

©Copyright 2023

William Andrew Byron

Developments in the application of Cyclotron Radiation Emission
Spectroscopy (CRES) towards the precise determination of
MeV-scale β -energy spectra

William Andrew Byron

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2023

Reading Committee:

Alejandro Garcia, Chair

Gray Rybka

Steven Sharpe

Program Authorized to Offer Degree:

Physics

University of Washington

Abstract

Developments in the application of Cyclotron Radiation Emission Spectroscopy (CRES) towards the precise determination of MeV-scale β -energy spectra

William Andrew Byron

Chair of the Supervisory Committee:
Alejandro Garcia
Physics

This thesis presents an apparatus for detection of cyclotron radiation yielding a frequency-based β^\pm kinetic energy determination in the 5 keV to 2.1 MeV range, characteristic of nuclear β decays. The cyclotron frequency of the radiating β particles in a magnetic field is used to determine the β energy precisely. The He6-CRES collaboration is working towards competitive sensitivity to tensor currents in the weak interaction via a precise determination of b_{Fierz} , which parameterizes Fierz interference. This thesis presents much of the work done to complete *Phase 1* of the He6-CRES experiment, whose overarching goal was to demonstrate that Cyclotron Radiation Emission Spectroscopy (CRES) can plausibly be used to achieve a 10^{-3} measurement of b_{Fierz} via the ratio of ${}^6\text{He}$ and ${}^{19}\text{Ne}$ spectra and to discover what experimental design, analysis, or simulation bottlenecks may hinder such a measurement. To our knowledge, this is the first direct observation of cyclotron radiation from individual highly relativistic β s in a waveguide. This work focuses on applying CRES to a precise measurement of b_{Fierz} but it also establishes the viability of the application of CRES to a variety of MeV-endpoint nuclei, opening its reach to searches for new physics beyond the TeV-scale via precision β -decay measurements.

TABLE OF CONTENTS

	Page
List of Figures	vi
List of Tables	x
Chapter 1: Theory and Motivation for Advances in Precision β Spectroscopy . . .	1
1.1 Weak interaction in the Standard Model	2
1.1.1 Standard Model overview	2
1.1.2 Charged-current weak interactions in the Standard Model	5
1.1.3 Charged-current weak interactions beyond the Standard Model	11
1.2 Energy spectrum and b_{Fierz}	16
1.2.1 Energy spectrum	16
1.2.2 Fierz interference term	17
1.2.3 Energy spectrum corrections	19
1.3 Experimental limits on exotic currents	22
1.4 Fierz term b_{Fierz} for ${}^6\text{He}$, ${}^{19}\text{Ne}$	26
1.5 Summary of theoretical motivation	29
Chapter 2: Overview of the CRES Technique	30
2.1 Simple description of CRES	30
2.2 Characteristics of CRES electrons	31
2.3 Characteristics of CRES signals	35
2.4 Characteristics of CRES spectrogram	36
Chapter 3: Overview of the He6-CRES Experiment	39
3.1 The He6-CRES experiment in a nutshell	39
3.1.1 Experimental goals of <i>Phase 1</i> of the He6-CRES experiment	39
3.1.2 <i>Phase 1</i> of the He6-CRES experiment	41

