - 1}. \end{aligned}$$

The derivation for the BPT follows the same pattern as for the IS. However, now we are interested in the number of decays that occur during each trapping cycle, which should be directly proportional to the population of the trap. So the number of decays detected at time t after BPT injection number j (total number of injections d) is proportional to the following expression:

$$\begin{aligned}
N_{BPT}(t + (j - 1)ct_b) &= N_{IS}e^{-t/\tau} \sum_{m=0}^{j-1} e^{-mct_b/\tau} \\
&= N_{IS}e^{-t/\tau} \frac{e^{ct_b/\tau} - e^{-(j-1)ct_b/\tau}}{e^{ct_b/\tau} - 1}.
\end{aligned} \tag{3.4}$$

The expanded expression is

$$N_{BPT}(t + (j - 1)ct_b) = \tau(1 - e^{-t_b/\tau}) \frac{1 - e^{-ct_b/\tau}}{e^{t_b/\tau} - 1} \frac{e^{ct_b/\tau} - e^{-(j-1)ct_b/\tau}}{e^{ct_b/\tau} - 1} e^{-t/\tau}. \tag{3.5}$$

We can simply integrate the expression to get a quantity proportional to the total number of decays D_j for a given BPT cycle j . Only one term depends on the time after injection t and we integrate from $t = 0$ to the next injection, which occurs at $t = ct_b$

$$D_j = \tau(1 - e^{-t_b/\tau}) \frac{1 - e^{-ct_b/\tau}}{e^{t_b/\tau} - 1} \frac{e^{ct_b/\tau} - e^{-(j-1)ct_b/\tau}}{e^{ct_b/\tau} - 1} \int_0^{ct_b} e^{-t/\tau} dt. \tag{3.6}$$

Since the ion cloud takes some time to cool down after injection, we would typically want to remove those decays from our analysis. We can include this in our equation by integrating from a time different than 0, our veto time v (typically 20-30 ms for the BPT).

$$D_j = \tau(1 - e^{-t_b/\tau}) \frac{1 - e^{-ct_b/\tau}}{e^{t_b/\tau} - 1} \frac{e^{ct_b/\tau} - e^{-(j-1)ct_b/\tau}}{e^{ct_b/\tau} - 1} \int_v^{ct_b} e^{-t/\tau} dt. \tag{3.7}$$

Using this final expression, the different trapping parameters can be optimized for the particular species of interest. One detail not included in the derivation here is the fact that, during the 800 ms background measurement, ions are stored in the IS until the BPT is ready for injection again. This means that the first BPT injection likely contains a larger number of ions than the other injections. Additionally, the lifetime of the ions may be modified due to a finite trap lifetime.

The model presented, with the inclusion of a larger first BPT injection, can be compared against the ^8Li data, using the 2α coincidence to increase statistics. As shown in Figure 3.6,

there is fairly good agreement with the model, though the number of events during the first capture is not modeled as well. A fit of the model is performed to estimate the trap lifetime, resulting in a trap lifetime of 83 ± 29 seconds. Given that the entire measurement cycle is only 12 seconds, this is clearly not a precision measurement of the trap lifetime, but it does indicate that the trap lifetime is more than long enough.




Figure 3.6: Trap cycle model comparison with the ^8Li , including a larger first injection and a finite trap lifetime. Events where two α were detected are used. See text for details.

3.3.1 Impact on background measurement

The triple correlation of the two α and the β from the ^8Li decay provides a very clear experimental signature, but it is possible for some background events to be present. There are two cases of background to consider: 1) background uniform in time and 2) a background dependent on the trapping cycle, perhaps from untrapped ions that land on the DSSDs. Each of these has a different scale factor associated with it. In past BPT experiments, the background was taken to be uniform, so it is interesting to see the impact of a better estimate of the untrapped background. The background measurement occurs after the d BPT cycle.

This lasts for some length of time l after the last injection has been held for ct_b .

In the uniform case, we merely need to scale the measured background by the ratio of the background measurement to the live time. For the entire BPT cycle, this scale factor is

$$R_{\text{uniform}} = \frac{\text{live time}}{\text{background time}} = \frac{d(ct_b - v)}{l}. \quad (3.8)$$

In the second case of untrapped ions, we use the model results to investigate a trapping-dependent background. We assume that some constant fraction of ions injected into the BPT become untrapped and serve as a background, perhaps ions that charge-exchange, become neutral, and land on trap elements or detectors, or ions that escape the potential during injection or thermalization. This background is then dependent on the population in the BPT by some additional untrapping efficiency η (set to 1 for simplicity), and we can use the expression already derived for the population in the BPT. The background measurement occurs after the last cycle d so we need to adjust the integral as well. The measured number of background counts is

$$B_m = \tau(1 - e^{-t_b/\tau}) \frac{1 - e^{-ct_b/\tau}}{e^{t_b/\tau} - 1} \frac{e^{ct_b/\tau} - e^{-(d-1)ct_b/\tau}}{e^{ct_b/\tau} - 1} \int_{ct_b}^{ct_b+l} e^{-t/\tau} dt. \quad (3.9)$$

Now we need to find the number of background events expected during the live time. The background B per BPT cycle j is

$$B_j = \tau(1 - e^{-t_b/\tau}) \frac{1 - e^{-ct_b/\tau}}{e^{t_b/\tau} - 1} \frac{e^{ct_b/\tau} - e^{-(j-1)ct_b/\tau}}{e^{ct_b/\tau} - 1} \int_v^{ct_b} e^{-t/\tau} dt. \quad (3.10)$$

However, we should also include the untrapped ions that remain from the previous BPT trapping cycle since these will not get ejected. The expression is then more properly

$$B_j = \tau(1 - e^{-t_b/\tau}) \frac{1 - e^{-ct_b/\tau}}{e^{t_b/\tau} - 1} \left(\frac{e^{ct_b/\tau} - e^{-(j-1)ct_b/\tau}}{e^{ct_b/\tau} - 1} \int_v^{ct_b} e^{-t/\tau} dt + \frac{e^{ct_b/\tau} - e^{-(d-1)ct_b/\tau}}{e^{ct_b/\tau} - 1} \int_{ct_b+l+(j-1)ct_b+v}^{ct_b+l+jct_b} e^{-t/\tau} dt \right). \quad (3.11)$$

The importance of including this contribution can be seen in Figure 3.7. Note that this second factor is not included in the expression for B_m because, for the BPT, the full trapping cycle is much longer than the lifetime and essentially all background counts from the previous trapping cycle will have decayed after the first few injections. For a longer-lived species, this correction could be included. The untrapped population from cycles further in the past could also be included if the trapping cycle is much shorter than the lifetime.




Figure 3.7: Untrapped ion background from current cycle and previous cycle, demonstrating the need to account for untrapped ions leftover from the previous measurement cycle.

The total number of background counts over the cycle is the sum over each injection j .

$$B = \sum_{j=1}^d B_j$$

Finally, the background scale factor for the untrapped ion case is then

$$R_{\text{untrapped}} = \frac{B}{B_m}.$$

A comparison, shown in Figure 3.8, between these two assumptions for the background source shows that the scale factor from the untrapped ion case is larger, up to about 40% larger at the 70 BPT injections per measurement cycle used in the ^8Li experiment. This scale factor is therefore taken as the more conservative case.




Figure 3.8: Background scale factors as a function of the total number of BPT injections and the number of injections to the IS. The other parameters are from the most recent ^8Li run; the selected number of IS captures was 4, which gives a ratio slightly higher than the uniform background assumption (black dashed line) but results in a higher data rate efficiency.

3.4 Experiment analysis

The analysis of the ^8Li data is on-going. Thanks to previous iterations of this measurement, much of the analysis infrastructure is in place, though there are some specific details that must be accounted for with this data set. Several aspects of the data analysis are presented here, though the full, blinded analysis is still on-going.

3.4.1 Simulations

Details of the simulations used in the analysis have been presented previously [41] and are largely unchanged from prior work. Some key aspects are presented here. The different simulation pieces follow the diagram shown in Figure 1.17.

The event generator is based on code written by Nicholas D. Scielzo [32] and updated by Aaron T. Gallant. A decay is created following the differential rate given in Ref. [21], which includes the recoil-order form factors. As mentioned in section 1.5, these form factors have been both experimentally measured [55–58] and theoretical calculations have been performed to determine their size [54]. The final state distribution of the decay is also included, which determines where the decay lands in the excitation energy spectrum of ^8Be and its associated Q value. Incorporating these effects leads to the generation of most events in the spectrum.

Bremsstrahlung must also be accounted for; in about 6.5% of events, a photon is created as the β exits the ^8Be nucleus [41]. The additional steps and correlations required to add a photon to the event generation are described in Ref. [22]. Additional details can be found in Ref. [41].

To simulate β -decays created from a tensor interaction, a similar procedure is used. Ref. [21] provides the decay rate assuming the SM; to change from A events to T events, the sign (\pm) of each term containing a $\beta - \nu$ correlation is swapped (details provided in Ref. [41]). Given that the tensor interaction has been constrained to such low contributions, only the leading order terms matter.

After creating a decay event, its origin inside of the ion cloud is determined. The ion cloud is assumed to be a 3-dimensional Gaussian whose size and position is determined from fits to data. The decay origin is placed inside of this distribution and the particle momenta propagated to the position of the detectors.

For the β particles, it is necessary to perform a simulation of scattering using Geant4. This effect is modeled with a detailed Geant4 simulation using the “option3” standard electromagnetic physics list [102–104]. A full CAD model of the trap is included to model the scattering. Small materials, such as bolts and wires, are not included in the simulation to reduce the complexity required; these are also located far from the ion cloud and have negligible impact on the scattering. More information about the Geant4 simulations, including comparisons to the previous ^8Li data set, can be found in Ref. [41].

Finally, the detector response must be included in the simulation. A full characterization of the DSSDs has been completed in this work and can be found in chapter 4. All detector effects contributing to the detector response are included in the discussion there.

The final simulated events are stored in a ROOT file for analysis and comparison to data. By performing a fit to a linear combination of the axial-vector events and the tensor events, a limit on the tensor fraction in the ^8Li data can be obtained.

3.4.2 Timing information

As noted in section 2.7, the BPT events contain timing information about each front strip trigger relative to the first trigger of an event. The relativistic β particle will always hit the DSSD well before an α , as shown in Figure 3.9. In principle, a silicon detector with good enough timing resolution (a few ns) should be able to easily discriminate between them. Some additional details about the intrinsic timing resolution of the DSSDs are provided in section 4.4.3.

The situation in the BPT is decidedly more complicated. Because the DSSDs have considerable RF pickup (section 2.6.1), notch filters (section 2.6.2) and a long shaping time




Figure 3.9: Time of flight for α (blue) and β (orange) particles as a function of kinetic energy, assuming a 7.5 cm distance from the trap center to the DSSD surface.

of 4 μs (section 2.7) are used to filter out the RF signal from the shaped energy pulse. However, as detailed in section 2.7, a leading edge discriminator triggers *on the shaped pulse* to avoid constant triggering on the RF pickup. This effectively destroys any timing resolution that the DSSDs would have and also introduces a considerable walk to the trigger as a function of the particle energy. The walk can be seen in Figure 3.10, where a walk of up to 5 μs can be observed for low energies. This is in line with the shaping time of 4 μs .

Though they should arrive at the DSSD first, β particles are observed to trigger first in only a few percent of events. The timing resolution can be seen to be several hundred nanoseconds at best, as well. Therefore, the BPT is unable to discriminate between α and β particles on the basis of timing information and must rely on energy cuts to identify particles. Timing information can be used to reject random coincidences with calibration source α particles, however, which effectively removes the vast majority of them from the data set. Additional energy and position cuts remove the rest of these random coincidences.

In addition to the relative timing between particle hits, the time after capture and ejection




Figure 3.10: Timing information from triple correlation events (only one run is shown). Timing is relative to the trigger of the first particle; a considerable walk of up to 5 μs can be seen as a function of particle energy. The blank space around 700 keV comes from a particle identification cut.

is also recorded. During capture and ejection, the electrodes are quickly pulsed by a relatively large voltage ($\sim 100V$); consequently, a large amount of noise triggers are observed in the data during these times. However, there is additional noise seen up to about 10 μs after a capture or ejection pulse. This noise is only seen in conjunction with a pulse, though the reason for the long timescale relative to the fast pulse is unclear. Data during the first 10 μs after capture is unusable to the high level of noise. This does not particularly affect the analysis however, since the ion cloud is cooling during this time, and this data would be discarded for the final measurement.

3.4.3 DAQ linearity

During the ^8Li campaign, the linearity of the DAQ was tested using two PB-5 precision pulsers [138], one for the front strips and one for the back strips since these require opposite polarity signals. These pulsers were coupled to the pulser input of the MPR-32 preamplifiers,

as shown in Figure 2.15, with each pulser signal split between the 8 inputs such that all channels received simultaneous signals. The two pulsers were triggered so that pulses were sent simultaneously to both front and back channels. The DAQ then triggered normally on the front strip signals, as would occur for an α event. It is easy to separate pulser signals from α signals due to the high multiplicity of these events.

During the data campaign, a pulser test was conducted approximately every four hours. This involved sending ~ 3000 test pulses (100 Hz for 30 seconds) at 24 different energies across the ADC range; such a test took approximately 15 minutes. These tests were automated by a Python script to ensure consistency. By looking at the DAQ response to these test pulses, the linearity and stability of the DAQ electronics can be tested. In between these tests, the pulsers were disconnected from the rest of the DAQ due to an apparent increase in noise when they were coupled. The source of this increased noise is not clear, but it could be due to poor grounding between the pulser and the preamplifiers.

In addition to these regular pulser tests, a final linearity test was conducted at the end of the data campaign using 137 peaks with extremely fine spacing (5 seconds per energy at 250 Hz). This test was performed to check the effect of the baseline restorer setting (section 2.7) used in this measurement, to ensure that this did not alter the linearity or affect the efficiency of events. This effect was observed during initial testing with strong baseline restoration. A much lower baseline restoration strength was used in the ^8Li measurement, but the finely-spaced pulser test is essential to determine any potential impact from the baseline restoration.

Figure 3.11 shows the time structure of the applied pulses for one standard pulser test; this is used to separate the peaks to fit them with Gaussian distributions individually. This is particularly important in the case of the finely-spaced pulser test, where the individual peaks have energy overlap. The pulser peaks are applied starting at lower energies and increasing to higher energies. The first peak is below the trigger threshold for some strips, hence the fewer triggers in Figure 3.11 for the first interval. At higher pulser peaks, close to or beyond




Figure 3.11: Time structure of a pulser test. Vertical axis has arbitrary units but scales with the number of triggers (blue). The green line is used to determine the time cuts (orange lines) to separate the pulses.

the end of the ADC range, some strips begin to double trigger and have other non-linear effects, leading to the higher number of triggers seen. Time cuts were automatically determined for each individual pulser test and used for data sorting later. Using the time cuts, the pulser peak at each energy were individually fit for every strip. The resulting fits can be compared with the pulser voltage (an indication of the amount of injected charge) to determine whether there is a linear relationship.

An example of this procedure for one strip is shown in Figure 3.12. There is a clear non-linearity for the upper $\sim 2\%$ of ADC channels, above about channel 4000. This is observed on all strips, though the exact location of the non-linearity is not constant. From this, events above channel 4000 should be discarded, as it is too difficult to accurately determine the energies in this range.

In addition, note the residuals shown in 3.12b; these have a very strange structure. The odd residual structure comes from a mistake in the operation of the pulser. Internally,




Figure 3.12: (a) Linearity evaluation from a pulser test, showing significant deviation about ADC channel 4000. (b) Residuals show a strange behavior that comes from pulser attenuation settings used, as detailed in the text.

the pulser has several different attenuation options (presumably through resistors) that can attenuate a higher voltage to a lower one; this can give more precise small voltages. Unfortunately, due to my misunderstanding, these attenuation options were changed during the pulser tests rather than using a single attenuation. This was intended to give more accurate pulses, but the pulser linearity itself changes with the use of different resistors, which makes sense in hindsight.




Figure 3.13: (a) Linearity evaluation from a pulser test restricted only to peaks with identical attenuation settings (colors). (b) Residuals.

Restricting analysis of the pulser data to those peaks with identical attenuation settings shows very good linearity with sensible residuals, as seen in Figure 3.13. Unfortunately again, several attenuation settings have only 2 or 3 peaks in them, making it impossible to use a quadratic fit to determine the size of the non-linearity. For the last two attenuation settings (yellow and blue), however, it is possible to use a quadratic fit. The non-linearity in these regions is determined to be negligible ($\lesssim 0.02$ keV). In addition, the linear coefficient at all of the attenuation settings is found to be similar. However, attempting to correct the peak energies for the different attenuation factors is too difficult to perform.

In addition to the linearity of the system, which is found to be quite good, the pulser data gives an indication of the stability of the DAQ over the whole experiment. The pulser stability is actually found to be worse than the stability determined from the calibration peaks. Therefore, it is surmised that the pulser itself drifted during the experiment, and the stability from the calibration sources is used instead.

3.4.4 On-going data analysis

The full analysis of this new ^8Li data set is on-going, with results expected within 1-2 years. The analysis is blinded. In addition, thanks to the reduction of β scattering from the new trap, fits to simultaneously extract a limit on tensor currents and the recoil-order parameters are being investigated. It is hoped that this experimental data will be able to improve the precision to which the recoil-order parameters are known, which will lead to overall improvements in the uncertainty on tensor currents.

Chapter 4

DSSD AND CALIBRATION SOURCE CHARACTERIZATIONS

The BPT has an ultimate measurement precision goal of $\Delta|C_T/C_A|^2 \leq 1 \times 10^{-3}$. The latest systematic uncertainty with ^8Li is $\Delta|C_T/C_A|^2 = 2.8 \times 10^{-3}$ [38]. Of this total systematic uncertainty, cuts to the data and knowledge of the detector response each account for 0.9×10^{-3} . Therefore, in order to reach the precision goal, it is necessary to have a full understanding of the detector response of the DSSDs used in the experiment. Most recently, another experiment at the Isotope mass Separator On-Line facility (ISOLDE) studying the decay of ^8B with DSSDs has also studied the detector response with similar detectors [73, 74]. Some additional comments on comparison to other experiments are noted in section 1.6.1.

This chapter provides an overview of the general silicon detector response to α particles in section 4.2. A characterization of the calibration α sources used is provided in section 4.3. The DSSDs in the BPT experiment are described in detail in section 4.4. (For information about the DSSD timing relevant for the BPT experiment, see section 3.4.2.) An experiment to directly characterize the DSSD detector energy response is described in section 4.5.

Portions of this chapter are adapted from ATLAS experiment proposal #1905 [139]. I appreciate the help of everyone involved in this experiment, particularly Daniel Santiago-Gonzalez, Ivan Tolstukhin, Heshani Jayatissa, and Ryan Tang, who provided assistance with the beam monitor and the data acquisition system. Thanks also to the ATLAS operators who made this challenging experiment possible.

Note that the 2013 version of the Stopping and Range of Ions in Matter program [1] (SRIM) and the stopping powers within are used throughout this work.

	α energy (keV)	intensity (%)
^{148}Gd	3182.690(24)	100
^{244}Cm	5762.64(3)	23.1(1)
	5804.77(5)	76.9(1)

Table 4.1: Relevant α energies and intensities from ^{148}Gd [144] and ^{244}Cm [143].

4.1 Calibration α sources

The α sources used for calibration in this work are ^{148}Gd and ^{244}Cm spectroscopy-grade α sources manufactured by Eckert & Ziegler. Two sources of each isotope were purchased in 2015 with activity of roughly 100 nCi [140, 141]. The sources had their active material electrodeposited onto a platinum backing, with active diameter of 5 mm. These sources were then each cut into four quadrants for use in the BPT; this was not performed precisely so different foils may have slightly different active areas. These commercial, spectroscopy-grade α sources have a much thinner dead layer than those produced by drop-deposition in-house [87], which greatly improved the energy calibration possible in the experiment [41, 142]. The energy distributions of α particles from the sources are investigated in section 4.3. The α energies and intensities from these sources are listed in Table 4.1. Note that ^{244}Cm has additional α energies at much lower intensities that are not considered in this work [143]. The uncertainties on the emitted α energy are much smaller than uncertainties on the detector response and calibration in this work and are, in general, ignored.

4.2 Silicon detector response to α particles

For monoenergetic incoming α particles, the resulting spectrum recorded with a silicon detector is asymmetric and cannot be approximated by a Gaussian at the level of precision for this experiment. The sources of this asymmetry have been largely understood for several decades (see e.g. Refs. [145, 146]), and an analytic fitting function has been developed to approximate the line shape of the detector response [147]. For the typical peak fitting

performed in this work (e.g. calibrations), the analytic fitting function used was taken from Ref. [147] and is composed of a Gaussian distribution and two left-sided exponential distributions to approximate the symmetric and asymmetric components of the detector response, respectively. The fitting function takes the form

$$f(u) = \sum_{i=1}^m \frac{A_i}{2} \left\{ \left(\frac{1-\eta}{\tau_1} \right) \exp\left(\frac{u-\mu_i}{\tau_1} + \frac{\sigma^2}{2\tau_1^2} \right) \operatorname{erfc}\left(\frac{1}{\sqrt{2}} \frac{u-\mu_i}{\sigma} + \frac{\sigma}{\tau_1} \right) + \frac{\eta}{\tau_2} \exp\left(\frac{u-\mu_i}{\tau_2} + \frac{\sigma^2}{2\tau_2^2} \right) \operatorname{erfc}\left(\frac{1}{\sqrt{2}} \frac{u-\mu_i}{\sigma} + \frac{\sigma}{\tau_2} \right) \right\}, \quad (4.1)$$

where m is the total number of peaks and the remaining variables are parameters for the fit. For a full derivation of this function and an explanation of the parameters and their physical interpretation, see Refs. [145–147]. Note that these fit parameters are in general energy-dependent; the parameters also depend on the thickness of the dead layers that α particles traverse (thicker dead layers result in a broader distribution) and other detector characteristics. Therefore, the parameters are treated independently for the ^{148}Gd and ^{244}Cm calibration sources; the parameter σ corresponding to the width of the symmetric component may also be treated independently for the different peaks resulting from the different dead layers. This analytic formula is only appropriate for use in simple treatments, such as fitting calibration spectra to extract the peak energy.

To approximate the more complicated detector response for the BPT, a numerical convolution combines the distributions of each of the physical processes that contributes to α energy loss. These functions model the details of both the source and the detector in order to characterize the physical processes that reduce the α energy and broaden the peaks in the spectrum. The final energy spectrum $H(h)$ is then the convolution of the energy distribution from the source $N(E_s)$ and the detector response function $R(E_s, h)$

$$H(h) = \int_0^\infty N(E_s) R(E_s, h) dE_s. \quad (4.2)$$

	Contribution	Shape	Determined by
Source distribution			
$N(E_s)$	energy loss and straggling due to emission depth	asymmetric	experiment, SRIM
Detector response			
D_{dead}	detector dead layer	Gaussian	SRIM
D_{niel}	non-ionization	asymmetric	Ref. [148]
D_{fano}	e-h pair creation (Fano factor)	Gaussian	Ref. [149]
D_{elec}	electronic noise	Gaussian	reference pulser

Table 4.2: Contributions to the energy loss distribution and detector response function. Adapted from Ref. [145].

Table 4.2 list the various energy loss mechanisms and the distributions that model them; note that these distributions correspond to a symmetric Gaussian component and two asymmetric (tailed) components, similar to equation 4.1.

For a particular incoming α energy, these distributions are largely energy-dependent, meaning that the variation in the response over the $\sim 20 - 30$ keV FWHM spread of a monoenergetic incoming α can be treated as a convolution of independent distributions. This is an important fact since the moments of the convolution of independent distributions is simply the sum of the individual moments. Therefore, the mean energy loss of each of these distributions can simply be added to find the total mean energy loss. Practically, this is very useful since, after performing the numerical convolution, the mean of the resulting distribution can simply be aligned to the sum of the means of the input distributions. This method ensures that the absolute energy of the response is kept, as opposed to only the line shape.

The following subsections provide an overview of the physical processes affecting the detector response, which were simulated in SRIM where noted. The intermediate and final distributions for a ^{148}Gd calibration source are shown in Figure 4.1. The different curves represent the source distribution at different stages of the convolution. The measured calibra-




Figure 4.1: Intermediate and final distributions for a ^{148}Gd source, assuming a 100 nm Si dead layer and 8 keV electronic noise σ . See text for details.

tion source distribution $N(E_s)$ (section 4.3) is shown in black. This distribution is convolved with a 100 nm Si dead layer energy loss distribution D_{dead} (section 4.2.1) and is shown in red. This is then convolved with the non-ionizing energy loss distribution D_{niel} (section 4.2.2) and is shown in blue. Finally, this is convolved with distributions for the electron-hole pair statistics D_{fano} and electronic noise D_{elec} (sections 4.2.3 and 4.2.4) and is shown in green.

This procedure shows that there is about 30 keV of lost energy for this calibration source compared to the emitted α energy; this is an important shift as the BPT experiment would like to have an energy calibration known to 5 keV or better. Note that accurately accounting for both the asymmetry in the distribution and width of the electronic noise Gaussian is required, since these shift the final peak location, as can be seen by the 2 keV shift between the blue and green distributions.

4.2.1 Dead layers

As α particles transit the window and doped layer (the dead layer) of a silicon detector, some energy is deposited into this layer. As there is no drift field in this region, ionization produced is not transported, and the energy is therefore “lost” from the resulting charge collection. This energy loss is statistically distributed and depends on the thickness and composition of the dead layer, as well as the incident α energy. Due to the wide range of α energies in the decay of ${}^8\text{Li}$ and ${}^8\text{B}$, a good knowledge of the detector dead layers is essential.

The detector dead layer distribution D_{dead} can be modeled directly in SRIM by transporting α particles of various energies through the appropriate thicknesses of the relevant materials (in this work: Si, Al, SiO_2 , and H_2O). The assumed densities of the detector materials has a direct effect on the stopping power and thus on the thicknesses measured in later sections. The densities of silicon and aluminum were taken to be 2.32 g/cm^3 and 2.7 g/cm^3 , respectively. Silicon dioxide has a wide range of densities depending on its crystal structure; in silicon devices, the SiO_2 layer is typically formed via thermal oxidation with a density of 2.2 g/cm^3 [150].

For an alpha particle passing through the dead layer, the energy loss distribution is practically Gaussian and is modeled as such. Particles passing through the dead layer at an angle different from normal can be simply treated as passing through a thicker dead layer. The mean of the Gaussian is controlled by the stopping power of the material and the thickness of the dead layer; total stopping powers (both electronic and nuclear stopping) and projected ranges from SRIM are shown in Figure 4.2. Since the stopping power is a function of the energy of the α , it can be numerically integrated as the α loses energy as it traverses the material to provide a final energy loss.

SRIM reports that for He ions, the program’s stopping powers match experiments to within 2%. Since, in this work, only the product of dead layer thickness and stopping power is required, all dead layer thicknesses reported assume SRIM stopping powers and uncertainty on the stopping powers are not incorporated into the reported uncertainties on

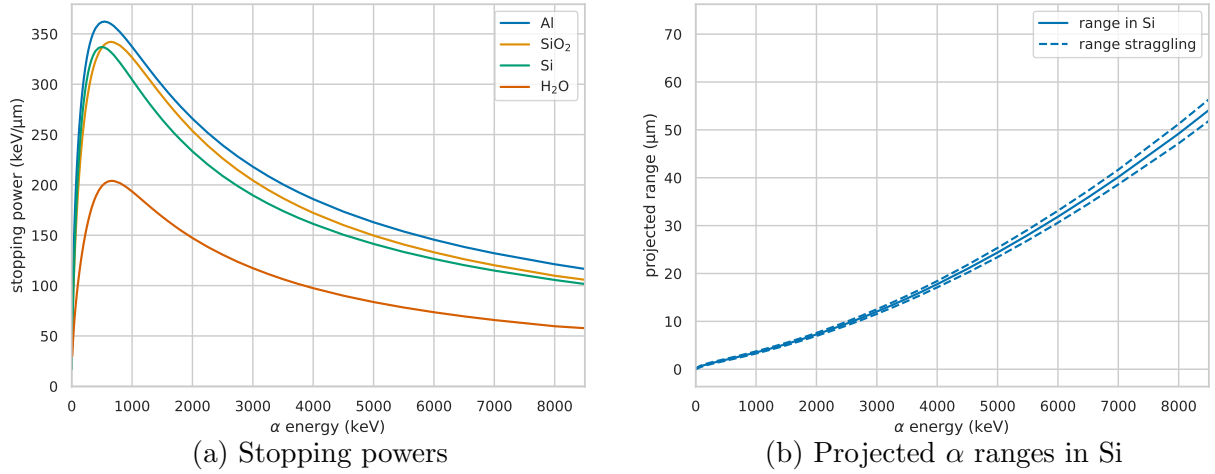


Figure 4.2: Total stopping powers (including both electronic and nuclear stopping) and projected ranges and range straggling from SRIM for α particles of materials relevant to this work.

the dead layers.

For compounds where hydrogen bonds have a sizable impact on the stopping, such as water, a compound correction factor for SRIM is used. It is entirely unclear from SRIM documentation what correction factor should be used for solid water (the phase of the target also impacts the scattering). However, Figure 4.3 taken from the SRIM website [151] suggests a correction factor 0.95 may be appropriate. In this work, the density of ice of 0.917 g/cm^3 is used. These uncertainties are likely unimportant since an H₂O dead layer has so far only been identified during an experiment to characterize the DSSDs used in the BPT (section 4.5), but not during the BPT ^8Li measurement itself.

4.2.2 Non-ionizing energy losses

After an α enters the depleted portion of the silicon detector, the particle's energy can produce detectable ionization in the form of e-h pairs, forming the D_{ion} distribution. The detector volume has an electric field applied to it to collect the e-h pairs and produce a signal that can then be amplified. When an α particle enters the detector volume, it loses energy to electronic ionization and excitation, as well as to recoils with silicon nuclei. With




Figure 4.3: Helium stopping in water vapor and solid, suggesting a correction factor of 0.95 for solid water. From Ref. [151]

enough recoil energy, these silicon nuclei can in turn produce ionization or dissipate the energy as phonons. Energy from the alpha particle that is dissipated as phonons is not collected by the detector and, thus, the energy of the alpha particle can be underestimated. To understand the distribution of energy losses to ionization, it is equivalent to model the non-ionizing energy loss distribution (NIEL). The vast amount of energy of the alpha is given to ionization (> 99% according to SRIM calculations), but the long-tailed asymmetry in the NIEL is important to reproduce the asymmetry seen in the detector response.

To model the NIEL distribution, alphas from 300 keV to 8700 keV in 100 keV increments were propagated through up to 50 μm of silicon (depending on the range of the α) in SRIM. SRIM itself does not record the NIEL distribution, but, using its output files, the distribution can be reconstructed. In particular, SRIM reports the average non-ionizing energy loss for a particular simulation; by running many simulations with a single α each, the reported average loss can be built into the NIEL distribution. Fits to these simulations were performed using a Burr (Type III) function which reproduced the shape well. The resulting distributions of non-ionizing energy losses for several different α energies can be seen in Figure 4.4. There is




Figure 4.4: Non-ionizing energy loss distributions for several α energies simulated by SRIM.

a relatively weak dependence on incoming α energy.

The NIEL distribution has been investigated by other authors (including Refs. [68, 142, 148, 152]). The more recent of these works also use SRIM, but Ref. [68] reports only the average energy loss. Ref. [142] used a different procedure with SRIM to produce the distributions. A comparison of several distributions is shown in Figure 6 of Ref. [142]; agreement is very good between these distributions for the average energy lost. Interestingly, the simulations performed in the course of this work (Figure 4.4) have a lower average energy loss but agrees well with the peak of energy loss.

A more recent calculation of the NIEL distribution, Ref. [148], used a version of SRIM modified by the program's author to output the NIEL distribution directly. The reported distributions have slightly higher average energy loss and larger tails. As an example, Figure 4.5 compares the NIEL distribution of a 1500 keV α as determined by this work and Ref. [148].¹ However, the distributions of Ref. [148] have a larger peak energy loss compared to

¹Note that in Figure 8 of Ref. [148], there is an error in the scaling of the vertical axis; it should be divided by a factor of 2 [153].

Ref. [142], though the average is similar.




Figure 4.5: Non-ionizing energy loss distribution of a 1500 keV α from this work and Ref. [148].

It is not immediately clear which distribution is more accurate, but it is possible that the distribution of Ref. [148] should be used, given that it simulated the NIEL with a version of SRIM modified by the program’s author. In Ref. [75], a heavier tail was added to the detector response to account for that seen in the data; this could have been attributed in part to the NIEL distribution. However, in Refs. [38, 41], a NIEL distribution more similar to this work was used instead. During the analysis performed in section 4.6, the NIEL of Ref. [148] was adopted.

Note that Ref. [148] was published two years after the calibration source characterization was completed (see section 4.3). The source characterization used the NIEL distribution reported here (Figure 4.4), not that of Ref. [148]. The source characterization also found that a heavier tail was required to explain the source distribution, though this was attributed to the source composition. It is possible that this should have been attributed to the NIEL instead. However, it is beyond the scope of this work to revisit the source characterization

to test the distribution of Ref. [148]. It is possible that the question of the more accurate NIEL distribution may need to be revisited in the future.

4.2.3 Electron-hole pair statistics

The distribution D_{fano} represents the number of created electron-hole pairs for an electron with a particular energy E_e . After the alpha particle enters the detector and ionizes electrons, these electrons create e-h pairs in the detector that are collected. This process is Gaussian, as detailed by Ref. [146] and references therein, with mean w and standard deviation

$$\sigma_w = \sqrt{wE_e F} \quad (4.3)$$

where F is the Fano factor. For photons in silicon at room temperature, $w = 3.67 \pm 0.02$ eV/e-h pair and $F = 0.13$ (see Ref. [146] and references therein for additional details). The BPT has detectors at 80 K; the ionization energy was measured at 90 K to be $\epsilon_0 = 3.81 \pm 0.02$ eV/e-h pair for photons and α particles [154]. In addition, see section 4.2.5 for a non-linearity in the silicon response related to electron-hole pair creation.

4.2.4 Electronic noise and energy resolution

The electronic noise distribution D_{elec} is modeled as a Gaussian with a mean of zero (no energy loss, only spreading). The electronic noise can be determined through the use of a reference pulser at several points in the energy region of interest in the DAQ. The reference pulser used in this work is a Berkeley Nucleonics Corporation model no. PB-5 [138]. In general, the width of the noise Gaussian varies over the spectrum, so mapping out the energy dependence using the pulser is essential.

See section 4.4.1 for an averaging procedure to overcome the electronic noise resolution limit by comparing simultaneous measurements of electrons and holes in a DSSD.

4.2.5 Silicon response non-linearity (pulse height defect)

It has been known for several decades that the energy response of silicon detectors depends on the Z of the incoming ion in what is typically called the “pulse height defect.”² This pulse height defect has been suggested to come from two competing effects: 1) the average e-h pair production energy ϵ , which decreases for higher electronic stopping powers, and 2) e-h recombination in the plasma created along heavy ion tracks. These effects have been studied for around fifty years in various papers (see e.g. [152, 156, 157] and references therein).

In one of the first papers³ to examine this effect with low- Z ions [152], the number of electron-hole pairs produced as a function of incident particle energy is modeled as

$$N_{\text{eh}}(E_0) = \int_0^{E_0} \frac{dE}{\epsilon_0 - kS(E)}. \quad (4.4)$$

Here $\epsilon_0 = 3.67 \pm 0.02$ eV/e-h pair for photons in Si at room temperature [152, 154] and k is a constant, independent of Z , that represents the average length traversed to create one e-h pair. The value of k was measured to be $k = 2.8 \pm 0.3 \times 10^{-4}$ nm/e-h pair for Si [159]. The electronic stopping power is included as $S(E)$. This model was found to explain the dependence of the detector pulse height on Z , although it does not account for plasma recombination, which was assumed to be an important competing factor for heavier ions [152].

However, this model also introduces a non-linearity in the Si response due to the inclusion of the electronic stopping power, which varies as a function of particle energy. The non-linearity is small and can typically be ignored for experiments where the range of ion energies

²The term “pulse height defect” has been used to refer to several distinct effects by various authors at different points in time over the last fifty years (see ch. 11 of Ref. [155]). The effects described in this section are separate from the dead layer losses and non-ionizing losses described in the previous sections, which also depend on the ion Z and must also be accounted for when different species are detected with the same detector.

³Eq. 4.4 is nearly identical in form and interpretation to that of Birks’ Law for scintillators, which describes quenching and recombination effects resulting in a non-linear response [158]. Surprisingly, this similarity in model is not referenced in Ref. [152], which was written 35 years later, and seems to have been overlooked by later authors as well.

is not large. For the purposes of calibration, the resulting non-linearity can more easily be expressed through a series expansion [156]. Given that $\frac{k}{\epsilon_0}S(E) \sim 10^{-2}$, equation 4.4 may be expanded as

$$N_{\text{eh}}(E_0) = \frac{1}{\epsilon_0}E_0 + \frac{1}{\epsilon_0} \int_0^{E_0} \frac{k}{\epsilon_0}S(E)dE + \dots, \quad (4.5)$$

where higher order terms contribute $\lesssim 0.05\%$ for α particles in Si at energies relevant to this work.

In the decays of ${}^8\text{Li}$ and ${}^8\text{B}$, α particles have energies over a considerable range (from $< 1 - 8$ MeV), and this intrinsic non-linearity impacts the detector response. A method for calibration including the silicon non-linearity is described in Ref. [156]. The impact of this non-linearity for the BPT was considered in Ref. [142] (see Figure 8 therein). This silicon non-linearity is included in the analysis of the previous BPT experiments (Refs. [26, 38, 75]). It is not clear, however, that this non-linearity is accounted for in other experiments studying ${}^8\text{B}$, for example, such as Refs. [65, 67, 68, 73]. For example, Ref. [66] makes the approximation, at least for α particles, that the number of e-h pairs is linearly related to the amount of deposited energy (minus the non-ionizing portion of the energy). Other recent authors have argued that this plasma recombination effect may be overstated and did not identify a non-linearity in the silicon response [160]. Previous authors have stated that they have used SRIM to determine the pulse heights of α particles at different energies; this is interpreted to mean that SRIM was used to simulate the NIEL distribution. It is emphasized that SRIM does not compute the number of e-h pairs produced, nor does it account for plasma recombination effects.

In Figure 4.6, an estimate of the impact on the apparent energy of an α is shown. From a hypothetical calibration, it is estimated that could be a significant non-linearity introduced from this effect, on the order of 10-15 keV. This estimate is quite similar to the results presented in Figure 8 of Ref. [142], though that analysis used 3 calibration points. A calibration that corrects for the silicon non-linearity using the approximation in equation

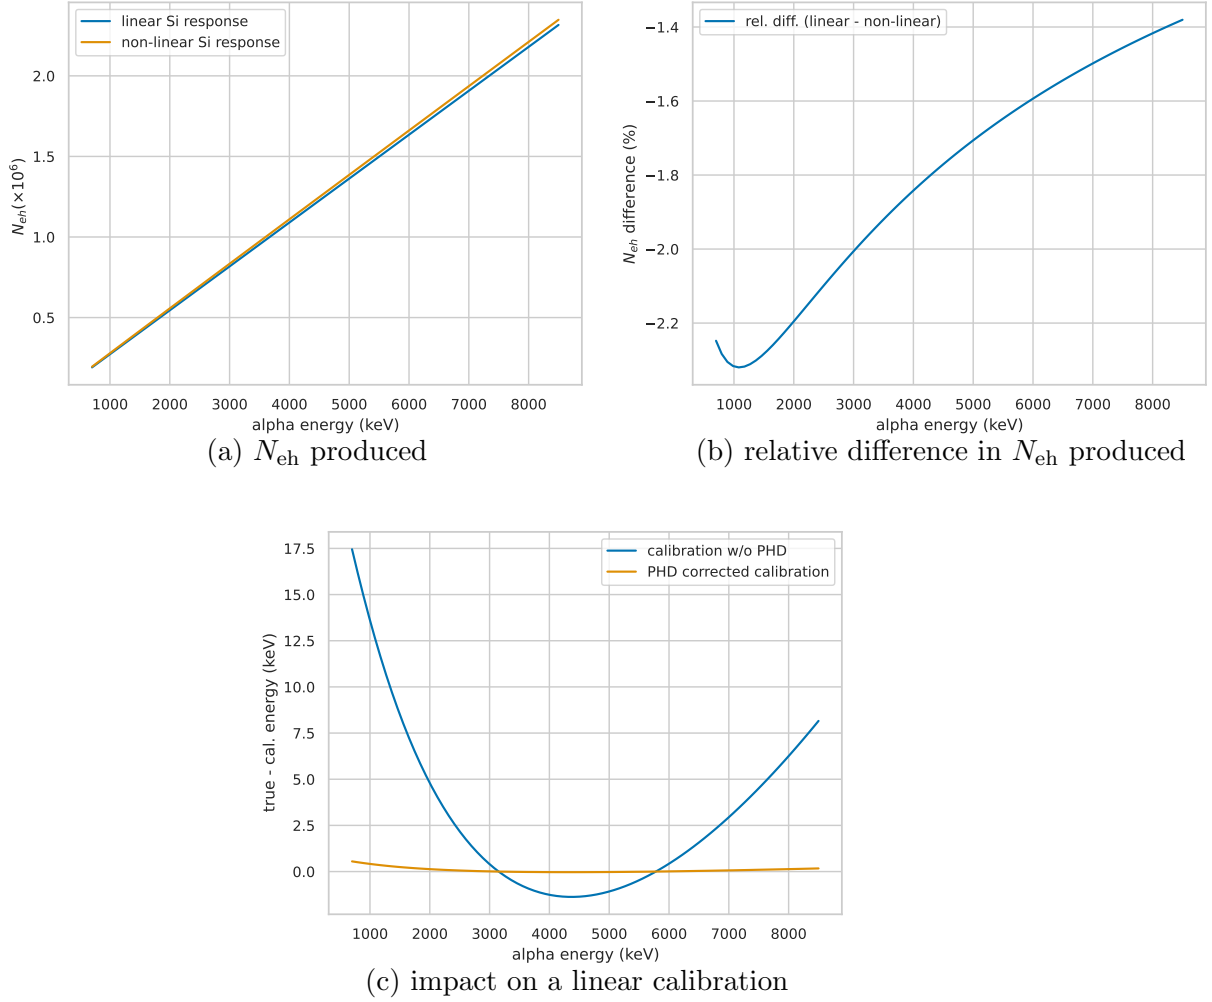


Figure 4.6: The number of electron-hole pairs N_{eh} produced by an alpha in silicon depends non-linearly on its initial energy, according to the model of Ref. [152]. (a) N_{eh} as a function of initial alpha energy according to the linear and non-linear models. (b) The relative difference in N_{eh} between these two models. (c) A hypothetical linear calibration to energy with ^{148}Gd and ^{244}Cm ignoring the pulse height defect (blue), demonstrating that this non-linear effect can have a significant impact on the BPT calibration. The calibration corrected for the pulse height defect (orange) is much improved.

4.5 is shown in 4.6c, demonstrating much better agreement. The disagreement seen at lower energies comes from the missing higher order terms in the approximation.

Note that since the band gap in silicon is temperature-dependent, the ionization energy is also temperature dependent. The BPT has detectors at 80 K. The ionization energy was measured at 90 K to be $\epsilon_0 = 3.81 \pm 0.02$ eV/e-h pair for photons (as well as α particles) [154], though Ref. [155] gives the ionization energy at 77 K to be $\epsilon_0 = 3.76$ eV/e-h pair. In previous BPT analyses, the change in ionization energy was not considered and the room temperature value was used. Fortunately, the impact on the silicon non-linearity is very small, and the absolute number of e-h pairs is irrelevant after calibration. Assuming a room temperature ionization energy instead of the 80 K value was estimated to have a shift of $\lesssim 0.15$ keV on the calibrated energies across the 0.7 - 8.5 MeV range. This is more than an order of magnitude below other calibration uncertainties.

In order for the BPT analysis to be accurate and to reach the stated precision goal, it is likely necessary to be certain that this non-linearity is properly included and that the evidence supports its inclusion. In the experiment described in section 4.5, identifying this possible non-linearity was attempted. However, due to uncertainties in the absolute beam energies, it was not possible to conclusively confirm or dismiss this possibility. For future BPT measurements, this effect may need to be studied further.

4.3 Calibration source characterization

The calibration sources (section 4.1) have a thin deposit of active material from which α particles are emitted. However, this active layer still has some thickness that degrades the energy of particles emitted from deeper in the source. The source energy distribution $N(E_s)$ can be modeled by a Gaussian describing the energy loss of an alpha particle produced at a particular depth, convolved over the thickness of the source [145]. This Gaussian describes both the average energy loss and the energy straggling, the statistical spread in energy loss, which are both proportional to the emission depth. As noted in Ref. [145], with adjustment

to account for the angular dependence of the travel distance of the emitted alphas, this can be modeled as

$$N(E_s) = \frac{N_0 \cos(\theta)}{d_s} \int_0^{\frac{d_s}{\cos\theta}} \frac{1}{\sqrt{2\pi\Omega^2(x)}} e^{-\frac{(\bar{E}(x)-E_0)^2}{2\Omega^2(x)}} dx, \quad (4.6)$$

where N_0 is the number of detected particles, d_s is the thickness of the source, θ is the angle between the normal to the source and the normal to the detector, $\bar{E}(x)$ is the mean energy of the α particles with initial energy E_0 exiting the source from depth x , and $\Omega(x)$ is the energy straggling variation as a function of emission depth. These parameters can be estimated by simulating via SRIM alpha emissions from various depths in a source. The resulting distributions of transmitted energies as a function of emission depth can be fitted and simple, parameterized expressions for $\bar{E}(x)$ and $\Omega(x)$ can be found.

The source distribution was measured in 2018 in a small experiment using a Canberra 25 mm² passivated implanted planar silicon (PIPS) detector (model no. PD 25-11-300 AM). This work was summarized in Ref. [41] and some additional details are reported here.

The calibration sources were placed at a fixed distance and rotated with respect to the PIPS detector so that α particles were emitted at different angles and therefore through different thicknesses of the source. Three different angles were tested: 0° (normal to detector surface), 45°, and 60°. The PIPS detector has a nominal depletion depth of 300 μ m and a nominal dead layer of <50 nm. This detector dead layer was not measured in this experiment and a 5% uncertainty on the dead layer thickness was assigned. The DAQ was a CAEN DT5724 digitizer. A small vacuum chamber was used; the sources were mounted to a rotatable post and the PIPS detector was placed in a fixed position. The setup is shown in Figure 4.7. A reference pulser was used to estimate the electronic noise and track any drift in the system. Four of the BPT calibration sources (two ¹⁴⁸Gd, two ²⁴⁴Cm) were characterized using this system as a consistency check.

A model was created to account for the change in the source distribution at these different angles. In modeling the source, it is difficult to know the actual thickness of the electroplated




Figure 4.7: Setup of calibration sources and PIPS detector for source characterization. The sources are mounted to a rotatable post; other hardware is present from a parallel experiment.

material and its composition. The material that is deposited on the surface of the source can be made of several different species, in addition to the active material, and its thickness can depend on the electroplating time and other factors (see e.g. [161, 162]). These details do not matter very much to the BPT experiment, as only the energy loss distribution of the source is needed, so the source can be treated as composed of a single species.

The spectrum from the experiment indicated that a highly asymmetric distribution was required to explain the source spectrum.⁴ An illustration of the adopted model is shown in Figure 4.8. The source dead layer (green) was found to be unnecessary to explain the distribution and set to zero. This has an exponential distribution with most of the active material near the surface. Two parameters are needed for this model: the source thickness and the length scale of the exponential distribution.

A fitting routine was created to minimize these parameters using the angle-dependent spectra. Uncertainty in the geometry of the setup was estimated and included. Note that this parameterization of the source should not be taken to necessarily correspond to the physical dimensions of the source, given that the source may be composed of several different

⁴In Ref. [148], published in 2020, a NIEL distribution with a larger tail is presented and could account for some of this tail seen in the source distribution: see section 4.2.2 for a discussion. It is beyond the scope of this work to revisit the source distribution model to attempt to incorporate this.




Figure 4.8: Illustration of the source distribution model adopted for the calibration sources: two parameters were needed, the thickness and length scale of the exponential distribution. The source dead layer with no active material (green) was found to be unnecessary and set to zero.

elements with different thicknesses. The length scale parameters for all four sources were treated independently but determined to be about 21 ± 1 nm, with good agreement between all. After testing, it was found that the source thickness could be fixed to 200 nm with little impact on the resulting spectrum, which makes sense given that the length scale parameter is much smaller than this. The resulting source distribution for ^{148}Gd can be seen in Figure 4.1.

Most importantly, using this model, the impact on the calibration peak was determined. Compared to a monoenergetic α source, there is a shift in the peak energy of 3-5 keV for ^{148}Gd and ^{244}Cm with the BPT DSSDs; this shift is also slightly different between ^{148}Gd and ^{244}Cm which introduces a change in the calibration slope. Figure 4.9 shows this shift assuming a Si detector dead layer of 100 nm with 8 keV electronic noise σ . The points are averaged between the two ^{148}Gd sources and two ^{244}Cm sources measured. The two ^{148}Gd sources agreed very well in their final shifts and source distributions and have an uncertainty of 0.2 keV. The two ^{244}Cm sources had slightly worse agreement, and this is incorporated linearly into the total uncertainty of 0.45 keV. Given that there were two ^{244}Cm and two ^{148}Gd sources originally ordered (section 4.1), this is reasonable. There is a slight dependence on the electronic noise; between 6 keV and 12 keV electronic noise σ , there is about a 1 keV




Figure 4.9: Decrease in the peak energy with the modeled calibration source distribution compared to a monoenergetic α , assuming a 100 nm Si detector dead layer and 8 keV electronic noise σ .

shift in the position of the peak and the uncertainty on the peak position can increase by about 30%. The electronic noise effect can be taken into account through the proper convolution of the distributions and is generally better than 12 keV.

For a future, even more precise determination of the tensor contribution, the calibration source model may need to be revisited and each of the sources in the experiment measured and characterized.

4.4 DSSD description

The DSSDs are 1 mm thick Micron Semiconductor Ltd. BB7(DS) 7P/2M detectors with 100 nm junction side windows and 2 mm strip pitch [119]. The detectors are mounted to a custom ceramic PCB to maintain UHV. Front strips (junction side) and back strips (ohmic side) are perpendicular to each other; event coincidence between a front strip and back strip is used to identify the pixel that an α struck. Each strip spans about 2° for the α particles

from the ${}^8\text{Li}$ decay. The DSSDs cover about 25% of the total solid angle. These detectors were specifically chosen for their large area, thin windows to reduce alpha energy losses, and relatively large thickness to sample the energy of the beta from the decay. During operation, the DSSDs are biased to +180 V on the ohmic side, which is at least 30 V over full depletion. The detectors are cooled to ~ 80 K, which reduces the leakage current to $\lesssim 10$ nA and improves the resolution. For calibration, ${}^{148}\text{Gd}$ and ${}^{244}\text{Cm}$ sources are used (section 4.1). Due to the high Q-value of the decay and the broad excited state in ${}^8\text{Be}$, alphas from the subsequent break-up can range in energy from less than 1 MeV to around 8 MeV. A picture of a DSSD mounted to the detector assembly of the BPT is shown in Figure 4.10.




Figure 4.10: A DSSD mounted to the detector assembly of the BPT; junction side shown.

The detector response is a multi-dimensional function that depends on the incoming alpha energy, the entrance angle, and the position of the alpha relative to the surface features of each front strip. The response of the detector is a statistical distribution of pulse heights, both in the incident front and back strips, as well as the neighboring strips due to charge-sharing. The detector response of a different model of a DSSD has recently been studied by another experiment at ISOLDE studying the decay of ${}^8\text{B}$ [73, 74].

When an α enters the DSSD (range $< 50 \mu\text{m}$ at 8 MeV), an electron-hole cloud is created, with the number of pairs corresponding to the amount of energy deposited in the active material of the detector. The electrons drift to the back strips and the holes drift to the front strips, following the electric field lines. Since the back strips are perpendicular to the front strips and the detectors are 1 mm thick, the electron cloud may be shared between two neighboring strips if the α struck near the boundary between two back strips. This back strip charge-sharing occurs for roughly 12% of the α particles, making it a very significant effect. Back strip charge sharing is studied in greater detail in section 4.5.

The front strip surface features of the DSSD are also extremely important in accurately understanding the detector response; these are shown in Figure 4.11 and a comparison between the manufacturer-provided values and the measured values is given in Table 4.3. Micron Semiconductor Ltd. kindly provided additional information and an estimation of the uncertainties of several surface features [163]. A $\approx 100 \text{ nm}$ dead layer is formed by the p+ implantation (ion-implanted boron) and covers most of the strip surface. In Ref. [41], this dead layer was measured to be $100 \pm 3.5 \text{ nm}$ by using a calibrated ^8Li spectrum and examining the energy shift over the surface of the detector. A more accurate measurement of this dead layer using a direct α beam characterization was attempted but was ultimately not successful in extracting a more precise measurement, as detailed in section 4.5.

In addition to the 100 nm dead layer, several different layers exist at the interstrip gaps. At the edges of each strip, a thicker p+ implantation layer and an aluminum strip provide the electrical contact for biasing and read-out. A silicon dioxide (SiO_2) layer insulates between the strip electrodes. This interstrip gap region greatly impacts the detector response function as described in detail in section 4.4.2. Figure 4.12 shows the detector response of a single pixel on the detector to the calibration ^{148}Gd and ^{244}Cm sources in channels. A number of different features are labeled and will be discussed in detail in the rest of the chapter.




Figure 4.11: Dimensions of the DSSD junction side surface features, not to scale. Approximate dimensions only; see Table 4.3.




Figure 4.12: DSSD response from a single pixel (single front strip and single back strip) in ADC channels. Two calibration alpha sources were used as indicated on the figure; a low-energy background peak from cosmic rays is also present. A: normal calibration line; B: back strip charge-sharing from electron cloud drift; C: front strip charge-sharing from interstrip gap; D: incomplete charge collection, presumably without front strip charge-sharing and with opposite polarity pulse on neighboring strip.

surface feature	stated dimensions [163]	measured dimensions	measurement type	dimensions from α beam (this work)	section for details
entrance window thickness	100 nm	100 ± 3.5 nm	^8Li α spectrum energy shift [41]		4.6.3
Al thickness	300 ± 30 nm	270 nm	profilometer	273.7 ± 6.9 nm	4.6.4
Al width	30 ± 1 μm	30 μm	profilometer, optical interferometry	29.2 ± 0.6 μm	4.6.4
periphery p+ implant thickness	500 ± 100 nm			263.5 ± 5.8 nm	4.6.4 and 4.6.6
periphery p+ implant width	50 ± 1 μm	50 μm	profilometer, optical interferometry	50.6 ± 0.9 μm	4.6.4 and 4.6.6
SiO ₂ thickness	-	650 nm	profilometer	see Fig. 4.28	4.6.4 and 4.6.6
SiO ₂ width	25 μm	25 μm	profilometer, optical interferometry	24.6 ± 0.8 μm	4.6.4 and 4.6.6

Table 4.3: Dimensions of key surface features of the BPT DSSDs. Profilometer and optical interferometry measurements performed by Tsviki Hirsh and Tal Zaharoni (Soreq NRC, Israel). Details of the α beam characterization of the DSSD are presented in section 4.5 and section 4.6.

4.4.1 Resolution improvement from simultaneous e-h measurements

After a charged particle creates an e-h cloud in the silicon detector, these charges are separated and drift towards opposite sides of the detector (in a typical silicon detector geometry, such as the DSSDs in use in the BPT). Charges are collected by the anode and cathode of the detector, and the current is detected by a charge-sensitive pre-amplifier that outputs a voltage corresponding to the total charge collected. This charge-collection and amplification process is affected by electronic noise, typically modeled as a Gaussian distribution with zero energy loss. Usually (and in the case of the BPT), this electronic noise dominates the resolution of detector. Note, however, that the electrons and the holes are collected separately and therefore have independent electronic noise contributions for a particular event. This means that by measuring the collected energy of the electrons and holes separately and taking an average for each event, it is possible to achieve a better resolution than by measuring either alone.




Figure 4.13: Improvement in resolution from taking a weighted average compared to a simple average, assuming two independent Gaussian noise contributions.

In the DSSDs used with the BPT, the electrons and holes are already measured separately since the electrons and holes are collected on the back strips and front strips respectively. One complication arises from the fact that the electronic noise contributions on front and back strips can be quite different (from RF pick-up, poor grounding, etc.). This is easily overcome, though, by taking an average weighted by the electronic noise contribution of the strips. The weighted average is simple to perform and is the theoretical best resolution, as seen in Figure 4.13.

Figure 4.13 assumes two independent Gaussian distributions with different σ_1, σ_2 and means of zero. A simple Monte Carlo is conducted by drawing events and computing the average or weighted average. A full maximum likelihood is also computed, as this corresponds to the theoretical best limit but is expensive to compute. (Note that in the large N limit, the weighted average corresponds to the maximum likelihood treatment.) In the weighted average and maximum likelihood cases, knowledge of the $\sigma_{1,2}$ is required. In the case of the BPT, this can be easily estimated from the spectra. The weighted average is simple to perform and aligns with the theoretical best improvement in resolution, up to about 30% improvement. For a small difference in resolution between the front and back strips (up to about 25%), a simple average also improves the resolution to a similar degree. At larger differences in resolution (larger than 75%), however, the average can actually degrade the resolution. Note that in Ref. [74], only a simple average was used between the front and back strips of their DSSD; this is a fine approximation due to the very similar resolutions in most of their detectors. For one of their detectors, an averaging was not performed due to a resolution worse by a factor of 2; if a weighted average had been used instead, a 10% improvement in resolution could have been achieved.

One caveat is that the weighted averaging procedure can bias the averaged result if e.g. either the front or back strip has a poor calibration. However, this can be overcome by re-calibrating the averaged spectra. Finally, note that this procedure only improves the resolution from the electronic noise. Energy loss distributions common to the front and

back strips (dead layer losses, non-ionizing losses, and the Fano factor conversion between deposited charge and e-h pair production) cannot have their resolutions improved. Given that the electronic noise typically dominates the resolution, though, the improvement can still be substantial. As an example of this improvement, Figure 4.14 demonstrates the resolution improvement with a 1550 keV α beam with a single pixel (see section 4.5 for experiment details). The electronic noise σ was determined through a fit with the detector response model; the front energy has $\sigma = 7.0$ keV, the back energy has $\sigma = 6.6$ keV, and the average energy has $\sigma = 4.9$ keV, an improvement of 26% and very close to the theoretical limit.




Figure 4.14: Resolution improvement from a simple average of front strip and back strip energies demonstrated with a 1550 keV α beam.

4.4.2 Interstrip gap region

Key detector effects for the BPT come from the interstrip gap region of the DSSD. These regions of the detector have several different layers of different materials that the α particles traverse and lose varying amounts of energy in. Due the wide range of α energies in the experiment, the amount of energy lost in the layers is highly dependent on the incoming

α energy. Between each of the 32 front strips of the DSSD is a 25 μm region of silicon passivated with a layer of SiO_2 . On each side of this region are the electrode contacts for two strips, which have a 50 μm wide p+ implant layer with the Al contact on top (Figure 4.11). These regions are approximately 1.25% of the detector surface by area. Using a combination of optical microscopy, optical interferometry, and a Bruker DektakXT profilometer [164], Tsviki Hirsh and Tal Zaharoni (Soreq NRC, Israel) were able to measure the interstrip surface feature dimensions directly, shown in Figure 4.15 and Figure 4.16. Note that in the optical microscopy image in Figure 4.15, SiO_2 is optically transparent at the wavelength of light used and therefore appears to have a lower height. These measurements show good agreement with the manufacturer’s specifications, as listed in Table 4.3. An interesting feature is the 10 μm wide, 30 nm deep trench near the Al, which is not present in the design specifications. This corresponds to the boundary of the thicker p+ implant.




Figure 4.15: DSSD interstrip gap using optical interferometry. This picture was taken at the far edge of a strip, hence the rounded corners. Note that SiO_2 is optically transparent at the wavelength of light used and therefore appears to have a lower than expected height. (Credit to Tsviki Hirsh and Tal Zaharoni, Soreq NRC, Israel.)

One of the key detector effects investigated in Ref. [41] was the presence of the “extra




Figure 4.16: Heights and widths of the DSSD interstrip gap features using a Bruker Dek-takXT profilometer [164]. (Credit to Tsviki Hirsh and Tal Zaharoni, Soreq NRC, Israel.)

dead layer,” corresponding to the combination of the thicker p+ implant and the Al contact. In Ref. [41], this was estimated to be made of a 300 nm Al layer and a 270 nm Si layer; these thicknesses correspond fairly well to the stated and measured thicknesses. Note, however, that there are an additional two dead layers that must be considered for a full characterization. Due to the 10 μm periphery of thicker p+ implant outside of the Al contact, an α can pass directly through this layer without passing through the Al layer. These two dead layers have similar surface area coverage (Al & thick p+: 3.0%, only thick p+: 2.0%) but quite different energy losses. Finally, an α can pass through the SiO_2 layer, which covers 1.25% of the surface.

The situation however is further complicated by the charge-sharing that occurs between neighboring strips when alphas strike the interstrip gap. Alphas that strike this interstrip gap region and the surrounding area create electron-hole pairs. Due to the distorted electric field lines near the interstrip gap, and depending on the specific position and penetration depth of the alpha, the holes may either be shared between the neighboring front strips or may get trapped at this SiO_2 interface region, leading to incomplete charge collection on the front strip. In addition holes that are trapped at the interface region can induce an opposite polarity event in the neighboring strip. (Note that the analog electronics that the BPT uses do not trigger on or record this opposite polarity pulse.) Additionally, electrons can be trapped at this interface region, leading to a lower-than-expected collection at the back strip.

Importantly, this incomplete charge collection on the front strips means that the calibration coefficients are not correct for these events, leading to a systematic energy defect for alphas that strike this interstrip gap region. This effect has been observed in the literature (see e.g. [165–168]), as shown in Figure 4.17. Incomplete charge collection has also been confirmed in the DSSDs used by the BPT, as seen in Figure 4.12.⁵ The size of this effect

⁵It is not clear that opposite polarity events have been observed in the BPT detectors since the DAQ used does not record such events. However, these events have been inferred from the energy loss present in the spectrum.

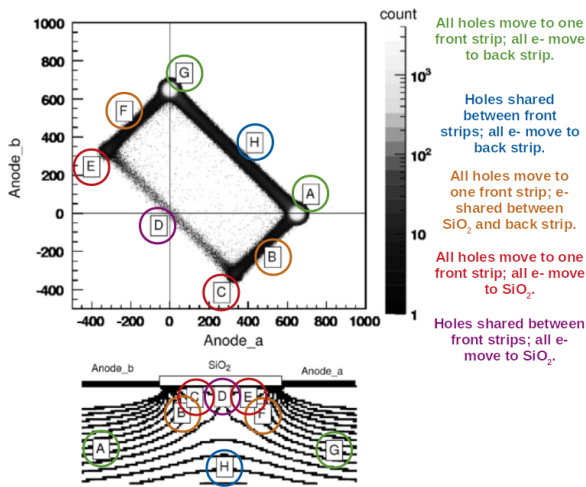
is dependent on the specific position that the alpha hits relative to the interstrip gap center, the energy of the alpha, and possibly the angle at which the alpha enters, making it a complicated effect to study. Events where the holes are shared between front strips can be removed in analysis with a cut, which has been performed in previous BPT analyses. However, it is impossible to remove events that have induced an opposite polarity pulse in the neighboring strip instead due to the BPT analog electronics which do not record such events, making the study of such events important.

Unfortunately, however, this thesis will not be able to put to rest the question of these opposite polarity events. In section 4.5, an experiment that attempted to characterize these events and the full detector response is described, but, for reasons unknown, these opposite polarity events were only recorded when in coincidence with a specific back strip and two specific front strips. As there is no known reason why these events should be linked to a particular back strip, it is assumed that there may have been some problem with the trigger set up of the DAQ and these events were not properly recorded. Further study of these events may be required.

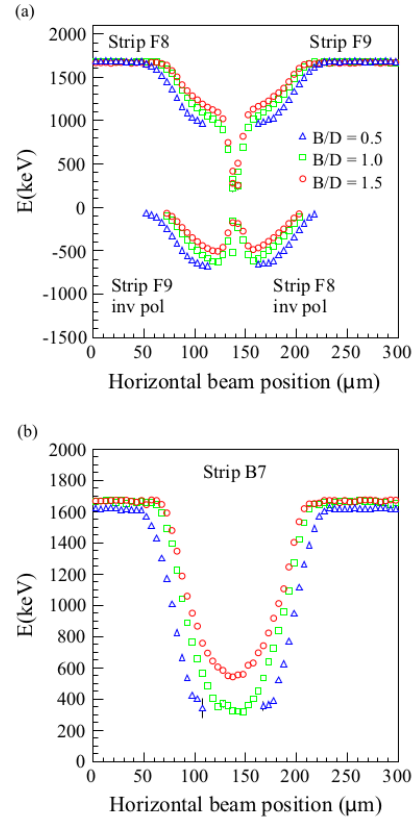
4.4.3 Timing resolution

The intrinsic timing resolution of the DSSDs can be on the order of several ns. Note, however, as detailed in section 3.4.2, that the timing resolution is terrible when used in the BPT due to the RF pickup mitigation strategies used. Briefly, the intrinsic timing resolution of the DSSDs will be discussed here, with BPT-specific modifications of the resolution left to section 3.4.2.

The timing resolution of a silicon detector depends on the electron and hole mobilities, which essentially quantifies how quickly the charge can be collected. Note that mobilities for electrons and holes can differ by a factor of 2-3 [155]. At room temperature, the electron mobility for silicon is $1350 \text{ cm}^2/\text{V}\cdot\text{s}$, while the hole mobility is $480 \text{ cm}^2/\text{V}\cdot\text{s}$. However, at liquid nitrogen temperature, the mobilities increase by a factor of ~ 15 , improving the



(a)



(b)

Fig. 3. Energy signal amplitude profiles for the 998 μm thick DSSSD with a proton beam energy of 1700 keV and a scanning area centered on the front inter-strip F8–F9 and on strip B7. B/D indicates the ratio between the applied bias and the full depletion voltage. (a) Inter-strip behavior for strips F8 and F9. (b) Strip B7 response versus the horizontal beam position.

Figure 4.17: (a) Distortion of the electric field in the interstrip gap region can cause holes to be trapped at the SiO_2 interface. This diagram of simulated data demonstrates how the different regions of the field distortion lead to different pulse heights in the neighboring strips. Taken from Ref. [165], color annotations added. (b) Demonstration of inverted polarity signals in neighboring front strips using a micro-proton beam across the interstrip gap. The lower figure shows how the back strip also suffers from a loss of electrons that are taken up by the SiO_2 interface. Taken from Ref. [166].

intrinsic resolution by approximately this much as well. At 77 K, the electron mobility becomes $2.1 \times 10^4 \text{ cm}^2/\text{V}\cdot\text{s}$, while the hole mobility is $1.1 \times 10^4 \text{ cm}^2/\text{V}\cdot\text{s}$.

The DSSDs are 1 mm thick and have a bias of +180 V. For an α particle, e-h pairs are created within the first $\sim 50 \text{ }\mu\text{m}$ of the front surface. Holes drift to the front of the detector and electrons drift to the back. The expected time scale for the holes to be collected would be something like 0.25 ns. For electrons, this would be more like 5 ns. Shaping of the signals in the charge-sensitive preamplifier (and in a shaping amplifier present for the BPT) will result in a longer time scale. Note that for the experiment discussed in section 4.5, no significant difference between the rise time of signals on the front strip and on the back strip was observed. This may be due to the fact that the preamplifier used, a mesytec MPR-32, has a risetime of 12 ns [123].

For the BPT, β particles are also detected by the DSSDs. These are primarily minimum ionizing particles and traverse the whole 1 mm thickness of the DSSD while depositing charge. This creates e-h pairs throughout the detector. The expected drift times are then similar between the electrons and holes, but they will be accumulated over a longer period of time than the expected short bunch from an α signal. This may result in a differently-shaped pulse that could be discriminated against, though this would likely require digital electronics and waveform recording, which the BPT does not currently use.

4.5 α beam characterization of a DSSD

An α -beam was used to directly characterize a DSSD at several different energies from 1.2 MeV (the lowest possible energy at ATLAS) to 8.5 MeV. This experiment was proposed as ATLAS experiment #1905. This proposal was allocated about one week of beam-time, and the experiment was conducted in two separate segments in January 2021 and March 2021. ⁶

⁶ATLAS is designed for accelerating heavy ion beams for use in nuclear reactions. Using an α beam at low-energies to characterize a detector is an atypical experiment for the facility and would ordinarily likely not be granted beam-time. However, due to restrictions on outside users imposed during the COVID-19 pandemic at this time period in early 2021, the facility was available to perform this experiment. I greatly

The measurement goals for the detector characterization experiment were as follows [139]:

- 1) A complete characterization of the DSSD detector response function for alphas in the energy range 1.2 – 8.5 MeV.
- 2) An investigation of the energy and angle dependence of the response of the DSSD interstrip gap to alphas in the energy range 1.2 – 8.5 MeV, including the inverted polarity signals generated in neighboring strips.
- 3) A measurement of the non-linearity of the response of silicon detectors to alphas in the energy range 1.2 – 8.5 MeV.

In order to perform these measurements, the α beam was required to have a low intensity ($\lesssim 15$ pps/mm²) with a large beam spot ($\gtrsim 1$ cm) in order to homogeneously cover several pixels of the detector surface with a rate that would not saturate the data acquisition system. The beam energy was required to be known absolutely and to have very little spread. In order to study the possible silicon non-linearity, the beam energy needed to be determined by a method other than using a silicon detector. This was attempted by performing a time of flight measurement of a higher-energy beam and then scaling two magnetic dipole field settings based on this energy. In the analysis of the data, the actual beam energy was found to be within 50 keV of the stated beam energy and the beam energy spread was estimated to be less than 10 keV FWHM from the analysis. These beam settings are acceptable for most of the detector characterization, but, unfortunately, the beam energy was not accurate enough to make any definitive statement about a non-linearity in the response of silicon. (See section 4.6.7 for details.) Even though the dipole magnet settings were scaled based on the time of flight, the different energy beams may have entered the magnets at slightly different positions and been affected by slightly different fields.

Of particular interest was a measurement of the detector window, which causes the alphas to lose energy that is not deposited as electron-hole pairs in the detector and therefore is not

appreciate the flexibility and efforts of the facility and its staff.

collected. The manufacturer quotes the window to be 100 nm thick, but a small uncertainty in this value can cause a significant shift in the calibration due to the large angle at which alphas enter the detector in the BPT. This angle is as large as 30° , corresponding to an increase in the dead layer thickness of up to about 15%. However, as explained later, it seems that water vapor may have frozen onto the surface of the detector, increasing the apparent dead layer and preventing an accurate measurement of its thickness. In addition, due to an apparent malfunction or misunderstanding of the DAQ triggers, opposite polarity events were only recorded when in coincidence with a specific back strip and two specific front strips. These opposite polarity events should not depend on which back strip collected charge, so most likely, the DAQ was not set up as intended.

Therefore, the data from ATLAS experiment #1905 is used primarily to explore the energy and angle dependence of charge-sharing and determine the probability of these different kinds of events. The data can also be used to better estimate some of the relative thicknesses of the different layers, though not their absolute thicknesses. Details of the experiment are provided in the next section, followed by an analysis of the data.

4.5.1 Experiment details

The α beam was delivered to the ATSCAT area at ATLAS where an MCP-based beam monitor and the detector characterization chamber were assembled (Figure 4.18). This beam monitor was a prototype designed for use at ATLAS; thanks to Ivan Tolstukhin and Daniel Santiago-Gonzalez (ANL) for their help with the instrument. The beam monitor was required in order to tune the beam, determine that the beamspot is sufficiently large, assess the spatial homogeneity of beam, and ensure that the rate is low enough so as not to damage the DSSD downstream. It is important that the beam is spatially homogeneous so that the detector area of interest is illuminated uniformly; this can be accomplished by defocusing the beam to a rather large beamspot (e.g. radius $\gtrsim 1$ cm) just upstream of the detector. In the beam monitor, the beam passes through a thin film causing electrons to be knocked

off. These electrons are then focused with a permanent magnet towards a position sensitive MCP. By viewing the processed signals from the MCP on an oscilloscope, the shape and size of the α beam can be seen in real time which aids tuning greatly. A defocused beam is shown in Figure 4.19; note that due to mismatched impedance or a cable defect, the edges of the image appear artificially flattened. Importantly, note that the beam monitor is mounted on a motorized linear feedthrough and was removed during the actual measurement of the DSSD response; as the beam passes through the foil of the beam monitor, it loses energy and its energy spread increases, both of which were unacceptable for this experiment. After the beam monitor was removed, a gate valve was opened to expose the DSSD to the beam for the data-taking process, which should last roughly an hour. This process was repeated for each beam energy.




Figure 4.18: Set-up of experiment #1905 (see text for details).

The detector assembly was mounted to a differentially-pumped rotary platform that will allow in-vacuum rotation in increments of about 0.9° (Figure 4.20). The DSSD was attached to the cold finger of a liquid nitrogen dewar and was cooled to ~ 80 K. A detector cover was placed on top of the detector to reduce the amount of residual gas that freezes on the




Figure 4.19: A defocused α beam as imaged by the MCP-based beam monitor and viewed on an oscilloscope.

detector (although this proved not to be sufficient, as described later). A Leybold 350i turbomolecular pump was attached to the chamber, and the pressure reached $\sim 1 \times 10^{-8}$ mbar while the assembly was cold. Note that the ATLAS beam line pressure immediately upstream of the detector chamber was $\sim 1 \times 10^{-6}$ mbar. Although a gate valve separated these two areas most of the time, during measurements with the beam, this was opened and a clear pressure increase in the detector chamber was observed. This higher pressure may have contributed to the water freezing problem discussed in section 4.6.

At each of the beam energies, the detector was rotated to several different angles, from 0° to 30° , which is the maximum angle that alphas enter the detector in the BPT using trapped ions. A mask directly upstream of the DSSD exposed a $7 \text{ mm} \times 10 \text{ mm}$ area on the detector, which caused several pixels to be totally illuminated by the beam while minimizing the trigger rate.

A digitizer system from the Multi-Sampling Ionization Chamber (MUSIC) experiment was used to record the waveforms of these events; this digitizer has up to 40 channels, a 14 bit analog-to-digital converter (ADC), and samples at a rate of 100 MS/s. Only 10 channels were used in this experiment. Five front strips and five back strips were recorded




Figure 4.20: Model of the DSSD vacuum chamber assembly. The differentially-pumped rotary platform is at the top. Cooling is provided by an attached liquid nitrogen dewar (not shown) and a copper rod along the rotation axis. This assembly allows in-vacuum rotation while maintaining cooling. The detector mask is shown at the front.

on the illuminated section of the detector. However, the outer strips were only used when triggering in coincidence with one of the inner strips in order to ensure the full illumination of a pixel. It is important to acquire signals from these neighboring strips in order to identify those interstrip gap events that cause a negative polarity pulse in a neighboring strip. The digitizer triggered on both the negative and positive polarity pulses (though there seems to have been problems with this, as discussed in the analysis later). During the experiment, the trigger rate was typically a few hundred Hz per channel. For calibration, a ^{148}Gd and a ^{244}Cm spectroscopy-grade alpha sources (section 4.1) were mounted to the detector assembly at a fixed position relative to the detector. The experiment attempted to use a reference pulser to monitor drifts in the DAQ, but this introduced a large amount of noise during the experiment that was not able to be resolved quickly. Therefore, no reference pulser peaks are present; drifts were instead evaluated using the calibration peaks (section 4.6.2).

During the first run in January 2021, the rotary stage had not yet been delivered and so the experiment was conducted at a fixed angle. This data set has limited utility but does provide a check on the results obtained during the March 2021 run. During this second

run, the DSSD previously used was found to be malfunctioning and so a back-up DSSD was swapped in. This means that these two runs therefore provide a check on the consistency of results across two different detectors. The nominal beam energies in the March run were: 1200, 1500, 2500, 6500, 7500, and 8500 keV. The January run included these energies as well as 3500 keV α beam.

4.6 Analysis and results

Analysis was primarily conducted on the March data set to evaluate the angle-dependent response of the detector; unless otherwise specified, details and figures are related to this March data set.

At the different energies, the α beam was centered on different spots on the detector as seen in Figure 4.21. The beam pattern matches the detector mask opening of 7 mm \times 10 mm (front strip span \times back strip span), with at least 4 front strips illuminated and at least 4 back strips illuminated, as each strip is 2 mm wide. Unfortunately, excessive noise on back strips 21 and 22 prevented their use with the DAQ. This ultimately reduced the number of pixels that were fully illuminated and that had charge-sharing events recorded. Four pixels, outlined in yellow in Figure 4.21, were fully illuminated and had charge-sharing recorded with neighboring strips. These four pixels were used for analyses that required either full illumination (such as the fraction of events passing through each dead layer) or complete charge-sharing information. For estimates of the primary dead layer thickness, two additional pixels were used, (f-14, b-24) and (f-14, b-25). Calibration was performed by strip as described in section 4.6.2.

Note that these thickness determinations rely on an accurate stopping power and material density. Because this analysis is only interested in reproducing the detector response, not in determining the absolute thickness of the different dead layers, uncertainties on the stopping powers and material densities are not included. Provided that the detector response is properly reproduced, these effects are not important. Thickness results that differ from the

stated or measured dimensions may be due to inaccuracies in the stopping powers or in the material densities.

It was found that the 7500 keV and 8500 keV beam energy spectra of the March data set had heavy tails that make estimating dead layer thicknesses and fractions from these spectra unreliable. These tails were not reproduced by typical NIEL losses, and it is thought that this energy spread was present in the beam.




Figure 4.21: Beam spots at 1500 keV and 7500 keV. Pixels inside of the yellow box are fully illuminated by the beam, along with all charge-sharing, and were primarily used for the analysis.

A comparison between the nominal beam energies and the beam energies as determined by the DSSD for the March run is shown in Table 4.4. See section 4.6.7 and Figure 4.41 for a discussion of the possible silicon response non-linearity.

4.6.1 Data cuts and classification

Each impinging α results in charge collected on a single pixel (one front strip and one back strip) or, if charge-sharing occurs, on up to two neighboring front strips and two neighboring back strips. Because this characterization of the detector response is interested in the fraction

Nominal energy (keV)	energy determined by DSSD (keV)
1200	1240.9 ± 2.4
1500	1550.3 ± 2.5
2500	2550.3 ± 2.0
6500	6480.2 ± 2.1
7500	7482.2 ± 2.8
8500	8431.3 ± 3.0

Table 4.4: Nominal beam energies and energy determined by DSSD for the March run of ATLAS experiment #1905.

Label	Note
f1b1	1 front strip, 1 back strip (single pixel)
f1b2	1 front strip, 2 back strips
f2b1	2 front strips, 1 back strip
f2b2	2 front strips, 2 back strips

Table 4.5: Events are classified by the number of neighboring strips with deposited energy, as shown.

of events that include charge-sharing, each event is classified by the number of strips in which charge is deposited. Table 4.5 shows the labels used for these events based on the charge-sharing. In addition to these labels, some single pixel events with front-back energy disagreement were classified as charge-sharing events where the shared energy was low enough to be missed by the trigger. This cut is detailed below.

To overcome the electronic noise resolution limit (section 4.4.1), a simple average was taken between the front energy and the back energy for single pixel events. This simple averaging (compared to a weighted average) was sufficient because the noise level was very similar between front and back strips. See Figure 4.14 for an example spectrum demonstrating this resolution improvement.

For some beam energies, the α rate was high enough that a strip did not have sufficient time to return to its baseline voltage before another event occurred (pileup). This affects the recorded energy of the event. To remove these events, a cut was placed on the baseline energy in the 1 μ s prior to the trigger such that it was sufficiently close to the nominal

baseline. In addition, events were removed where the total deposited energy on the front strips or the back strips was more than 200 keV larger than the nominal beam energy, as these were most likely due to pileup.

To accurately characterize the charge-sharing of the DSSD and to understand these distributions at very low energies close to the noise pedestal, it was necessary to place a noise cut of 20 keV. For each event, a strip with an energy below the noise cut was removed from the event and the remaining strips were analyzed. The exact position of the noise cut can have an impact on the charge-sharing fractions primarily; a range of reasonable cuts was evaluated for the impact on the uncertainty on the charge-sharing fraction but the impact was negligibly small.

Some charge-sharing events where only a small fraction of the total energy was deposited in a neighboring strip would not pass this noise cut or would otherwise not have triggered due to noise fluctuations. In an attempt to evaluate this, a threshold of 50 keV was placed; a single-pixel event with an energy disagreement between the front and back strips beyond this threshold was assigned as a charge-sharing event. The location of this cut impacts the charge-sharing fraction, but it was found that a 50 keV cut was reasonable and had only a small impact on the charge-sharing fraction, below the other uncertainties.

Finally, since negative polarity events were assumed to not have been properly recorded (see section 4.5), these events were discarded from the few pixels that did record such events.

4.6.2 Calibration

Data was calibrated using two α sources attached to the detector assembly (^{148}Gd and ^{244}Cm , see section 4.1 and note that only the higher energy ^{244}Cm α was used for calibration). These sources are relatively weak and only a few hours elapsed between runs with exposure to the α beam; this means that the calibration peaks contained a relatively low number of counts. For this reason, data was calibrated by strip instead of by pixel. Only α events that struck the pixels of interest were used for calibration to limit the angle of incidence

between the source and the detector, which would affect the calibration. The position of the sources was measured by hand and agreed with the computer-aided design (CAD) model of the detector assembly; this model was then used to find the angle between each source and the relevant pixels. For simplicity in the calibration and due to the low number of counts, the average angle across each front strip and back strip was used, resulting in increases in the apparent dead layer of between 1.7% and 4.8%. The increase across the relevant pixels in each strip varied by at most 1%, corresponding to a difference in the apparent dead layer of about 1.5 nm. This is roughly a factor $2\times$ smaller than the uncertainty on the dead layer thickness (see section 4.6.3 for the dead layer determination).

Calibration was split into two time periods to account for a small but noticeable disagreement between calibration coefficients during two different segments of the experiment. This was checked by comparing calibration coefficients obtained with all calibration files to coefficients obtained with individual files. This change in regime coincided with a DAQ crash that required a full reboot of the digitizers and may have changed some trigger settings or otherwise slightly modified the system response. The first calibration segment contained nominal beam energies of 7500 and 8500 keV; the second segment contained beam energies of 1200, 1500, 3500, and 6500 keV.

A linear calibration was assumed as only two calibration points were available. The calibration α peaks were fit with a function matching equation 4.1. Using the fitted parameters, the peak of the function was found. To account for the source distribution (section 4.3), the peak shift resulting from using a realistic source distribution instead of a monoenergetic source distribution (see Figure 4.9) was computed, assuming electronic noise of 6 keV σ and for a dead layer of 153 nm Si-equivalent (see section 4.6.3). The uncertainty on the peak was evaluated by drawing parameters from a multivariate normal distribution using the covariance matrix (the inverse Hessian) of the fit then evaluating the peak of the function. This method properly accounts for the correlations between the various parameters of the fit. The primary dead layer was taken to be composed of 100 nm of Si and 85 ± 5 nm of

H₂O (section 4.6.3). The thickness of water was randomly drawn from a Gaussian during the fitting routine. The uncertainty on the peak shift from the realistic source distribution was taken as 0.2 keV for ¹⁴⁸Gd and 0.45 keV for ²⁴⁴Cm (section 4.3). This fitting routine was repeated 1000 times, and the standard deviation of the resulting distribution of peaks was taken as the uncertainty. The total uncertainty at the nominal beam energies and calibration energies is shown in Figure 4.22.

As described, the total uncertainty includes statistical uncertainty from fitting the calibration peaks and systematic uncertainties from the realistic source distribution and from the uncertainty on the dead layer thickness. Additionally, a 15% uncertainty on the average NIEL is added in quadrature, following Ref. [142]; this contributes about 2 keV uncertainty. In contrast to Ref. [142], this work ignores uncertainties on the silicon non-linearity (“pulse height defect”) as a goal of the analysis was to determine its existence assuming a linear calibration. The absolute size of the uncertainties on the calibration are rather small (including the aforementioned effects) and primarily affect the determination of the beam energy; there is not a strong impact on the determination of dead layer thicknesses.

The front strip - back strip energy difference for single pixel events provides some indication about the consistency of the calibration, as shown in Figure 4.23.

4.6.3 Primary DSSD dead layer

The primary dead layer thickness of the DSSD window (the p+ implant of nominal 100 nm thickness) was measured by looking at shift in the peak energy when the DSSD is at different angles to the incident beam. The effective thickness of the dead layer increases with the secant of the angle between the beam and the detector. The shift in energy is largest at the lower beam energies, making these energies the most useful in measuring the dead layer thickness. In addition, at the higher beam energies, there is not sufficient energy loss separation between the dead layer of the thicker p+ implant and the primary dead layer. The primary dead layer thickness was determined from the nominal 1500 keV α beam energy for




Figure 4.22: Total energy uncertainty at 1σ from calibration for the α beam characterization of a DSSD for each strip.




Figure 4.23: Mean front-back energy difference for single pixel events.

these reasons, as well as the fact that the majority of the ^8Li and ^8B α spectrum is close to this energy. The other beam energies were used to confirm that the result agreed reasonably well.

In order to measure the primary dead layer thickness, events where an alpha hit a single pixel (one front strip and one back strip) for the center pixels were fit with a function matching equation 4.1. Using the fitted parameters, the peak of the function was found. The uncertainty on the peak was evaluated by drawing parameters from a multivariate normal distribution using the covariance matrix (the inverse Hessian) of the fit then evaluating the peak of the function. This was repeated 1000 times, and the standard deviation of the resulting distribution of peaks was taken as the statistical uncertainty. This method properly accounts for the correlations between the various parameters of the fit. Typical statistical uncertainties on the energy of the peaks are 0.05 keV or lower.

Figure 4.24 shows the peak energy as a function of angle with the nominal 1500 keV α beam for six different pixels. As determined by the DSSD, the deposited energy corresponds to about 1510 keV; the true energy of the alpha beam would be about 1550 keV, accounting for the average energy losses from the dead layer and non-ionizing distribution. Note that the pixels agree within about 2 keV, which roughly agrees with the uncertainty of the calibration. In the fits to the peak shifts, the angle at which the beam was normal to the detector was allowed to vary but was common to all pixels. The fit result of $40.3 \pm 0.9^\circ$ agrees well with the nominal angle of 40° . It was considered that the plane of the detector may not have been perpendicular to the plane of the floor, which could have resulted in some small shift in the data. This was also allowed to vary and was found to be $0.0 \pm 0.9^\circ$.

Unfortunately, it is apparent from the data that the dead layer found here is much thicker than stated (c.f. Table 4.3) or previously measured [41]. It is thought that water vapor from the beam line (at $\sim 10^{-6}$ mbar pressure) may have frozen onto the surface of the detector when the gate valve was opened to allow beam to hit the DSSD. Although the detector assembly had some shielding in place to try to prevent this, it evidently was not enough.

Furthermore, the shift in the peak energy does not follow a line as well as one might expect, as can be seen from the residuals. This could be due to an underestimation of the uncertainty of the peak energy (though this seems appropriately calculated) or as the result of a shifts from the DAQ or a very small deviation in beam energy e.g. from unstable power supplies. In any case, this measurement is unable to accurately determine the thickness of the dead layer. Instead, the result of Ref. [41] of 100 ± 3.5 nm is used in the final detector response.




Figure 4.24: Dependence of the peak energy as a function of angle with the nominal 1500 keV α beam.

However, the apparent dead layer thickness is applied to the calibration of the data and so an estimate is required. Note that since this measurement looks at the change in the position of the peak, a relative offset due to a less accurate calibration does not affect the result. For simplicity, the thickness was estimated using the stopping power of silicon only, as opposed to a combination of water and silicon. The average slope of the shifts is found to be -40.5 ± 0.9 . The average slope is quite consistent between different fitting approaches, but the uncertainty on the result is very sensitive to the exact peak locations and was difficult to estimate consistently. The uncertainty is therefore estimated by taking the average of the

individual slope uncertainties; this estimate is slightly larger (about 15%) than the variance of the individual slopes. From the slope of the shifts, the dead layer thickness was found to be 153.0 ± 3.3 nm (Si-equivalent) with stopping powers taken from SRIM [1]. The dead layer thickness is found to be consistent within uncertainties when using only the stopping power or when including a full detector response to account for asymmetric distributions that could affect the peak position.

In Ref. [41], the primary dead layer was measured in situ in the BPT using the calibrated ^8Li α spectrum, since α particles from the decay impinge on the DSSDs at angles between 0° and about 30° . This measurement gave a result of 100 ± 3.5 nm. Taking this measurement as the basis for the primary detector dead layer, the dead layer measured here can be split into a composition of Si and H_2O . SRIM was used to calculate the stopping power of water (assuming a density 0.917 g/cm³ and a compound correction factor of 0.95). At 1550 keV, the energy loss corresponds to a dead layer composed of the nominal 100 ± 3.5 nm Si and about 85 ± 5 nm H_2O as shown in Figure 4.25. This layer construction is used in the calibration for the analysis of this experiment (section 4.6), including in Figure 4.24. This different layer construction affects the calibration only a small amount (about 0.5 keV at the highest energy of 8.5 MeV) with negligible impact on the determination of charge-sharing fractions.

Note that water freezing onto the surface of the DSSDs is a concern in the BPT experiment as well. However, the situation there is somewhat better than in this beam experiment. In the BPT, the inner surface of the detectors faces an enclosed volume whose sides are held at ~ 80 K, so water molecules are likely to freeze to the outer surface before they reach the inside of the trap volume. The trap is cooled down from the outsides towards the detector surface as well, so water present in the vacuum chamber is likely to freeze onto these surfaces first before the DSSDs are below the freezing point. In addition, the BPT vacuum chamber and entire beam line upstream of the chamber reach pressures of better than 1×10^{-7} mbar prior to cooling, at least an order of magnitude (if not two) better than the ATLAS beam




Figure 4.25: Stopping power comparison between a layer composed of only silicon and one composed of silicon and water.

line. As mentioned, Ref. [41] measured the dead layer in the BPT using a calibrated ^8Li α spectrum and obtained a thickness of 100 ± 3.5 nm, consistent with no water being present. The presence of water in this beam experiment indicates, though, that every BPT experiment should perform a similar measurement against the data to ensure that water has not frozen onto the detector surface.

4.6.4 Additional DSSD dead layers

In addition to the primary 100 nm Si dead layer of the DSSDs, there are 3 additional dead layers that contribute to the energy spectrum at the few percent level (see section 4.4): a layer formed from a thicker p+ implant region at the periphery of the front strips, an Al electrode on top of this thicker region, and an insulating SiO_2 region between the front strips. To determine the thickness of the additional dead layers, a similar approach as the primary dead layer was attempted, by looking at the change in peak position as a function of angle. However, this method proved unreliable due to overlap in peak positions and possible peak

broadening that affect the location of the peaks. Slopes from the determined peak positions corresponded to unrealistically large dead layers.

Instead, using the results for the primary dead layer (section 4.6.3), the thickness of the additional dead layers is determined directly from the energy spectrum of α particles that hit a single pixel (no charge-sharing) with the resolution enhancement from averaging the energy of the front and back strips. Using a complete detector response function (section 4.2), 8 parameters were simultaneously fit to a spectrum: (1) the beam energy, (2) the electronic noise σ , (3)-(5) the dead layer thicknesses, and (6)-(8) the fraction of events through each additional dead layer. The primary dead layer was taken as 100 nm Si. In addition, all dead layers were taken to have 85 nm of H₂O on top of them. The beam was assumed to be monoenergetic; some beam broadening may be contained in the electronic noise σ parameter since this parameter broadens the detector response with a Gaussian of mean zero.

Figure 4.26 shows the energy spectrum for 1550.3 ± 2.5 keV α particles showing the contributions from each of the different dead layers. There is generally good agreement between the detector response model and the data. The tails of the spectrum are not reproduced as well as other features; this could be due to some energy tail in the beam or deviations in the dead layers. In any case, these tail inaccuracies make up a very small amount of the response and the most important features are reproduced well.

As with the primary dead layer, the nominal 1500 keV data set was primarily used for the analysis, which was then validated with the remaining energies. This is due to the large separation between the peaks resulting from the different dead layers and the fact that this energy corresponds to the ^8Li and ^8B α peak, so it is most important to have a good characterization of the DSSD at this energy. At higher energies, peak separation was not as clear, particularly for the periphery p+ implant. An identical procedure at 1200 keV nominal beam energy resulted in dead layer thicknesses within the statistical uncertainty as the 1500 keV beam energy.

Fits were performed for each pixel at each angle, as well as a combined, simultaneous fit




Figure 4.26: Energy spectrum for 1552.0 ± 1.3 keV α particles that deposited charge in a single front strip and single back strip (no charge-sharing) showing the contributions from each of the different dead layers. All layers were also assumed to have 85 nm of H₂O on top of them. This spectrum is from a single pixel.

across all angles. All parameters were assumed to be identical across each angle but were allowed to vary between pixels. This accounts for possible imperfections and differences between pixels (dust, for example), and can provide an estimate of the uncertainty. This treatment is compared with performing the independent fit for each angle as shown in Figure 4.27, for the periphery Si p+ implant, the Al electrodes, and the SiO₂ insulating gap. There is generally good agreement between the independently-fitted colored points and the simultaneous fit to all angles. The four pixels were treated independently and there is also good agreement among them. The periphery Si p+ implant is determined to have a thickness of 263.5 ± 5.8 nm, and the Al electrode has a thickness of 273.7 ± 6.9 nm. These results agree reasonably well with those obtained in Ref. [41], which suggests that the treatment of frozen H₂O was performed accurately.

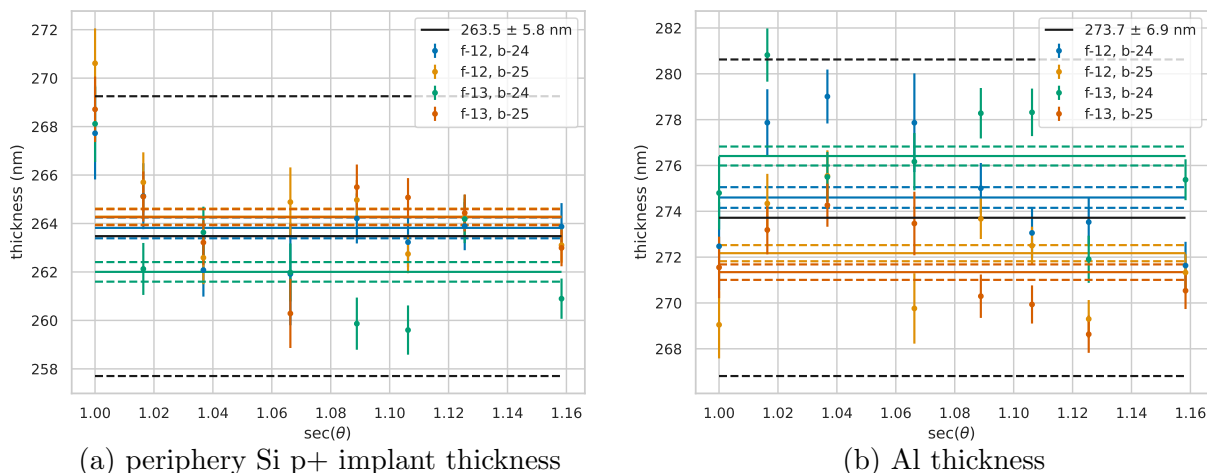


Figure 4.27: Thickness determination for (a) the periphery Si layer and (b) the Al layer using the nominal 1500 keV α beam, assuming 85 nm of H₂O on top. Each color represents a different pixel. The points were fit independently, and colored horizontal lines show the result for a simultaneous fit across the different angles; uncertainties on these are statistical only. Black horizontal lines show the average and total uncertainty.

The SiO₂ interstrip gap is the primary site of front strip charge sharing, as observed in section 4.6.6). Due to noise thresholds, the single pixel spectrum includes some events that share only a small amount of charge with a neighboring front strip. This is likely responsible for the extended tail of the SiO₂ peak seen in Figure 4.26. Using the single pixel spectrum

to determine the thickness of the SiO₂ dead layer therefore gives a biased estimate towards thicker dead layers. Instead, the back strip energies of the front strip charge sharing events can be used to measure the thickness of the SiO₂ layer. Figure 4.37a shows examples of these spectra where the SiO₂ dead layer is clearly and cleanly visible. Using these spectra to determine the thickness, results for both the March and January data sets are shown in Figure 4.28. Splines are produced from each set of data and the average is taken, as shown in black; uncertainties are estimated by bootstrapping.

There is a clear energy dependence on the determined thickness of the SiO₂; an angle dependence was also observed at the level of about 2%. The physical dimension of the SiO₂ layer was measured to be 650 nm (section 4.4.2). There is some agreement between the measured physical dimension at the 8500 keV α energy, but at lower energies, there is up to a 10% disagreement with the measured dimension. Some reasons for the disagreement could include an incorrect stopping power or density of SiO₂. It may also be the case that there is some electron trapping near this interface that would appear as an increased dead layer; the specific amount of this charge trapping would depend on the shape of the electric field and the interplay with front strip charge sharing and is difficult to estimate.

The thickness of the periphery p+ implant can also be determined from the back energies of front strip charge sharing events. This is only possible at the higher α energies (see section 4.6.6), but the values obtained agree well with the thickness determined from the single pixel spectra, as shown in Figure 4.29.

For all parameters, the statistical uncertainty on their determination is smaller than the standard deviation of the results for the four pixels of interest. Therefore, the sample standard deviation is taken as the uncertainty instead, to account for possible differences between the pixels across the detector. The systematic uncertainties on the 100 nm primary Si dead layer and the 85 nm H₂O dead layer are anti-correlated; to account for this, the uncertainty on the Si-equivalent dead layer of 3.3 nm was used instead (section 4.6.3). The uncertainty on the primary dead layer was added linearly, scaled for the difference in stopping




Figure 4.28: Determination of the SiO_2 dead layer thickness from the back energies of front strip charge-sharing. Uncertainties include statistical uncertainty and the systematic uncertainty from the primary dead layer. A spline (black) is produced from the average with uncertainties estimated by bootstrapping.




Figure 4.29: Thickness of the periphery p+ implant from the back energies of front strip charge-sharing. Results agree well with the thickness determined from the single pixel spectra at 1500 keV nominal beam energy.

power of the other materials, since the measurement of the other dead layers is relative to the primary dead layer. Uncertainty on the angle was taken to be negligible as the 0.9° uncertainty would affect the effective thickness of the largest angle point by about 1% and the other points by much less. As noted in section 4.6, uncertainties on the stopping powers or material densities are not included, as this analysis is only interested in reproducing the detector response.

The fraction of events through each layer is determined slightly differently depending on each layer. The Al electrode dead layer is measured using only the fits from this section. The measurements of the SiO_2 and periphery p+ implant layers depend on the results of section 4.6.6.

For the Al electrode layer, a fit at each beam energy is performed. This is similar to the fitting procedure described above, but includes a limited number of parameters, excluding dead layers that do not have sufficient separation in the spectrum. This gives a fairly

consistent result for the fraction of events passing through the Al electrode. A weighted average of the points gives 2.92 ± 0.06 % of events passing through the Al dead layer, as shown in Figure 4.30. This corresponds to a width of 29.2 ± 0.6 μm . All of these events deposit energy in a single front strip. The 8500 keV point is likely biased due to a lower energy separation between the primary dead layer and the Al electrode dead layer, but it has little impact on the weighted average.




Figure 4.30: Fraction of events passing through the Al electrode dead layer.

As discussed in section 4.6.6, events where an α passes through the SiO_2 layer are split between single pixel events and front strip charge-sharing events. (Back strip charge sharing events have the same response as for single pixel events.) For every energy, the fraction of events that are single pixel and front strip charge sharing which pass through the SiO_2 layer are added. This gives a nearly constant result across all energies, as shown in Figure 4.31a. From a weighted average, the fraction of all events passing through the SiO_2 layer is 1.23 ± 0.04 %. This corresponds to a width of 24.6 ± 0.8 μm , which closely matches the physical dimensions. As discussed in section 4.6.6, the increasing fraction of front strip charge sharing

events from the periphery p+ implant indicates that the maximum fraction of these events from the SiO₂ layer is reached at the higher energies. The fraction of events from this layer that have a single pixel response is then determined by subtraction from this maximum at the lower energies. Negative fractions are not allowed. The result of this procedure is shown in Figure 4.31b. Note this procedure uses only the results of Figure 4.40. The points shown are from the fits of the spectra performed previously in this section and agree well, validating the procedure and the total fraction of the DSSD surface covered by SiO₂.




Figure 4.31: (a) Fraction of all events passing through the SiO₂ layer. (b) Fraction of all events passing through the SiO₂ layer that deposit energy in a single front strip. Compare to Figure 4.40 and see text for details.

To determine the fraction of events passing through periphery p+ implant only (not including those events which pass through both p+ implant and the Al electrode), a similar procedure is used. In contrast to the SiO₂ layer, it cannot be assumed that the maximum fraction of front strip charge sharing events of the periphery p+ implant is reached (Figure 4.40). Instead, since no front strip charge sharing events from the periphery p+ implant are recorded at the lower energies, the fraction of these events in the single pixel spectrum is taken as the maximum fraction. A weighted average of these three points gives $2.14 \pm 0.07\%$ of events that pass through the periphery p+ implant. This corresponds to a total width of $50.6 \pm 0.9 \mu\text{m}$ for the periphery p+ implant, accounting for the width of the Al

electrode. Combining this fraction of events with the results of Figure 4.40 for the periphery p+ implant gives the fraction of these events that deposit energy in a single pixel, as shown in Figure 4.32.




Figure 4.32: Fraction of events that pass through the periphery p+ layer and deposit energy in a single pixel. Compare to Figure 4.40 and see text for details.

A summary of the different event fractions is given in section 4.6.8.

4.6.5 Back strip charge-sharing

Back strip charge-sharing is less complicated than front strip charge-sharing and is presented first. These events are labeled f1b2, as noted in Table 4.5. Events where an α deposited energy in two back strips were assigned to the pixel where more energy was deposited.

A comparison of the summed back energy spectrum with the single pixel spectrum is shown in Figure 4.33. Note that the front energies follow a nearly identical distribution between these event types, but that some charge is lost on the back strips when charge-sharing occurs. Some loss is expected, as the insulating interstrip gap between back strips

can trap some drifting electrons.⁷ The average charge loss from back strip charge-sharing is shown in Figure 4.34. The points in this figure have been corrected with the calibration disagreements between front and back strips shown in Figure 4.23. For use in simulations later, a simple line is fit to the average of the individual pixels, with uncertainties estimated from the sample standard deviation to account for potential variations between the strips.




Figure 4.33: Spectrum of back strip charge-sharing (f1b2) events compared to single pixel events for the nominal 1500 keV α beam. These spectra are the sum of the four pixels of interest.

Figure 4.35 shows the back strip energy spectrum of a single for charge-sharing events, (prior to the event-by-event summation shown in Figure 4.33). Note that the peaks at zero energy and close to the maximum energy; events near the maximum deposited most of their energy in the back strip of interest and almost zero energy in the neighboring back strip. It was found that a beta distribution modeled the charge-sharing distribution well. For an alpha of given energy E' (after accounting for energy loss from dead layers and NIEL), the amount of energy shared with a strip E is then

⁷See Ref. [166] for simulated field lines on back strips of a similar DSSD.




Figure 4.34: Average charge loss in back strip charge-sharing (f1b2) events as a function of α energy, corrected for calibration disagreements between individual front and back strips (see Figure 4.23). A simple line (in black) is fit to the mean of the individual pixels, with uncertainties shown (dashed).

$$f(x; \alpha, \beta) = \frac{1}{B(\alpha, \beta)} x^{\alpha-1} (1-x)^{\beta-1}, \quad (4.7)$$

where x is E/E' , B is the beta function, and $\alpha = \beta = 0.5$.⁸ The distribution is convolved over the range of deposited energies from the α beam through the various dead layers. Corrections are made to account for the charge-sharing energy loss (Figure 4.34) and a 20 keV noise cut. Finally, the model is broadened by the electronic noise and Fano factor statistics (sections 4.2.3 and 4.2.4).




Figure 4.35: Back strip energy spectrum of a single pixel for back strip charge-sharing events with (a) nominal 1500 keV α and (b) nominal 8500 keV α . (Note that this spectrum is the summation of several pixels to increase statistics.) A model based on the beta distribution (Eq. 4.7) is shown in orange; see text for details.

The fraction of events that shared charge between two back strips was found to be independent of angle, although a strong energy dependence was observed. Figure 4.36 shows the back strip charge-sharing fraction as a function of beam energy. A smoothed spline through the averages of these points is shown in black which is used in the detector response model to estimate the charge-sharing fraction at other α energies. Uncertainties the fraction of events are relatively small, but their effect on the spline is shown by the dashed lines. Uncertainty on the beam energy is negligibly small.

⁸The beta distribution is defined on the interval $[0,1]$, hence the definition of x .




Figure 4.36: Back strip charge sharing fraction; a smoothed spline in black is overlaid.

Between the nominal 1200 keV and 8500 keV α beams, there is about a factor of 2 increase in the fraction of events that experience back strip charge-sharing. Such a large increase was somewhat unexpected and difficult to explain accurately without detailed detector simulations that are beyond the scope of this work. A likely contribution comes from the increase in range of an 8500 keV α (50 μm) compared to a 1200 keV α (4 μm). Additionally, there is an increase in the lateral straggling (1.5 μm from 0.3 μm). Given the increased range, increased lateral straggling, and the greater number of e-h pairs created in the DSSD, it is reasonable that the fraction of charge-sharing should increase, though it is difficult to quantify the increase.

4.6.6 Front strip charge-sharing

Because α particles penetrate the interstrip gap of the DSSD, the charge-sharing response naturally depends the energy of the α due to the stopping-power dependence on energy. Events where energy was deposited in two front strips and one back strip are labeled f2b1,

as noted in Table 4.5.




Figure 4.37: Spectrum of front strip charge-sharing (f2b1) events compared to single pixel events for (a) the nominal 1500 keV α beam and (b) the nominal 7500 keV α beam. These spectra are the sum of the four pixels of interest.

A comparison of the front strip charge sharing summed energy spectrum with the single pixel spectrum is shown in Figure 4.37 for both the nominal 1500 keV and 7500 keV α beam. Note that the summed front strip energies have a lower average than the back strip energies of these events. This indicates an average charge loss, most likely a loss of holes that get trapped in the interstrip gap region. The average charge loss energy is shown in Figure 4.39. A simple line is fit to the data for use in simulations; this may not describe the data particularly well at the lower energies, but this is unlikely to have any meaningful effect. Similarly to back strip charge-sharing, a model based on the beta distribution (Eq. 4.7) describes the fraction of charge in each front strip well, as shown in Figure 4.38. Disagreements between the model and distribution are likely due to noise cuts.

The back strip energies in the 1500 keV α spectrum (Figure 4.37a) correspond exactly to the SiO_2 interstrip gap dead layer, as compared to the single pixel spectrum. In the 7500 keV spectrum, two peaks are evident and correspond to the SiO_2 interstrip gap and the thicker periphery p+ implant. This indicates that the charge-sharing fraction and the width of the charge-sharing region on the DSSD is energy-dependent. Specifically, it seems



Figure 4.38: Front strip energy spectrum of a single pixel for front strip charge-sharing events with nominal 8500 keV α . (Note that this spectrum is the summation of several pixels to increase statistics.) A model based on the beta distribution (Eq. 4.7) is shown in orange.

that the higher energy α particles, which have a longer range, deposit charge in a region of distorted field that exists underneath the periphery p+ implant at roughly 20 μm below the surface. The lower energy alphas do not penetrate as deeply, and the field above this point would fully collect the charge.

The relative contributions to front strip charge sharing events from the SiO_2 and the periphery p+ implant can be extracted from the f1b2 back energy spectra. These contributions are shown in Figure 4.40. The 3500 keV points were taken from January data. A logistic model was used to describe the SiO_2 , with uncertainties estimated from bootstrapping. Interpolation was used to describe the periphery p+ implant; uncertainties were also estimated from bootstrapping. The contribution from the SiO_2 layer rises to a maximum at the higher energies. From section 4.6.4, the maximum fraction was determined to be $1.23 \pm 0.04\%$. At the higher energies, contributions from the periphery p+ implant begin to appear. Unfortunately, no beam energies were taken between 3500 - 6500 keV, which appears to be




Figure 4.39: Average charge loss in front strip charge-sharing (f2b1) events as a function of α energy, corrected for calibration disagreements between individual front and back strips (see Figure 4.23). A simple line (in black) is fit to the mean of the individual pixels, with uncertainties shown (dashed).

where the periphery p+ implant begins to contribute. Looking at the ^{244}Cm calibration peak, it seems that there is a contribution from the periphery p+ region at its energy 5.8 MeV, but there are not enough statistics to make a reliable determination.




Figure 4.40: Front strip charge sharing fraction of events by dead layer. Circles show the total front strip charge sharing; squares and circles show contributions from the SiO_2 layer and the periphery p+ implant. Lines show the adopted models for these contributions, with uncertainties estimated in dashed lines. The 3.5 MeV points were taken from January data.

The energy-dependence of the front interstrip charge-sharing was somewhat investigated in Ref. [166], which used 1700 keV and 6000 keV protons to study interstrip gap effects. It was found that the effective interstrip gap width *decreased* at the higher energy, which is the opposite of the finding here. There are a few relevant differences between these two studies, as follows:

- 1) Ref. [166] used a model W1 DSSD which has Al metallization over the entire front surface, while this work studies a BB7 DSSD that has only a periphery Al metallization. Notably, in the BB7 studied here, there is a 10 μm gap between the SiO_2 gap and the Al strip (see Figure 4.11). It is possible that the field lines near the interstrip gap are

different than those of Ref. [166], which could lead to this energy-dependent effect. It is beyond the scope of this work to model the electric field lines in the BB7 DSSD, but this may be the best explanation.

- 2) The DSSD in Ref. [166] was presumably at room temperature, while this study was collected at ~ 80 K. The hole mobility at 77 K is roughly 22 times higher than at room temperature [155], which could possibly contribute some effect. However, Figure 5 of Ref. [166] indicates that the effective interstrip gap width decreased at higher bias voltages, which also increases the drift velocity, suggesting that this may actually reduce the size of the effect.
- 3) The projected range of 1700 keV protons is 35 μm and for 6000 keV protons is 295 μm . The α range is between about 5 μm at 1200 keV to about 50 μm at 8500 keV (Fig. 4.2b). The creation of the e-h pairs along the particle track and the local electric field at the creation sites is very different, particularly at the higher proton energies, which penetrate much more deeply.

4.6.7 Silicon response non-linearity

The nominal beam energies and those determined from the DSSD are shown in Table 4.4. The nominal beam energies were determined by measuring the time of flight of the 7500 keV α beam and then scaling the magnetic field of a 22° bending magnet and a 17° bending magnet along the beam line for the other energies. During the March experiment, the time of flight was measured for the 8500 keV beam; however, shortly after this measurements, problems with both the DSSD and the beam line appeared, necessitating a two day downtime to fix the issues and replace the DSSD. Therefore, the time of flight for the 8500 keV beam may not be accurate and is not used.

The time of flight was measured for the 7500 keV α beam but was not measured for lower energies due to time constraints and a the time structure of the beam at the lower energies.

The time of flight measurement gave an energy of 7.50 ± 0.02 MeV. Using the recorded magnetic fields and scaling the energy from the time of flight, a comparison between the nominal energy and the DSSD energy is made. Figure 4.41 shows the determined beam energies from the DSSD and the nominal beam energies from the magnet scaling. The uncertainties shown include those from the DSSD and its calibration, but the largest contribution by far comes from the time of flight; this contribution to the uncertainty includes only the 20 keV uncertainty on the 7500 keV energy and the resulting uncertainty from the magnetic field scaling.




Figure 4.41: Energy difference between nominal beam energy and energy determined from the DSSD. See text for details about uncertainties.

There is clearly some discrepancy between the nominal energies, but its origin is not clear. The size of a non-linearity due to the silicon response is several times larger than expected (Figure 4.6). It is possible that some non-linearity (not due to the intrinsic silicon response) is missing in the calibration of the DSSD. Due to a noise issue (and limited time to fix it) during the experiment, it was not possible to use a precision pulser to check the linearity of the DAQ. Additionally, it is not clear what uncertainty to assign the nominal beam energies;

due to experimental time constraints, only a single beam energy was measured via time of flight and magnet settings scaled from this reading. Note also that Figure 4.41 includes only the uncertainty from the magnet scaling based on the measured time of flight uncertainty of the 7500 keV α beam. Depending on the exact beam path through the steering elements, it is possible that the beam energy may have been slightly shifted from its nominal value.

Due to these difficulties in interpretation, it is not possible to make a definitive statement about this possible silicon non-linearity from this experiment. Therefore, throughout the analysis of the α beam experiment, the beam energy as determined by the DSSD is used.

4.6.8 Summary of results

In this work, the DSSDs used in the BPT experiment have been fully characterized at α energies between 1200 keV and 8500 keV. A complete detector response model has been created, including all charge-sharing, which is found to be highly energy-dependent. Angle dependence of the response was investigated and determined to be well modeled by the usual $\sec(\theta)$ path length increase at all angles relevant to the BPT.

A summary of the DSSD dimensions measured in this experiment are presented in Table 4.3. Overall, the dimensions were found to be in excellent agreement with the manufacturer specifications and with prior measurements with physical and optical probes. The periphery p+ implant was found to be considerably thinner than manufacturer specifications, though this could be due to incomplete charge collection instead of no charge collection in this area, which would result in a thinner apparent dead layer.

Figure 4.42 provides a summary of the fraction of events passing through the different dead layers as a function of α energy and whether charge is deposited in a single front strip or two front strips. The remaining events pass through the entrance window dead layer of 100 ± 3.5 nm. Figure 4.36 shows the fraction of events undergoing back strip charge sharing as a function of α energy. The average charge loss in these events is shown in Figure 4.34. Figure 4.39 shows the average charge loss for front strip charge sharing events.




Figure 4.42: Event fractions for the different dead layers and the number of front strips in which charge is deposited.

This experiment was not able to measure the silicon response non-linearity due to uncertainties on the α beam energies (section 4.6.7). In addition, opposite polarity front strip charge sharing events were only observed when a particular back strip was hit, an unexpected behavior that might indicate a DAQ issue. However, given that the dimensions of the SiO₂ agree very well with manufacturer dimensions, it seems unlikely that these events are “missing.” Additional study of this particular DSSD model and these opposite polarity events may be required in the future.

Chapter 5

CONCLUSIONS AND FUTURE DIRECTIONS

This work has contributed several important experimental improvements to the measurement of $\beta - \nu$ angular correlation in ${}^8\text{Li}$ (and ${}^8\text{B}$). With reference to Figure 1.17 and the uncertainty budget in Table 1.2, this work has reduced the systematic uncertainties associated with both β -scattering and the detector response. The new BPT Mk IV trap (chapter 2) has successfully demonstrated a new rod design with graphite and glassy carbon materials that reduce β scattering by a factor of 4 compared to the previous BPT, while also reducing the required RF voltage. These are important experimental demonstrations which are of use to other ion traps for which the minimization of β scattering is important. A complete DSSD characterization has been performed (chapter 4) which has reduced detector uncertainties and identified previously unknown detector effects. Experimental determinations of the recoil order form factors will require use of data across the broad range of α energies. These measurements will also require combinations of observables other than the α energy difference, which is less susceptible to detector uncertainties. The identification of several energy-dependent detector effects is therefore of extreme importance to accurately and precisely measure these recoil-order terms.

Finally, the new, high-statistics ${}^8\text{Li}$ data set collected with the BPT Mk IV (chapter 3) will enable an improved limit to be set on electroweak tensor couplings, as well as a global fit with the recoil-order terms to be investigated. The analysis of this blinded data set is on-going, with results expected in the next couple of years. Coupled with the experimental advancements, this data set will allow for improved sensitivity to a tensor contribution once data analysis is completed.

In addition to the experimental improvements, this work has improved our collaboration's understanding of the correlation between the angular correlation coefficient a and the Fierz interference term b , as shown in section 1.4. The future sensitivity of the BPT Mk IV using ${}^8\text{Li}$ and ${}^8\text{B}$ has also been investigated, showing great promise for the next generation

of experiments using this technique. In the near future, the BPT should be able to reach an uncertainty of $\Delta|C_T/C_A|^2 \leq 10^{-3}$, though the theoretical understanding of the nuclear structure still requires improvement.

5.1 Future directions

There are several opportunities to improve limits on tensor couplings to both left- and right-handed neutrinos from measurements of β -decay spectra. For the next generation of experiments, the measurement goal is to reach an uncertainty on the Fierz interference term of $\Delta b = 1 \times 10^{-3}$, corresponding to $\Delta C_T^+/C_A^+ = 2 \times 10^{-3}$. This is an ambitious goal that will require many experimental and theoretical improvements.

Several experiments have or are planning to measure isotopes that undergo β -delayed particle decays, as these have an enhanced sensitivity (section 1.3.1). As shown in section 1.4, these experiments will need to take advantage of correlations between the angular correlation coefficient a and the Fierz interference term b to correctly interpret these measurements. In addition, it was shown that analyzing the 2D spectrum, either $(E_e, \theta_{\beta\nu})$ or (E_e, p_r) , provides maximum sensitivity for the decay, which is important due to the time and effort required to perform a measurement at the level of precision that the next generation of experiments is designed for.

For the BPT to reach this level of uncertainty will require considerable effort, but it is achievable. The decays of ${}^8\text{Li}$ and ${}^8\text{B}$ are disadvantaged by their large Q values when compared to isotopes like the neutron and ${}^6\text{He}$. However, the delayed α from these decays improves the sensitivity by a factor of 3, bringing them to a level competitive with other decays, particularly when systematic uncertainties are considered. Given the other bounds on left-handed tensor couplings from other observables [12] and the plans for future measurements with other isotopes, it may be that the best motivation for future measurements of ${}^8\text{Li}$ and ${}^8\text{B}$ would be to improve limits on right-handed tensor couplings. From the current tension in the global fit for C_T^- , this is a powerful motivation.

5.2 Experiment upgrades and variations

Though trapping ^8Li and ^8B has advantages for the measurement, using a Paul trap comes with some disadvantages. A key disadvantage, explored in depth in this work, is the RF pickup on the DSSDs. This worsens the detector resolution and requires mitigation strategies that eliminate the possibility of using time of flight to perform particle identification (section 3.4.2).

It may be possible to upgrade the DAQ to a digital acquisition system, which has been explored to some degree though is not reported here. Digital acquisition offers the significant advantage of waveform recording, allowing for offline analysis to fit and subtract off the RF pickup in a straightforward way. This should improve the resolution of the detectors and allow lower energy signals to be collected. However, any digital DAQ will require an appropriate filtering algorithm at the trigger level to avoid constant triggering on RF pickup. This is possible but requires filtering times on the order of μs that are longer than the usual filtering times for digital acquisition systems, though this is not thought to be a significant problem.

As noted in chapter 3, during the latest ^8Li data run as well as previous ^8Li and ^8B runs, it seems that the rate of ions produced has essentially been maximized for the current system. This is largely due to charge saturation of the gas catcher in the case of ^8Li . In the case of ^8B , it may be that additional water molecules are needed to transport the ions. Modifications to the system could be made. One idea in the case of ^8B is to use a water-permeable nafion tubing [134] to humidify the helium gas supply of the gas catcher at very small level.

One of the other challenges with the gas catcher is that the isotopes of interest exit primarily as molecules over a wide range of masses (section 3.1). Since the BPT can trap only a small mass range at a time during an experiment, some of the activity (up to $\sim 50\%$) must be sacrificed to select a particular mass peak. In principle, a trap is not required to

study the decays of ${}^8\text{Li}$ and ${}^8\text{B}$ with sufficient precision for this measurement. These decays break up into two α particles of $\sim\text{MeV}$ energies, which are high energy enough to escape from a very thin foil with minimal energy deposition. One solution to this problem may therefore be to switch from using a Paul trap to implanting in a thin film, since implantation is mass-independent and therefore all ions could be implanted. To contribute negligibly to the uncertainty of the measurement, such a film should be below 100 nm thick, as this is roughly the thickness of the DSSD dead layer.

In recent years, thin, free-standing polymer films have been developed with thickness of < 10 nm and areas of a few cm^2 [169]. Adapting these thin films for implantation would solve a number of problems for the BPT. For one, ions of any mass could be implanted; this would improve the statistics for ${}^8\text{B}$ particularly, required for an improvement to the determination of the unoscillated neutrino spectrum. Furthermore, implantation techniques would contribute no RF pickup to the detector response, dramatically simplifying the required electronics; this should improve the energy and timing resolution. Technical challenges remain to be worked out, such as coating these polymer films in a conductive material, perhaps by atomic layer deposition. In addition, there is a possibility for increase β scattering when using a thin film, which must be estimated via simulation to understand the feasibility of performing such a measurement.

Alternatively, it may also be possible to move the experiment to a facility with much higher ${}^8\text{Li}$ production, such as TRIUMF.

In addition to ${}^8\text{Li}$ and ${}^8\text{B}$, another isotope of interest for the BPT is ${}^{20}\text{Na}$. This isotope is a β -delayed α emitter with the same spin sequence and hence same factor of 3 enhancement but has not been used to search for exotic couplings before. Compared to ${}^8\text{Li}$ and ${}^8\text{B}$, ${}^{20}\text{Na}$ has a lower Q value to the state of interest at ~ 6.5 MeV, making it more competitive with other isotopes for measurements of left-handed neutrino tensor couplings C_T^+ . In addition, ${}^{20}\text{Na}$ contains both F and GT decays, making it possible to search for both S and T couplings (though it would be challenging to collect enough statistics to be competitive with limits on

S couplings from other decays). ^{20}Na also decays to a sharply defined state, as opposed to ^8Li and ^8B , and would not suffer from the uncertainty over an intruder state or state-dependent recoil-order form factors. These recoil terms are likely also smaller for ^{20}Na . A Paul trap will be required to detect all of the decay products, though, since the final ^{16}O nucleus is too massive to escape from a thin film without significant energy loss.

Appendices

Appendix A

AVAILABILITY OF PRODUCTS FROM THIS WORK

All code and scripts are stored in a private `gitlab.phy.anl.gov` repository and are available upon request. Additionally, CAD models and prints of all components of the BPT Mk IV are available upon request.

Appendix B

EXPECTED SENSITIVITIES USING THE ASIMOV DATA SET

The results of a particular experiment (often expressed as confidence limits on a parameter) depend on the exact data set obtained, including statistical fluctuations. It is therefore not possible to predict with certainty what the sensitivity of a given experiment will be. However, to give an idea and to plan future experiments, the *median* expected sensitivity at some confidence level can be estimated, which gives a good sense of comparison between different experiments. Ref. [47] presents a method for determining the significance of a particular data set and/or for estimating the expected sensitivity of a future experiment.¹ Based on likelihood tests, this method uses two key features: 1) the asymptotic behavior of likelihood-based test statistics in the large N limit, and 2) a “median” data set called the Asimov data set.² This method has been employed widely in high-energy physics and is particularly valuable because it avoids computationally-expensive Monte Carlo simulations when evaluating sensitivity over a parameter space. However, the Asimov data set is not widely used in nuclear physics, and Ref. [47] is rather opaque in places (at least for this author).

In this appendix, the results of Ref. [47] are applied to a trivial example as a simple illustration of the method and to set up a framework for computing the results presented in section 1.4. Note that, in contrast to many searches for rare physics processes in high-energy physics, our measurement does not contain a typical background and signal process where the number of events is Poisson distributed for each. Instead, we are looking for a change in the shape of an energy spectrum compared to the SM case and the number of events is constant, a slightly simpler case. Additionally, we ignore any nuisance parameters.

We begin with a short rehashing of the method of Ref. [47]; formulae and some text are

¹Thank you to Michael D. Hank for pointing me to this reference.

²From Ref. [47]: the name of the Asimov data set is inspired by the short story “Franchise,” by Isaac Asimov [170]. In it, elections are held by selecting the single most representative voter to replace the entire electorate.

adapted from this reference. We have an experiment that measures a single parameter μ in some histogram with some number of counts. From the particular data of this experiment, we would like to determine the true value of μ (μ_t) and the associated uncertainty. This can be done by performing a likelihood test for a particular μ , where the likelihood ratio $\lambda(\mu)$, has an associated test statistic t_μ . In the cases presented in this work, we use the multinomial likelihood ratio (equation 40.16 of Ref. [2], with terms removed as indicated in the subsequent text). The likelihood ratio and associated test statistic are then

$$t_\mu = -2 \ln \lambda(\mu) = 2 \sum_{i=1}^I n_i \ln \frac{n_i}{m_i(\mu)} \quad (\text{B.1})$$

where I is the total number of bins in the histogram, n_i is the number of counts in bin i for the data, and $m_i(\mu)$ is the expected number of counts in bin i assuming a particular value for the parameter of interest μ . Using Equation B.1, one can therefore find the likelihood ratios and test statistics for all values of μ , given the particular data set. These in turn can be used to determine the most likely $\hat{\mu}$ of the underlying physics. The test statistic indicates the level of disagreement between the data and a hypothesized value of μ , with larger values indicating more disagreement. To quantify the level of disagreement, the p -value is computed

$$p_\mu = \int_{t_{\mu,obs}}^{\infty} f(t_\mu|\mu) dt_\mu, \quad (\text{B.2})$$

where $f(t_\mu|\mu)$ is the probability density function (PDF) of t_μ under the assumption of a particular value of μ , and $t_{\mu,obs}$ is the value of t_μ computed from the data at the particular μ being tested. To determine the confidence band for $\hat{\mu}$, the p -value for every μ can be computed, then those μ with $p < \alpha$ can be excluded. This would then be referred to as a $1 - \alpha$ CL. (Often, $\alpha = 0.05$, and this would then be the 95% CL.) We require, then, $f(t_\mu|\mu)$ to understand the distribution of the t_μ and to find the p -value.

With reference to the results of two earlier papers [171, 172], it can be shown that t_μ follows a non-central chi-square distribution (for one degree of freedom for a single parameter

of interest)

$$f(t_\mu; \Lambda) = \frac{1}{2\sqrt{t_\mu}} \frac{1}{\sqrt{2\pi}} \left[e^{-\frac{1}{2}(\sqrt{t_\mu} + \sqrt{\Lambda})^2} + e^{-\frac{1}{2}(\sqrt{t_\mu} - \sqrt{\Lambda})^2} \right] \quad (\text{B.3})$$

where the non-centrality parameter is $\Lambda = \frac{(\mu - \mu')^2}{\sigma^2}$ in the large N limit. Here, the data are distributed according to μ' so that $\hat{\mu} = \mu'$ and the uncertainty on $\hat{\mu}$ is given by σ . (This result can be generalized to additional parameters.) Now that we know the shape of the distribution of the test statistic t_μ , we need only estimate Λ to obtain the distribution directly. Note that when $\mu = \mu'$, this becomes a standard chi-square distribution.

This is where the Asimov data set comes in. The Asimov data set is defined as the data set “such that when one uses it to evaluate the estimators for all parameters, one obtains the true parameter values” [47]. For the cases considered here, the Asimov data set is easily found by taking μ' and using it to produce the PDF of the data directly (without any randomization); this is essentially the expectation value for each bin i assuming μ' . Then, using the Asimov data set to compute a likelihood ratio, we find

$$t_\mu = -2 \ln \lambda_A(\mu) \approx \frac{(\mu - \mu')^2}{\sigma^2} = \Lambda. \quad (\text{B.4})$$

Importantly, the Asimov data set provides the median value of the test statistic. If we are interested in the expected sensitivity of a planned experiment, typically taken at the median, then we only need to use the Asimov data set assuming some μ' , compute the test statistic t_μ at every μ , and find the p -values associated with the t_μ . This then allows us to exclude μ based on their median p -values, providing the expected sensitivity.

B.1 A simple example

It can be a little more instructive to consider this in the case of a hypothetical experiment where we would like to estimate the projected sensitivity to some parameter. Our experiment again measures a single parameter μ in some histogram with some number of counts. It has

some resolution $s \equiv 1$. In the large N limit and assuming no bias, the PDF of the experiment is then simply a Gaussian of mean μ_t and standard deviation s . Assume that the SM $\mu_t = 0$, and we would like to understand our median sensitivity to a BSM μ . Intuitively, our expected 95% CL sensitivity would be $\mu_t \pm 2s/\sqrt{N}$.

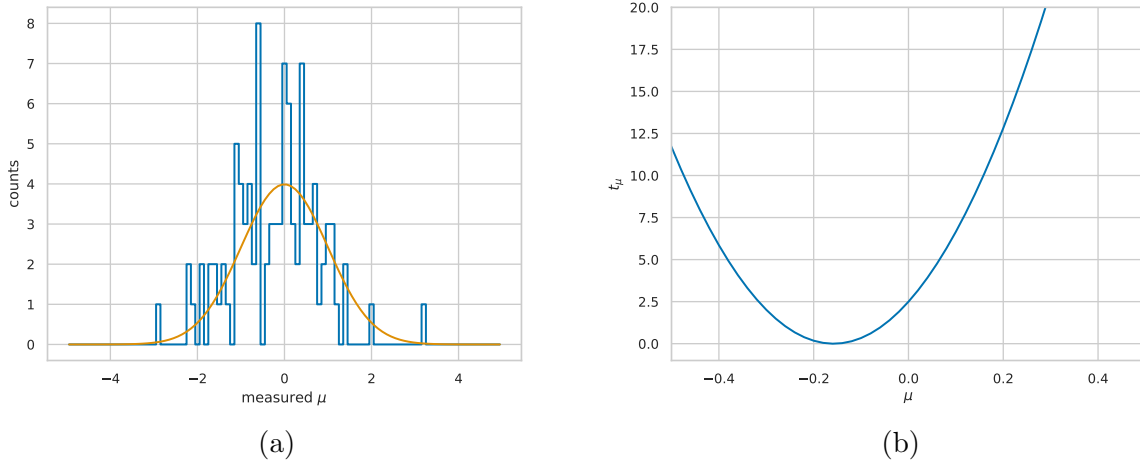


Figure B.1: (a) A typical single experiment in our Monte Carlo example. (b) The t_μ for this single experiment. Note that the $\hat{\mu}$ for this particular experiment is not at the SM value of zero.

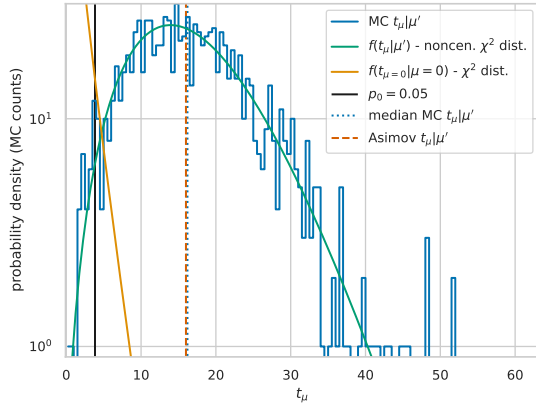
We test this in two ways: 1) by performing a Monte Carlo (MC), and 2) by using the Asimov data set. For the MC, let us run 1000 experiments of 100 counts each. A typical experiment might look like that in Figure B.1a, with the t_μ for this experiment given by Figure B.1b. By performing many experiments, we can build up $f(t_\mu|\mu)$ directly. Examples of $f(t_\mu|\mu)$ at several different values of μ are shown in Figure B.2. The MC results for t_μ are shown in blue and agree well with a non-central chi-square distribution (green). At $\mu = 0$, the distribution becomes the typical chi-squared distribution for a single degree of freedom (yellow); this is because here $\mu = \mu'$.

The median of the MC t_μ is found explicitly from the distribution (vertical blue dotted line). To determine the median 95% CL sensitivity, we compare the median MC t_μ to the $t_{\mu=0}$ whose $p = 0.05$. This t_μ is shown as the solid black line. Above this line, the integral of the yellow chi-square distribution contains 5% of its area. Essentially, this hypothesis test is


telling us whether the data agrees better with $\mu = 0$ (the SM) or with the particular μ under question at 95% CL. In Figure B.2, we can see that we would reject $\mu = -0.4$ but accept the other μ shown. The cut-off is at $\mu = -0.2$, where the t_μ is exactly on the $p = 0.05$ line (to within the accuracy of the simulation). This means that our median expected sensitivity at 95% CL would be $-0.2 < \mu < 0.2$. Our prediction for the range was $\pm 2s/\sqrt{N} = \pm 0.2$, exactly matching what we find by simulation (to within numerical fluctuations).

We now consider the second method using the Asimov data set. Here, we assume our data is exactly distributed according to the Gaussian PDF with mean $\mu' = \mu_t$, scaled by the number of events N ; this results in bins with fractional bin counts, which is fine for our purpose. The resulting distribution is shown in orange in Figure B.1a. Next, for every μ of interest, we compute the Asimov data set in the same way. We then compute the test statistic t_μ between the μ' Asimov data set and the μ Asimov data set. This immediately gives us the median t_μ and also the median p value. This can be seen in Figure B.2, where the Asimov t_μ is shown by the orange dashed line, which agrees very well with the median t_μ found from the MC. Moreover, the Asimov test statistic is equal to the non-centrality parameter Λ from equation B.4. This method is actually how the non-centrality parameter for the green distributions was computed, which is seen to perfectly match the MC distribution. Following the same comparison above for the desired p value, using the Asimov test statistics results in the same median expected sensitivity at 95% CL of $-0.2 < \mu < 0.2$, in line with our prediction.


Therefore, all one needs to do is to compute the Asimov test statistic and compute its p value using the standard chi-square distribution with the appropriate number of degrees of freedom (equal to the number of parameters in question). Rejecting those μ below the desired α will result in the expected sensitivity at the $1 - \alpha$ CL. The power of this method comes from the fact that no Monte Carlos are required, only the computation of the test statistic at each point of interest in the parameter space, making this method much faster than performing many simulations, especially for higher-dimensional problems, where one




(a) $\mu = -0.4$



(b) $\mu = -0.2$



(c) $\mu = -0.1$



(d) $\mu = 0.0$

Figure B.2: Distributions of t_μ for different values of μ . Note that the distributions are identical for the positive μ (symmetric about $\mu = 0$). See text for explanation.

would like to estimate sensitivities for multiple parameters of interest.

As a brief example of a higher-dimensional case, we consider a similar problem as before. In this case, we are interested in measuring μ_1, μ_2 and assume that we have Gaussian resolution of $s \equiv 1$ on each independently. (The covariance matrix is the identity matrix.) Our experiment will measure 100 events. Using the Asimov method, we can quickly compute the expected sensitivities at two different p -values of usual interest, 68% CL and 95% CL. We use the chi-squared distribution with two degrees of freedom to compute the p -values of the computed test statistics. The result is shown in Figure B.3. As before, this agrees with the expectation. In this plot, we have tested 60×60 parameter combinations. If this were computed using a Monte Carlo, we might expect to draw something like 360 million events to test the same parameter space, assuming 1000 experiments per point, and computed 3.6 million test statistics. Instead, we have drawn no events and computed merely 3600 test statistics using the Asimov treatment.




Figure B.3: Expected sensitivities at two different p -values using the Asimov treatment for estimating two independent, Gaussian-distributed parameters. See text for details.

Appendix C

MICROPARTICLE LINEAR QUADRUPOLE ION TRAP FOR DEMONSTRATIONS

Ion traps and related instruments have been used for more than a century (including the work of Paul and colleagues [173]). For the purposes of an undergraduate lecture or public demonstration, a linear quadrupole ion trap is a relatively simple, inexpensive device to build and can be operated easily and reliably, both by the presenter and the general public. Several different topics can be discussed related to charge, electrostatics and dynamics, and the applications of ion traps.

I collaborated with John A. Varriano (Christian Brothers University, Memphis, Tennessee, and my father) to build a demonstration ion trap based on a previously published design [174].¹ This small device traps *Lycopodium* (club moss) spores, which are about 20-30 μm in diameter and have a mass on the order of 10 ng [174]. Charge is given to the spores by the triboelectric effect: cloth is rubbed against a polytetrafluoroethylene (PTFE) rod which is then dipped into the spores. The spores can then be transferred from the rod into the trap.

The entire demonstration was fabricated for a few hundred dollars using a commercially-available high-voltage transformer. This transformer is a 6 kV, 30 mA current-limited transformer operated at 60 Hz, 115 VAC input, originally manufactured for neon signs or other gas displays [176]. The trap and associated equipment is shown in Figure C.1. A variac is used as the power input for the transformer and allows easy control of the voltage amplitude on the trap. A green laser pointer is used to illuminate the spores to make them easier to see by eye. The trap must be isolated from air currents which will disperse the light spores; a cardboard box with plastic viewing windows works well.

The linear quadrupole trap itself (Figure C.2) is made of four 1/8" stainless steel rods

¹A design for a demonstration surface ion trap has also been recently published [175], which may be of interest.




Figure C.1: The demonstration linear quadrupole ion trap and associated equipment.




Figure C.2: Spores (green dots) trapped in the microparticle linear quadrupole trap.

spaced 3/8" apart, with a length of about 6". A clear PTFE end piece serves as support for the rods and also allows the laser pointer optical access for illumination. Additionally, the PTFE end piece provides confinement along the axis of the trap, presumably by accumulating charge during the trap operation. Copper plates are placed above and below the quadrupole and can be biased to provide an electric field vertically. The bias on these plates can be used to raise and lower the entire ion cloud and/or to remove particular charge states.

The circuit diagram of the device is shown in Figure C.3. The large resistors serve as a safety measure to limit the current in the event of a short between the two AC outputs. Two optional DC connections are provided for biasing the top and bottom plates, if desired.

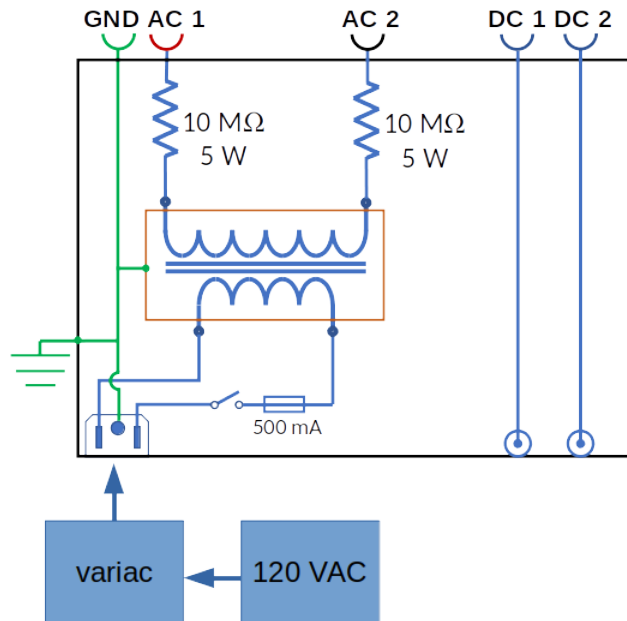


Figure C.3: Circuit diagram of the demonstration ion trap.

Figure C.2 shows a number of trapped spores, and Figure C.4 shows a single line of trapped spores. Audience members are able to change the trapping amplitude by adjusting the variac and can see the resulting effect on the ions. The extremely small size of the spores means that they are susceptible to air currents but also that drag forces from the air effectively “cool” their motion, resulting in structures that resemble Coulomb crystals. Additionally, it was observed that the spores seemed to induce air currents themselves at high

trapping amplitudes that cause large micromotions of the spores. This results in interesting vortex-like structures that do not exist in the typical ion traps used in atomic or nuclear physics. For audience members, this can be a great connection to hydrodynamics and soft matter physics, where the scale of the physics studied often necessitates understanding seemingly disordered systems.




Figure C.4: A single line of spores (green dots) trapped in the microparticle linear quadrupole trap. Smeared lines show the micromotion oscillation of the spores. The nut visible is roughly 8 mm wide.

Appendix D

COLLABORATION PICTURES

For posterity, pictures taken during the 2022 ^8Li data campaign of some of the collaboration members are included here. Collaborators for this experiment not pictured here include Nathan Callahan, Sergio Lopez, Graeme Morgan, Patrick O'Malley, Fabio Rivero, Gemma Wilson, and Rey Zite.



Figure D.1: From left: Jason A. Clark, Guy Savard, Aaron T. Gallant, Mary T. Burkey, Louis Varriano, and Nicholas D. Scielzo.



Figure D.2: From left: Louis Varriano, Bernhard Maaß, Jason A. Clark, Kay Kolos, Guy Savard, Daniel P. Burdette, Peter Müller, Adrian A. Valverde, W. Sam Porter, Aaron T. Gallant, and Zach Purcell.




Figure D.3: From left: Biying Liu, Adrian A. Valverde, Jason A. Clark, Maxime Brodeur, Brenden Longfellow, Dwaipayan Ray, and Alec Cannon.



Figure D.4: From left: Tsviki Hirsh and the author, looking extremely haggard after a month of taking data.

REFERENCES

- [1] J. F. Ziegler, M.D. Ziegler, and J.P. Biersack. “SRIM – The stopping and range of ions in matter (2010)”. en. In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 268.11–12 (June 2010), 1818–1823. ISSN: 0168583X. DOI: [10.1016/j.nimb.2010.02.091](https://doi.org/10.1016/j.nimb.2010.02.091).
- [2] R. L. Workman et al. “Review of Particle Physics”. In: *PTEP* 2022 (2022), p. 083C01. DOI: [10.1093/ptep/ptac097](https://doi.org/10.1093/ptep/ptac097).
- [3] B. Acharya et al. “Fundamental Symmetries, Neutrons, and Neutrinos (FSNN): Whitepaper for the 2023 NSAC Long Range Plan”. en. In: arXiv:2304.03451 (Apr. 2023). arXiv:2304.03451 [nucl-ex, physics:nucl-th]. URL: <http://arxiv.org/abs/2304.03451>.
- [4] T. D. Lee and C. N. Yang. “Question of Parity Conservation in Weak Interactions”. en. In: *Physical Review* 104.1 (Oct. 1956), 254–258. ISSN: 0031-899X. DOI: [10.1103/PhysRev.104.254](https://doi.org/10.1103/PhysRev.104.254).
- [5] C. S. Wu et al. “Experimental Test of Parity Conservation in Beta Decay”. en. In: *Physical Review* 105.4 (Feb. 1957), 1413–1415. ISSN: 0031-899X. DOI: [10.1103/PhysRev.105.1413](https://doi.org/10.1103/PhysRev.105.1413).
- [6] Allan Franklin. *Experiment, Right or Wrong*. Cambridge University Press, 1990. DOI: [10.1017/CB09780511527302](https://doi.org/10.1017/CB09780511527302).
- [7] Adam Falkowski, Martín González-Alonso, and Oscar Naviliat-Cuncic. “Comprehensive analysis of beta decays within and beyond the Standard Model”. en. In: *Journal of High Energy Physics* 2021.4 (Apr. 2021), p. 126. ISSN: 1029-8479. DOI: [10.1007/JHEP04\(2021\)126](https://doi.org/10.1007/JHEP04(2021)126).
- [8] J. D. Jackson, S. B. Treiman, and H. W. Wyld. “Possible Tests of Time Reversal Invariance in Beta Decay”. en. In: *Physical Review* 106.3 (May 1957), 517–521. ISSN: 0031-899X. DOI: [10.1103/PhysRev.106.517](https://doi.org/10.1103/PhysRev.106.517).
- [9] Martín González-Alonso, Oscar Naviliat-Cuncic, and Nathal Severijns. “New physics searches in nuclear and neutron β decay”. en. In: *Progress in Particle and Nuclear Physics* 104 (Jan. 2019). arXiv: 1803.08732, 165–223. ISSN: 01466410. DOI: [10.1016/j.pnnp.2018.08.002](https://doi.org/10.1016/j.pnnp.2018.08.002).
- [10] P. Delahaye et al. “The MORA project”. en. In: *Hyperfine Interactions* 240.1 (Dec. 2019), p. 63. ISSN: 0304-3843, 1572-9540. DOI: [10.1007/s10751-019-1611-x](https://doi.org/10.1007/s10751-019-1611-x).
- [11] P. Herczeg. “Beta decay beyond the standard model”. en. In: *Progress in Particle and Nuclear Physics* 46.2 (2001), 413–457. ISSN: 01466410. DOI: [10.1016/S0146-6410\(01\)00149-1](https://doi.org/10.1016/S0146-6410(01)00149-1).
- [12] K. K. Vos, H. W. Wilschut, and R. G. E. Timmermans. “Symmetry violations in nuclear and neutron β decay”. en. In: *Reviews of Modern Physics* 87.4 (Dec. 2015), 1483–1516. ISSN: 0034-6861, 1539-0756. DOI: [10.1103/RevModPhys.87.1483](https://doi.org/10.1103/RevModPhys.87.1483).
- [13] Oscar Naviliat-Cuncic and Martín González-Alonso. “Prospects for precision measurements in nuclear β decay in the LHC era”. en. In: *Annalen der Physik* 525.8–9 (Sept. 2013), 600–619. ISSN: 00033804. DOI: [10.1002/andp.201300072](https://doi.org/10.1002/andp.201300072).

- [14] Barry R Holstein. “Precision frontier in semileptonic weak interactions: introduction and overview”. In: *Journal of Physics G: Nuclear and Particle Physics* 41.11 (Oct. 2014), p. 110301. DOI: [10.1088/0954-3899/41/11/110301](https://doi.org/10.1088/0954-3899/41/11/110301). URL: <https://dx.doi.org/10.1088/0954-3899/41/11/110301>.
- [15] Ayala Glick-Magid et al. “Beta spectrum of unique first-forbidden decays as a novel test for fundamental symmetries”. en. In: *Physics Letters B* 767 (Apr. 2017), 285–288. ISSN: 03702693. DOI: [10.1016/j.physletb.2017.02.023](https://doi.org/10.1016/j.physletb.2017.02.023).
- [16] H Paul. “LEAST-SQUARES ADJUSTMENT OF THE COUPLING CONSTANTS”. en. In: *Nuclear Physics A* 154.1 (Sept. 1970), 160–176. DOI: [https://doi.org/10.1016/0375-9474\(70\)91077-8](https://doi.org/10.1016/0375-9474(70)91077-8).
- [17] M. González-Alonso and O. Naviliat-Cuncic. “Kinematic sensitivity to the Fierz term of β -decay differential spectra”. en. In: *Physical Review C* 94.3 (Sept. 2016), p. 035503. ISSN: 2469-9985, 2469-9993. DOI: [10.1103/PhysRevC.94.035503](https://doi.org/10.1103/PhysRevC.94.035503).
- [18] M. Beck et al. “Improved determination of the $\beta - \bar{\nu}_e$ angular correlation coefficient a in free neutron decay with the *a*SPECT spectrometer”. en. In: *Physical Review C* 101.5 (May 2020), p. 055506. ISSN: 2469-9985, 2469-9993. DOI: [10.1103/PhysRevC.101.055506](https://doi.org/10.1103/PhysRevC.101.055506).
- [19] J.D. Jackson, S.B. Treiman, and H.W. Wyld. “Coulomb corrections in allowed beta transitions”. en. In: *Nuclear Physics* 4 (Aug. 1957), 206–212. ISSN: 00295582. DOI: [10.1016/0029-5582\(87\)90019-8](https://doi.org/10.1016/0029-5582(87)90019-8).
- [20] M.E. Ebel and G. Feldman. “Further remarks on Coulomb corrections in allowed beta transitions”. en. In: *Nuclear Physics* 4 (Aug. 1957), 213–214. ISSN: 00295582. DOI: [10.1016/0029-5582\(87\)90020-4](https://doi.org/10.1016/0029-5582(87)90020-4).
- [21] Barry R. Holstein. “Recoil effects in allowed beta decay: The elementary particle approach”. en. In: *Reviews of Modern Physics* 46.4 (Oct. 1974), 789–814. ISSN: 0034-6861. DOI: [10.1103/RevModPhys.46.789](https://doi.org/10.1103/RevModPhys.46.789).
- [22] F. Glück. “Order- α radiative correction calculations for unoriented allowed nuclear, neutron and pion β decays”. en. In: *Computer Physics Communications* 101.3 (May 1997), 223–231. ISSN: 00104655. DOI: [10.1016/S0010-4655\(96\)00168-3](https://doi.org/10.1016/S0010-4655(96)00168-3).
- [23] Leendert Hayen and Albert R. Young. “Consistent description of angular correlations in β decay for Beyond Standard Model physics searches”. en. In: *arXiv:2009.11364 [nucl-ex, physics:nucl-th]* (Oct. 2020). arXiv: 2009.11364. URL: <http://arxiv.org/abs/2009.11364>.
- [24] Leendert Hayen. “Standard model $\mathcal{O}(\alpha)$ renormalization of g_A and its impact on new physics searches”. en. In: *Physical Review D* 103.11 (June 2021), p. 113001. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.103.113001](https://doi.org/10.1103/PhysRevD.103.113001).
- [25] O Kofoed-Hansen. “THEORETICAL ANGULAR CORRELATIONS IN ALLOWED BETA TRANSITIONS”. en. In: *Dan. Mat. Fys. Medd.* 28.9 (1954). URL: <http://publ.royalacademy.dk/books/414/2860?lang=da>.

- [26] A. T. Gallant et al. “Angular Correlations in the β Decay of ${}^8\text{B}$: First Tensor-Current Limits from a Mirror-Nucleus Pair”. en. In: *Physical Review Letters* 130.19 (May 2023), p. 192502. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.130.192502](https://doi.org/10.1103/PhysRevLett.130.192502).
- [27] C. H. Johnson, Frances Pleasonton, and T. A. Carlson. “Precision Measurement of the Recoil Energy Spectrum from the Decay of He^6 ”. en. In: *Physical Review* 132.3 (Nov. 1963), 1149–1165. ISSN: 0031-899X. DOI: [10.1103/PhysRev.132.1149](https://doi.org/10.1103/PhysRev.132.1149).
- [28] F. Glück. “Order- α radiative correction to ${}^6\text{He}$ and ${}^{32}\text{Ar}$ β decay recoil spectra”. en. In: *Nuclear Physics A* (1998).
- [29] E. G. Adelberger et al. “Positron-Neutrino Correlation in the $0^+ \rightarrow 0^+$ Decay of ${}^{32}\text{Ar}$ ”. en. In: *Physical Review Letters* 83.7 (Aug. 1999), 1299–1302. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.83.1299](https://doi.org/10.1103/PhysRevLett.83.1299).
- [30] A. Gorelov et al. “Scalar Interaction Limits from the $\beta - \nu$ Correlation of Trapped Radioactive Atoms”. en. In: *Physical Review Letters* 94.14 (Apr. 2005), p. 142501. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.94.142501](https://doi.org/10.1103/PhysRevLett.94.142501).
- [31] P. A. Vetter et al. “Measurement of the $\beta - \nu$ correlation of ${}^{21}\text{Na}$ using shakeoff electrons”. en. In: *Physical Review C* 77.3 (Mar. 2008), p. 035502. ISSN: 0556-2813, 1089-490X. DOI: [10.1103/PhysRevC.77.035502](https://doi.org/10.1103/PhysRevC.77.035502).
- [32] N. Scielzo et al. “Measurement of the $\beta - \nu$ Correlation using Magneto-optically Trapped ${}^{21}\text{Na}$ ”. en. In: *Physical Review Letters* 93.10 (Aug. 2004), p. 102501. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.93.102501](https://doi.org/10.1103/PhysRevLett.93.102501).
- [33] X Fléhard et al. “Measurement of the $\beta - \nu$ correlation coefficient $a_{\beta\nu}$ in the β decay of trapped ${}^6\text{He}^+$ ions”. en. In: *Journal of Physics G: Nuclear and Particle Physics* 38.5 (May 2011), p. 055101. ISSN: 0954-3899, 1361-6471. DOI: [10.1088/0954-3899/38/5/055101](https://doi.org/10.1088/0954-3899/38/5/055101).
- [34] G. Darius et al. “Measurement of the Electron-Antineutrino Angular Correlation in Neutron β Decay”. en. In: *Physical Review Letters* 119.4 (July 2017), p. 042502. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.119.042502](https://doi.org/10.1103/PhysRevLett.119.042502).
- [35] M. T. Hassan et al. “Measurement of the neutron decay electron-antineutrino angular correlation by the aCORN experiment”. en. In: *Physical Review C* 103.4 (Apr. 2021), p. 045502. ISSN: 2469-9985, 2469-9993. DOI: [10.1103/PhysRevC.103.045502](https://doi.org/10.1103/PhysRevC.103.045502).
- [36] G. Li et al. “Tensor Interaction Limit Derived From the α - β - $\bar{\nu}$ Correlation in Trapped ${}^8\text{Li}$ Ions”. en. In: *PHYSICAL REVIEW LETTERS* 110.9 (2013), p. 092502. DOI: [10.1103/PhysRevLett.110.092502](https://doi.org/10.1103/PhysRevLett.110.092502).
- [37] M. G. Sternberg et al. “Limit on Tensor Currents from ${}^8\text{Li}$ β Decay”. en. In: *Physical Review Letters* 115.18 (Oct. 2015). ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.115.182501](https://doi.org/10.1103/PhysRevLett.115.182501).
- [38] M. T. Burkey et al. “Improved Limit on Tensor Currents in the Weak Interaction from ${}^8\text{Li}$ β Decay”. en. In: *Physical Review Letters* 128.20 (May 2022), p. 202502. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.128.202502](https://doi.org/10.1103/PhysRevLett.128.202502).

- [39] P. Müller et al. “ β -Nuclear-Recoil Correlation from ${}^6\text{He}$ Decay in a Laser Trap”. en. In: *Physical Review Letters* 129.18 (Oct. 2022), p. 182502. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.129.182502](https://doi.org/10.1103/PhysRevLett.129.182502).
- [40] E.T.H. Clifford et al. “The decay of ${}^{20}\text{Na}$: Measurements of isospin mixing and the weak vector coupling constant as well as other new decay data”. en. In: *Nuclear Physics A* 493.2 (Mar. 1989), 293–322. ISSN: 03759474. DOI: [10.1016/0375-9474\(89\)90399-0](https://doi.org/10.1016/0375-9474(89)90399-0).
- [41] M. T. Burkey. “Searching for Tensor Currents in the Weak Interaction Using Lithium-8 β Decay”. PhD thesis. University of Chicago, 2019. DOI: [10.6082/uchicago.1697](https://doi.org/10.6082/uchicago.1697).
- [42] V Vorobel et al. “Investigation of $\beta - \nu$ angular correlation in the β^+ decay of ${}^{18}\text{Ne}$ and ${}^{14}\text{O}$ ”. en. In: *Czech. J. Phys.* (2002).
- [43] V. Vorobel et al. “Beta-neutrino angular correlation in the decay of ${}^{14}\text{O}$: Scalar coupling and interatomic interaction”. en. In: *The European Physical Journal A* 16.1 (Jan. 2003), 139–147. ISSN: 1434-6001, 1434-601X. DOI: [10.1140/epja/i2002-10067-y](https://doi.org/10.1140/epja/i2002-10067-y).
- [44] Praveen D. Shidling et al. “TAMUTRAP facility: Penning trap facility for weak interaction studies”. en. In: *Hyperfine Interactions* 240.1 (Dec. 2019), p. 40. ISSN: 0304-3843, 1572-9540. DOI: [10.1007/s10751-019-1584-9](https://doi.org/10.1007/s10751-019-1584-9).
- [45] V. Araujo-Escalona et al. “Simultaneous measurements of the β -neutrino angular correlation in ${}^{32}\text{Ar}$ pure Fermi and pure Gamow-Teller transitions using β -proton coincidences”. en. In: *Physical Review C* 101.5 (May 2020), p. 055501. ISSN: 2469-9985, 2469-9993. DOI: [10.1103/PhysRevC.101.055501](https://doi.org/10.1103/PhysRevC.101.055501).
- [46] W. Byron et al. “First observation of cyclotron radiation from MeV-scale e^{pm} following nuclear beta decay”. en. In: arXiv:2209.02870 (Sept. 2022). arXiv:2209.02870 [nucl-ex]. URL: <http://arxiv.org/abs/2209.02870>.
- [47] Glen Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. en. In: *The European Physical Journal C* 71.2 (Feb. 2011), p. 1554. ISSN: 1434-6044, 1434-6052. DOI: [10.1140/epjc/s10052-011-1554-0](https://doi.org/10.1140/epjc/s10052-011-1554-0).
- [48] W.J. Huang et al. “The AME 2020 atomic mass evaluation (I). Evaluation of input data, and adjustment procedures*”. en. In: *Chinese Physics C* 45.3 (Mar. 2021), p. 030002. ISSN: 1674-1137, 2058-6132. DOI: [10.1088/1674-1137/abddb0](https://doi.org/10.1088/1674-1137/abddb0).
- [49] F. C. Barker. “Delayed Alpha Spectra from the Beta Decay of ${}^8\text{Li}$ and ${}^8\text{B}$ ”. en. In: *Australian Journal of Physics* 42.1 (1989), p. 25. ISSN: 0004-9506. DOI: [10.1071/PH890025](https://doi.org/10.1071/PH890025).
- [50] R. B. Wiringa et al. “Charge-symmetry breaking forces and isospin mixing in ${}^8\text{Be}$ ”. en. In: *Physical Review C* 88.4 (Oct. 2013), p. 044333. ISSN: 0556-2813, 1089-490X. DOI: [10.1103/PhysRevC.88.044333](https://doi.org/10.1103/PhysRevC.88.044333).
- [51] F. C. Barker, H. J. Hay, and P. B. Treacy. “ 0^+ States of ${}^8\text{Be}$ ”. en. In: *Australian Journal of Physics* 21.3 (1968), p. 239. ISSN: 0004-9506. DOI: [10.1071/PH680239](https://doi.org/10.1071/PH680239).
- [52] F. C. Barker. “ 2^+ States of ${}^8\text{Be}$ ”. en. In: *Australian Journal of Physics* 22.3 (1969), p. 293. ISSN: 0004-9506. DOI: [10.1071/PH690293](https://doi.org/10.1071/PH690293).

- [53] Solveig Hyldegaard. “Beta-decay studies of ^8Be and ^{12}C ”. en. PhD thesis. Denmark: Aarhus University, 2010.
- [54] G. H. Sargsyan et al. “Impact of Clustering on the ^8Li β Decay and Recoil Form Factors”. en. In: *Physical Review Letters* 128.20 (May 2022), p. 202503. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.128.202503](https://doi.org/10.1103/PhysRevLett.128.202503).
- [55] R. E. Tribble and G. T. Garvey. “Induced weak currents and $\beta^\pm - \alpha$ angular correlations in $A = 8$ ”. en. In: *Physical Review C* 12.3 (Sept. 1975), 967–983. ISSN: 0556-2813. DOI: [10.1103/PhysRevC.12.967](https://doi.org/10.1103/PhysRevC.12.967).
- [56] Barry R. Holstein. “Electromagnetic effects and weak form factors”. en. In: *Physical Review C* 19.4 (Apr. 1979), 1467–1472. ISSN: 0556-2813. DOI: [10.1103/PhysRevC.19.1467](https://doi.org/10.1103/PhysRevC.19.1467).
- [57] R. D. McKeown, G. T. Garvey, and C. A. Gagliardi. “Beta-alpha angular correlations in mass 8”. en. In: *Physical Review C* 22.2 (Aug. 1980), 738–749. ISSN: 0556-2813. DOI: [10.1103/PhysRevC.22.738](https://doi.org/10.1103/PhysRevC.22.738).
- [58] T. Sumikama et al. “Test of the conserved vector current hypothesis by a β -ray angular distribution measurement in the mass-8 system”. en. In: *Physical Review C* 83.6 (June 2011), p. 065501. ISSN: 0556-2813, 1089-490X. DOI: [10.1103/PhysRevC.83.065501](https://doi.org/10.1103/PhysRevC.83.065501).
- [59] F. C. Gilbert. “Range Distribution of Alpha Particles Following the Decays of ^8Li and ^8B ”. en. In: *Physical Review* 93.3 (Feb. 1954), 499–500. ISSN: 0031-899X. DOI: [10.1103/PhysRev.93.499](https://doi.org/10.1103/PhysRev.93.499).
- [60] W. T. Winter. “Experimental Determination of the ^8B Neutrino Spectrum”. PhD thesis. University of California, Berkeley, 2007.
- [61] E. G. Adelberger et al. “Solar fusion cross sections. II. The p p chain and CNO cycles”. en. In: *Reviews of Modern Physics* 83.1 (Apr. 2011), 195–245. ISSN: 0034-6861, 1539-0756. DOI: [10.1103/RevModPhys.83.195](https://doi.org/10.1103/RevModPhys.83.195).
- [62] C E Ortiz et al. “Shape of the ^8B Alpha and Neutrino Spectra”. en. In: *Physical Review Letters* 85.14 (2000). DOI: [10.1103/PhysRevLett.85.2909](https://doi.org/10.1103/PhysRevLett.85.2909).
- [63] Chris Ortiz. “Scalar Currents in $0^+ \rightarrow 0^+$ Beta Decay and the ^8B Neutrino Spectrum”. PhD thesis. University of Notre Dame, 2000.
- [64] W. T. Winter et al. “Determination of the ^8B Neutrino Spectrum”. en. In: *Physical Review Letters* 91.25 (Dec. 2003), p. 252501. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.91.252501](https://doi.org/10.1103/PhysRevLett.91.252501).
- [65] W. T. Winter et al. “The ^8B neutrino spectrum”. en. In: *Physical Review C* 73.2 (Feb. 2006), p. 025503. ISSN: 0556-2813, 1089-490X. DOI: [10.1103/PhysRevC.73.025503](https://doi.org/10.1103/PhysRevC.73.025503).
- [66] O.S. Kirsebom et al. “Analysis of the response of silicon detectors to α particles and ^{16}O ions”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 758 (Sept. 2014), 57–61. ISSN: 01689002. DOI: [10.1016/j.nima.2014.05.005](https://doi.org/10.1016/j.nima.2014.05.005).
- [67] M. Bhattacharya, E. G. Adelberger, and H. E. Swanson. “Precise study of the final-state continua in ^8Li and ^8B decays”. en. In: *Physical Review C* 73.5 (May 2006), p. 055802. ISSN: 0556-2813, 1089-490X. DOI: [10.1103/PhysRevC.73.055802](https://doi.org/10.1103/PhysRevC.73.055802).

- [68] O. S. Kirsebom et al. “Precise and accurate determination of the ^8B decay spectrum”. en. In: *Physical Review C* 83.6 (June 2011). ISSN: 0556-2813, 1089-490X. DOI: [10.1103/PhysRevC.83.065802](https://doi.org/10.1103/PhysRevC.83.065802).
- [69] Oliver Sølund Kirsebom. “ ^8B Neutrinos and ^{12}C Resonances”. en. PhD thesis. Denmark: Aarhus University, 2010.
- [70] T. Roger et al. “Precise Determination of the Unperturbed ^8B Neutrino Spectrum”. en. In: *Physical Review Letters* 108.16 (Apr. 2012), p. 162502. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.108.162502](https://doi.org/10.1103/PhysRevLett.108.162502).
- [71] S. J. Freedman et al. “Comment on “Precise Determination of the Unperturbed ^8B Neutrino Spectrum””. en. In: *Physical Review Letters* 109.18 (Nov. 2012), p. 189201. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.109.189201](https://doi.org/10.1103/PhysRevLett.109.189201).
- [72] T. Roger et al. “Roger et al. Reply:” en. In: *Physical Review Letters* 109.18 (Nov. 2012), p. 189202. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.109.189202](https://doi.org/10.1103/PhysRevLett.109.189202).
- [73] S. Viñals et al. “Calibration and response function of a compact silicon-detector set-up for charged-particle spectroscopy using GEANT4”. en. In: *The European Physical Journal A* 57.2 (Feb. 2021), p. 49. ISSN: 1434-6001, 1434-601X. DOI: [10.1140/epja/s10050-021-00371-5](https://doi.org/10.1140/epja/s10050-021-00371-5).
- [74] S. Viñals i Onès. “Beta-decay of ^8B into highly excited states of ^8Be : Isospin mixing and proton-halo contributions”. PhD thesis. Universidad Complutense de Madrid, Nov. 2020.
- [75] B. Longfellow et al. “Determination of the ^8B neutrino energy spectrum using trapped ions”. en. In: *Physical Review C* 107.3 (Mar. 2023), p. L032801. ISSN: 2469-9985, 2469-9993. DOI: [10.1103/PhysRevC.107.L032801](https://doi.org/10.1103/PhysRevC.107.L032801).
- [76] J. Büscher et al. “ β -Decay studies with an implantation technique”. en. In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 266.19–20 (Oct. 2008), 4652–4656. ISSN: 0168583X. DOI: [10.1016/j.nimb.2008.05.084](https://doi.org/10.1016/j.nimb.2008.05.084).
- [77] L. Varriano et al., in progress. en. 2023.
- [78] N.D. Scielzo et al. “The Beta-decay Paul trap: A radiofrequency-quadrupole ion trap for precision studies”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 681 (July 2012), 94–100. ISSN: 01689002. DOI: [10.1016/j.nima.2012.04.035](https://doi.org/10.1016/j.nima.2012.04.035).
- [79] Ryan Matthew Yee. “Beta-delayed neutron spectroscopy using trapped radioactive ions”. en. PhD thesis. University of California, Berkeley, 2013.
- [80] Matt Sternberg. “Limits on Tensor Currents from ^8Li beta decay”. English. PhD thesis. University of Chicago, 2013.
- [81] Shane Caldwell. “A trapped-ion technique for beta-delayed neutron studies”. en. PhD thesis. University of Chicago, 2015.

- [82] Agnieszka Anna Czeszumaska. “Beta-delayed neutron studies of $^{137-138}\text{I}$ and $^{144-145}\text{Cs}$ performed with trapped ions”. en. PhD thesis. University of California, Berkeley, 2016.
- [83] Elizabeth Klara Marie Heckmaier. “Addressing the Reactor Antineutrino Anomaly with a New Detector Array for Fission Fragment Beta-Spectroscopy”. en. PhD thesis. University of California, Irvine, 2018.
- [84] Kevin Siegl. “IMPACT OF PRECISE β -DECAY INFORMATION ON TRAPPED ION β -DELAYED NEUTRON SPECTROSCOPY”. en. PhD thesis. University of Notre Dame, 2019.
- [85] Graeme Morgan et al. “The BEtA Recoil-ion trap (BEARtrap): An ion trap dedicated to β -delayed neutron studies using recoil-ion time-of-flight spectroscopy”. In: *2021 Fall Meeting of the APS Division of Nuclear Physics* (2021).
- [86] Mary T. Burkey et al. “Precision $\beta - \nu$ correlation measurements with the Beta-decay Paul Trap”. en. In: *Hyperfine Interactions* 240.1 (June 2019), p. 36. ISSN: 0304-3843, 1572-9540. DOI: [10.1007/s10751-019-1580-0](https://doi.org/10.1007/s10751-019-1580-0).
- [87] Gang Li. “Electron-Neutrino Angular Correlation Measurement in the Decay of ^8Li ”. en. PhD thesis. McGill University, 2012.
- [88] N.D. Scielzo et al. “A Novel Approach to β -delayed Neutron Spectroscopy Using the Beta-decay Paul Trap”. en. In: *Nuclear Data Sheets* 120 (June 2014), 70–73. ISSN: 00903752. DOI: [10.1016/j.nds.2014.07.009](https://doi.org/10.1016/j.nds.2014.07.009).
- [89] K. Siegl et al. “Recoil ions from the β decay of ^{134}Sb confined in a Paul trap”. en. In: *Physical Review C* 97.3 (Mar. 2018), p. 035504. ISSN: 2469-9985, 2469-9993. DOI: [10.1103/PhysRevC.97.035504](https://doi.org/10.1103/PhysRevC.97.035504).
- [90] A. Czeszumaska et al. “ β -delayed neutron emission studies of $^{137-138}\text{I}$ and $^{144-145}\text{Cs}$ performed with trapped ions”. en. In: *Physical Review C* 101.2 (Feb. 2020), p. 024312. ISSN: 2469-9985, 2469-9993. DOI: [10.1103/PhysRevC.101.024312](https://doi.org/10.1103/PhysRevC.101.024312).
- [91] B. S. Wang et al. “ β -delayed-neutron studies of $^{135,136}\text{Sb}$ and ^{140}I performed with trapped ions”. en. In: *Physical Review C* 101.2 (Feb. 2020), p. 025806. ISSN: 2469-9985, 2469-9993. DOI: [10.1103/PhysRevC.101.025806](https://doi.org/10.1103/PhysRevC.101.025806).
- [92] R. M. Yee et al. “ β -Delayed Neutron Spectroscopy Using Trapped Radioactive Ions”. en. In: *Physical Review Letters* 110.9 (Feb. 2013), p. 092501. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.110.092501](https://doi.org/10.1103/PhysRevLett.110.092501).
- [93] Wolfgang Paul, Nobel Lecture. en. 1989. URL: <https://www.nobelprize.org/prizes/physics/1989/paul/lecture/>.
- [94] Raymond E. March and John F. J. Todd. *Quadrupole ion trap mass spectrometry*. en. 2nd ed. Chemical analysis. Hoboken, N.J.: J. Wiley, 2005. ISBN: 978-0-471-48888-0.
- [95] Tommi Eronen, Anu Kankainen, and Juha Äystö. “Ion traps in nuclear physics—Recent results and achievements”. en. In: *Progress in Particle and Nuclear Physics* 91 (Nov. 2016), 259–293. ISSN: 01466410. DOI: [10.1016/j.pnpnp.2016.08.001](https://doi.org/10.1016/j.pnpnp.2016.08.001).

- [96] P. Delahaye et al. “The open LPC Paul trap for precision measurements in beta decay”. en. In: *The European Physical Journal A* 55.6 (June 2019), p. 101. ISSN: 1434-6001, 1434-601X. DOI: [10.1140/epja/i2019-12777-3](https://doi.org/10.1140/epja/i2019-12777-3).
- [97] R. F. Wuerker, H. Shelton, and R. V. Langmuir. “Electrodynamic Containment of Charged Particles”. en. In: *Journal of Applied Physics* 30.3 (Mar. 1959), 342–349. ISSN: 0021-8979, 1089-7550. DOI: [10.1063/1.1735165](https://doi.org/10.1063/1.1735165).
- [98] M. Drewsen and A. Brøner. “Harmonic linear Paul trap: Stability diagram and effective potentials”. en. In: *Physical Review A* 62.4 (Sept. 2000), p. 045401. ISSN: 1050-2947, 1094-1622. DOI: [10.1103/PhysRevA.62.045401](https://doi.org/10.1103/PhysRevA.62.045401).
- [99] Fernande Vedel. “On the dynamics and energy of ion clouds stored in an R.F. quadrupole trap”. en. In: *International Journal of Mass Spectrometry and Ion Processes* 106 (May 1991), 33–61. ISSN: 01681176. DOI: [10.1016/0168-1176\(91\)85011-A](https://doi.org/10.1016/0168-1176(91)85011-A).
- [100] F. Vedel et al. “Computed energy and spatial statistical properties of stored ions cooled by a buffer gas”. en. In: *Physical Review A* 27.5 (May 1983), 2321–2330. ISSN: 0556-2791. DOI: [10.1103/PhysRevA.27.2321](https://doi.org/10.1103/PhysRevA.27.2321).
- [101] Gang Li et al. “A compact cryogenic pump”. en. In: *Cryogenics* 75 (Apr. 2016), 35–37. ISSN: 00112275. DOI: [10.1016/j.cryogenics.2015.12.006](https://doi.org/10.1016/j.cryogenics.2015.12.006).
- [102] 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>.
- [103] J. Allison et al. “Geant4 developments and applications”. In: *IEEE Transactions on Nuclear Science* 53.1 (2006), pp. 270–278. DOI: [10.1109/TNS.2006.869826](https://doi.org/10.1109/TNS.2006.869826).
- [104] J. Allison et al. “Recent developments in Geant4”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 835 (2016), pp. 186–225. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2016.06.125>. URL: <https://www.sciencedirect.com/science/article/pii/S0168900216306957>.
- [105] Autodesk Inc. Autodesk Inventor. 2022. URL: <https://www.autodesk.com/>.
- [106] Andrii Tykhonov. cad-to-geant4 converter. 2022. URL: <https://github.com/tihonav/cad-to-geant4-converter>.
- [107] Scientific Instrument Services, Inc. SIMION. 2022. URL: www.simion.com.
- [108] M. Benali et al. “Geometry optimisation of a transparent axisymmetric ion trap for the MORA project”. en. In: *The European Physical Journal A* 56.6 (June 2020). arXiv:2004.09351 [nucl-ex, physics:physics], p. 163. ISSN: 1434-6001, 1434-601X. DOI: [10.1140/epja/s10050-020-00168-y](https://doi.org/10.1140/epja/s10050-020-00168-y).
- [109] H. W. Koch and J. W. Motz. “Bremsstrahlung Cross-Section Formulas and Related Data”. en. In: *Reviews of Modern Physics* 31.4 (Oct. 1959), 920–955. ISSN: 0034-6861. DOI: [10.1103/RevModPhys.31.920](https://doi.org/10.1103/RevModPhys.31.920).

- [110] HTW Hochttemperatur-Werkstoffe GmbH. 2022. URL: <http://www.htw-germany.com/>.
- [111] National Institute of Standards and Technology. “Cryogenic Technology Resources”. 2022. URL: <https://trc.nist.gov/cryogenics/index.html>.
- [112] George A. Beitel. “The Use of Graphite in High and Ultrahigh Vacuum: A Review”. In: *Journal of Vacuum Science and Technology* 8 (1971), pp. 647–657.
- [113] C Accettura et al. “Ultra-High Vacuum characterization of Molybdenum-Carbide Graphite for HL-LHC collimators”. en. In: *Journal of Physics: Conference Series* 1350.1 (Nov. 2019), p. 012085. ISSN: 1742-6588, 1742-6596. DOI: [10.1088/1742-6596/1350/1/012085](https://doi.org/10.1088/1742-6596/1350/1/012085).
- [114] MWI, Inc. “EDM Graphite”. 2022. URL: <https://mwi-inc.com/edm-graphite/>.
- [115] Daniel Burdette et al. “Development of the St. Benedict Paul Trap at the Nuclear Science Laboratory”. en. In: *Hyperfine Interactions* 240.1 (Dec. 2019), p. 70. ISSN: 0304-3843, 1572-9540. DOI: [10.1007/s10751-019-1606-7](https://doi.org/10.1007/s10751-019-1606-7).
- [116] TDK Electronics. “Ferrite Materials”. 2022. URL: <https://www.tdk-electronics.tdk.com/en/529404/products/product-catalog/ferrites-and-accessories/ferrite-materials>.
- [117] Accu-Glass Products Inc. “32 AWG, 50 Ohm Coaxial Cable”. 2022. URL: <https://www.accuglassproducts.com/uhv-kapton-insulated-wire-and-cable/32-awg-50-ohm-coaxial-cable>.
- [118] Stanford Research Systems. 2023. URL: <https://www.thinksrs.com/products/DS345.htm>.
- [119] Micron Semiconductor Ltd. “BB7 detector design”. accessed Jan. 2023. URL: <http://www.micronsemiconductor.co.uk/product/bb7/>.
- [120] Eljen Technology. “EJ-200, EJ-204, EJ-208, EJ-212”. accessed March 2023. URL: <https://eljentechnology.com/products/plastic-scintillators/ej-200-ej-204-ej-208-ej-212>.
- [121] Hamamatsu Photonics K. K. “PHOTOMULTIPLIER TUBES R877, R877-01”. accessed March 2023. URL: https://www.hamamatsu.com/us/en/product/optical-sensors/pmt/pmt_tube-alone/head-on-type/R877.html.
- [122] C. Müller-Gatermann. personal communication. 2023.
- [123] mesytec GmbH & Co. KG. “MPR-16/L MPR-32 / MPR-64 data sheet”. accessed March 2023. URL: <https://www.mesytec.com/products/datasheets/MPR-16.pdf>.
- [124] Ken Teh, “Data Acquisition at ATLAS”, ATLAS Users Workshop. en. 2006. URL: https://www.phy.anl.gov/atlas/User_Meeting/Teh_ATLAS-Workshop-Apr08.pdf.
- [125] Rene Brun and Fons Rademakers. “ROOT — An object oriented data analysis framework”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 389.1 (1997), 81–86. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(97\)00048-X](https://doi.org/10.1016/S0168-9002(97)00048-X).

- [126] mesytec GmbH & Co. KG. “MSCF-16 F data sheet”. accessed March 2023. URL: https://www.mesytec.com/products/nuclear-physics/MSCF-16_F_V.html.
- [127] mesytec GmbH & Co. KG. custom MSCF-16 quote. communication to Jason Clark. 2012.
- [128] Phillips Scientific “16 Channel Peak ADC CAMAC MODEL 7164, 7164H data sheet”.
- [129] L. Varriano et al. “ATLAS Proposal 1957: A new limit on the tensor contribution to the weak interaction with ^8Li ”. 2021.
- [130] Argonne National Laboratory, “ATLAS Facility Layout with Beam Locations”. 2022. URL: <https://www.anl.gov/atlas/reference/atlas-facility-layout-with-beam-locations>.
- [131] A.J. Koning. Calculation and NOV 2015. analysis of d+7Li reaction. ENDF/TENDL-15 MAT 328.
- [132] J. P. Schiffer et al. “Study of the (d, p) Reaction in the 1p Shell”. en. In: *Physical Review* 164.4 (Dec. 1967), 1274–1284. ISSN: 0031-899X. DOI: [10.1103/PhysRev.164.1274](https://doi.org/10.1103/PhysRev.164.1274).
- [133] Guy Savard. “Large radio-frequency gas catchers and the production of radioactive nuclear beams”. en. In: *Journal of Physics: Conference Series* 312.5 (Sept. 2011), p. 052004. ISSN: 1742-6596. DOI: [10.1088/1742-6596/312/5/052004](https://doi.org/10.1088/1742-6596/312/5/052004).
- [134] Perma Pure, “All About Nafion™ Tubing”. 2023. URL: <https://www.permapure.com/environmental-scientific/resources/all-about-nafion-and-faq/>.
- [135] Lowell S. Brown and Gerald Gabrielse. “Geonium theory: Physics of a single electron or ion in a Penning trap”. en. In: *Reviews of Modern Physics* 58.1 (Jan. 1986), 233–311. ISSN: 0034-6861. DOI: [10.1103/RevModPhys.58.233](https://doi.org/10.1103/RevModPhys.58.233).
- [136] G. Savard et al. “A new cooling technique for heavy ions in a Penning trap”. en. In: *Physics Letters A* 158.5 (Sept. 1991), 247–252. ISSN: 03759601. DOI: [10.1016/0375-9601\(91\)91008-2](https://doi.org/10.1016/0375-9601(91)91008-2).
- [137] Edmond P F Lee and Timothy G Wright. “Heats of formation of $\text{LiOH}(X^1\Sigma^+)$ and $\text{LiOH}^+(X^2\Pi)$: the ionization energy of LiOH ”. en. In: *Chemical Physics Letters* (2002).
- [138] Berkeley Nucleonics Corporation. “Model PB-5 Precision NIM Pulse Generator”. accessed Feb. 2023. URL: <https://www.berkeleynucleonics.com/model-pb-5-precision-nim-pulse-generator>.
- [139] L. Varriano. “ATLAS Proposal 1905: α beam characterization of a DSSD for the Beta-decay Paul Trap”. 2021.
- [140] Eckert & Ziegler. “Certificate of Calibration Alpha Standard Source (^{244}Cm)”. communications with N. D. Scielzo. 2015.
- [141] Eckert & Ziegler. “Certificate of Calibration Alpha Standard Source (^{148}Gd)”. communications with N. D. Scielzo. 2015.

- [142] T.Y. Hirsh et al. “The use of cosmic-ray muons in the energy calibration of the Beta-decay Paul Trap silicon-detector array”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 887 (Apr. 2018), 122–127. ISSN: 01689002. DOI: [10.1016/j.nima.2018.01.021](https://doi.org/10.1016/j.nima.2018.01.021).
- [143] Balraj Singh and E. Browne. “Nuclear Data Sheets for A = 240”. en. In: *Nuclear Data Sheets* 109.10 (Oct. 2008), 2439–2499. ISSN: 00903752. DOI: [10.1016/j.nds.2008.09.002](https://doi.org/10.1016/j.nds.2008.09.002).
- [144] Y. A. Akovali. “Review of Alpha-Decay Data from Doubly-Even Nuclei”. en. In: *NUCLEAR DATA SHEETS* 84.1 (May 1998), 1–114. DOI: [10.1006/ndsh.1998.0009](https://doi.org/10.1006/ndsh.1998.0009).
- [145] E Steinbauer et al. “A survey of the physical processes which determine the response function of silicon detectors to alpha particles”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 339.1–2 (Jan. 1994), 102–108. ISSN: 01689002. DOI: [10.1016/0168-9002\(94\)91787-6](https://doi.org/10.1016/0168-9002(94)91787-6).
- [146] E. Steinbauer et al. “Energy resolution of silicon detectors: approaching the physical limit”. en. In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 85.1–4 (Mar. 1994), 642–649. ISSN: 0168583X. DOI: [10.1016/0168-583X\(94\)95898-X](https://doi.org/10.1016/0168-583X(94)95898-X).
- [147] G. Bortels and P. Collaers. “Analytical function for fitting peaks in alpha-particle spectra from Si detectors”. en. In: *International Journal of Radiation Applications and Instrumentation. Part A. Applied Radiation and Isotopes* 38.10 (Jan. 1987), 831–837. ISSN: 08832889. DOI: [10.1016/0883-2889\(87\)90180-8](https://doi.org/10.1016/0883-2889(87)90180-8).
- [148] M. Jurado Vargas, A. Fernández Timón, and J.F. Ziegler. “Determination and analysis of ionization/excitation distributions for alpha particles in silicon detectors”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 971 (Aug. 2020), p. 164134. ISSN: 01689002. DOI: [10.1016/j.nima.2020.164134](https://doi.org/10.1016/j.nima.2020.164134).
- [149] P. Bauer and G. Bortels. “Response of Si detectors to electrons, deuterons and alpha particles”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 299.1–3 (Dec. 1990), 205–209. ISSN: 01689002. DOI: [10.1016/0168-9002\(90\)90777-4](https://doi.org/10.1016/0168-9002(90)90777-4).
- [150] K.E. Petersen. “Dynamic micromechanics on silicon: Techniques and devices”. en. In: *IEEE Transactions on Electron Devices* 25.10 (Oct. 1978), 1241–1250. ISSN: 0018-9383. DOI: [10.1109/T-ED.1978.19259](https://doi.org/10.1109/T-ED.1978.19259).
- [151] SRIM documentation/ “The Stopping and Range in Compounds”. accessed Feb. 2023. URL: <http://www.srim.org/SRIM/Compounds.htm>.
- [152] W.N. Lennard et al. “Nonlinear response of Si detectors for low-Z ions”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 248.2–3 (Aug. 1986), 454–460. ISSN: 01689002. DOI: [10.1016/0168-9002\(86\)91033-8](https://doi.org/10.1016/0168-9002(86)91033-8).

- [153] A. Fernández Timón. personal communication. 2022.
- [154] R.H. Pehl et al. “Accurate determination of the ionization energy in semiconductor detectors”. en. In: *Nuclear Instruments and Methods* 59.1 (Feb. 1968), 45–55. ISSN: 0029554X. DOI: [10.1016/0029-554X\(68\)90342-X](https://doi.org/10.1016/0029-554X(68)90342-X).
- [155] Glenn F. Knoll. *Radiation detection and measurement*. en. 4th ed. Hoboken, N.J: John Wiley, 2010. ISBN: 978-0-470-13148-0.
- [156] F. Munnik et al. “Stopping powers of 200–3000 keV ^4He and 550–1750 keV ^1H ions in Vvyns”. en. In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 119.4 (Dec. 1996), 445–451. ISSN: 0168583X. DOI: [10.1016/S0168-583X\(96\)00468-5](https://doi.org/10.1016/S0168-583X(96)00468-5).
- [157] C Lee and N.R Fletcher. “Parameterization of the pulse height defect and resolution for low-Z ions incident on silicon barrier detectors”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 432.2–3 (Aug. 1999), 313–317. ISSN: 01689002. DOI: [10.1016/S0168-9002\(99\)00230-2](https://doi.org/10.1016/S0168-9002(99)00230-2).
- [158] J B Birks. “Scintillations from Organic Crystals: Specific Fluorescence and Relative Response to Different Radiations”. en. In: *Proceedings of the Physical Society. Section A* 64.10 (Oct. 1951), 874–877. ISSN: 0370-1298. DOI: [10.1088/0370-1298/64/10/303](https://doi.org/10.1088/0370-1298/64/10/303).
- [159] P. Bauer and G. Bortels. “Response of Si detectors to electrons, deuterons and alpha particles”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 299.1–3 (Dec. 1990), 205–209. ISSN: 01689002. DOI: [10.1016/0168-9002\(90\)90777-4](https://doi.org/10.1016/0168-9002(90)90777-4).
- [160] Y. Zhang and W. J. Weber. “Role of energy partitioning on electron-hole recombination, trapping, and detection in silicon detectors”. en. In: *Physical Review B* 82.7 (Aug. 2010), p. 075202. ISSN: 1098-0121, 1550-235X. DOI: [10.1103/PhysRevB.82.075202](https://doi.org/10.1103/PhysRevB.82.075202).
- [161] A.M. Beesley et al. “Evolution of chemical species during electrodeposition of uranium for alpha spectrometry by the Hallstadius method”. en. In: *Applied Radiation and Isotopes* 67.9 (Sept. 2009), 1559–1569. ISSN: 09698043. DOI: [10.1016/j.apradiso.2009.03.031](https://doi.org/10.1016/j.apradiso.2009.03.031).
- [162] C. G. Méndez et al. “Nanoscopic study of chemical species during uranium electrodeposition for alpha spectrometry sources”. en. In: *Journal of Materials Science* 45.18 (Sept. 2010), 5061–5070. ISSN: 0022-2461, 1573-4803. DOI: [10.1007/s10853-010-4389-4](https://doi.org/10.1007/s10853-010-4389-4).
- [163] S. Walsh, Micron Semiconductor Ltd. communications with N. D. Scielzo. 2015, 2016.
- [164] Bruker. “DektakXT”. accessed Jan. 2023. URL: <https://www.bruker.com/fr/products-and-solutions/test-and-measurement/stylus-profilometers/dektakxt.html>.

- [165] S. Takeda et al. “Development of double-sided silicon strip detectors (DSSD) for a Compton telescope”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 579.2 (Sept. 2007), 859–865. ISSN: 01689002. DOI: [10.1016/j.nima.2007.05.305](https://doi.org/10.1016/j.nima.2007.05.305).
- [166] L. Grassi et al. “Study of the inter-strip gap effects on the response of Double Sided Silicon Strip Detectors using proton micro-beams”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 767 (Dec. 2014), 99–111. ISSN: 01689002. DOI: [10.1016/j.nima.2014.08.009](https://doi.org/10.1016/j.nima.2014.08.009).
- [167] D. Torresi et al. “Influence of the interstrip gap on the response and the efficiency of Double Sided Silicon Strip Detectors”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 713 (June 2013), 11–18. ISSN: 01689002. DOI: [10.1016/j.nima.2013.02.027](https://doi.org/10.1016/j.nima.2013.02.027).
- [168] J. Yorkston et al. “Interstrip surface effects in oxide passivated ion-implanted silicon strip detectors”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 262.2–3 (Dec. 1987), 353–358. ISSN: 01689002. DOI: [10.1016/0168-9002\(87\)90873-4](https://doi.org/10.1016/0168-9002(87)90873-4).
- [169] Ya Huang et al. “Large-area, transferable sub-10 nm polymer membranes at the air–water interface”. en. In: *Nano Research* 11.7 (July 2018), 3833–3843. ISSN: 1998-0124, 1998-0000. DOI: [10.1007/s12274-017-1957-9](https://doi.org/10.1007/s12274-017-1957-9).
- [170] I. Asimov, “Franchise,” in *Isaac Asimov: The Complete Stories, vol. 1*, Broadway Books, New York, 1990.
- [171] S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. en. In: *The Annals of Mathematical Statistics* 9.1 (Mar. 1938), 60–62. ISSN: 0003-4851. DOI: [10.1214/aoms/1177732360](https://doi.org/10.1214/aoms/1177732360).
- [172] Abraham Wald. “Tests of statistical hypotheses concerning several parameters when the number of observations is large”. en. In: *Transactions of the American Mathematical Society* 54.3 (1943), 426–482. ISSN: 0002-9947, 1088-6850. DOI: [10.1090/S0002-9947-1943-0012401-3](https://doi.org/10.1090/S0002-9947-1943-0012401-3).
- [173] Wolfgang Paul and Helmut Steinwedel. “Notizen: Ein neues Massenspektrometer ohne Magnetfeld”. In: *Zeitschrift für Naturforschung A* 8.7 (1953), pp. 448–450. DOI: [doi: 10.1515/zna-1953-0710](https://doi.org/10.1515/zna-1953-0710).
- [174] Kenneth G. Libbrecht and Eric D. Black. “Improved microparticle electrodynamic ion traps for physics teaching”. en. In: *American Journal of Physics* 86.7 (July 2018), 539–558. ISSN: 0002-9505, 1943-2909. DOI: [10.1119/1.5034344](https://doi.org/10.1119/1.5034344).
- [175] V Rybin et al. “DIY electrodynamic trap for physics education”. In: *Physics Education* 57.1 (Dec. 2021), p. 015023. DOI: [10.1088/1361-6552/ac3b93](https://doi.org/10.1088/1361-6552/ac3b93).
- [176] Information Unlimited, “Transformer High Voltage w/Switch 6kv 30ma 60Hz”. en. 2023. URL: <https://www.amazing1.com/products/transformer-high-voltage-w-switch-6kv-30ma-60hz.html>.