3.1.3	Thesis scope	48
3.2	Transport, vacuum, and monitoring systems	48
3.3	Production, intensity, and purity of radioactive sources	51
3.3.1	${}^6\text{He}$ source	51
3.3.2	${}^{19}\text{Ne}$ source	53
3.4	He6-CRES apparatus	58
3.4.1	Decay cell and magnetic trap	58
3.4.2	RF system	65
3.4.3	Cryo-system and temperature distribution	71
3.5	DAQ	74
3.6	Magnet	76
3.7	β monitor	80
3.8	DAQ user interface and SQL database	82
3.9	RF noise and SNR characteristics	85
3.9.1	RF noise background equations	85
3.9.2	Simplest working model of the He6-CRES RF system	87
3.9.3	Signal injection for RF calibration	89
3.9.4	Comparison of cascade analysis to measurement	93
3.9.4.1	Cascade analysis: assumptions and limitations	93
3.9.4.2	Network analyzer measurements	95
3.9.5	Comparison	97
3.9.6	Discussion of other relevant RF measurements	98
3.9.6.1	${}^{83\text{m}}\text{Kr}$ mean SNR	98
3.9.6.2	Noise floor with terminator only	99
3.9.6.3	Changing the waveguide temperatures	100
3.9.7	Considerations to optimize SNR	101
3.9.8	Reflection model of He6-CRES noise floor	102
3.9.9	Conclusion of RF noise and SNR characteristics	104
Chapter 4:	Initial Observations of CRES Events from ${}^{83\text{m}}\text{Kr}$, ${}^6\text{He}$, and ${}^{19}\text{Ne}$	109
4.1	Initial observations of CRES events from ${}^{83\text{m}}\text{Kr}$	109
4.2	Initial observations of CRES events from ${}^{19}\text{Ne}$ and ${}^6\text{He}$	110

Chapter 5:	Monte Carlo Simulation	119
5.1	Simple Monte Carlo of ratio experiment	119
5.1.1	Ratio measurement sensitivity to b_{Fierz}	121
5.1.2	Sensitivity to main magnetic field determination	124
5.1.3	Sensitivity to monitor instability	127
5.1.4	An alternate way to present a ratio experiment	129
5.1.5	Finer frequency binning to improve efficiency cancellation	130
5.1.6	Wall-effect efficiency	131
5.2	Monte Carlo including CRES events born below the detection window	136
5.3	Particle-level Monte Carlo	140
Chapter 6:	Data Analysis and Preliminary Ratio Measurement	149
6.1	Event reconstruction	149
6.1.1	Katydid	150
6.1.2	Python post-processing	151
6.2	Building the ratio	153
6.3	Preliminary ratio measurement	156
6.3.1	Ratio measurement experimental process	156
6.3.2	Comparing observed and predicted ratios	161
6.3.2.1	Including events from below	161
6.3.2.2	Cutting events from below	166
6.4	Discussion of systematics	176
Chapter 7:	Deep Learning Image Segmentation for CRES Event Classification	181
7.1	Introduction to UNET for CRES	183
7.1.1	Convolutional neural networks	183
7.1.2	UNET for CRES event reconstruction	184
7.1.3	UNET architecture	185
7.1.4	Loss function: cross-entropy loss	189
7.1.5	Gradient descent algorithm: Adam	189
7.1.6	Machine learning metrics	190
7.2	Overview of the three studies conducted	192
7.2.1	Simulated training data	193

7.2.2	Preprocessing data with max pooling	194
7.2.3	UNET implementation details	195
7.3	Study 1: Constant SNR tracks, no sidebands	197
7.3.1	Specifics of simulated training data	197
7.3.2	Model hyperparameters and training configuration	198
7.3.3	Study results	198
7.4	Study 2: Tracks with oscillating SNR, no sidebands	199
7.4.1	Specifics of simulated training data	199
7.4.2	Model hyperparameters and training configuration	200
7.4.3	Study results	200
7.5	Study 3: Constant SNR tracks, sidebands	201
7.5.1	Specifics of simulated training data	201
7.5.2	Model hyperparameters and training configuration	201
7.5.3	Study results	202
7.6	Conclusion of UNET studies	203
Chapter 8:	Conclusions	216
8.1	Experimental Summary	216
8.2	Towards a 10^{-3} determination of b_{Fierz} with CRES	216
8.3	Final remarks	219
Appendix A:	Chirp SNR Optimization	220
A.1	The concept	220
A.1.1	A balance	220
A.1.2	A more precise model	221
A.1.3	Is this the best we can do?	223
Appendix B:	Kapton Adsorption of Helium	225
B.1	Background	225
B.1.1	General governing equations	225
B.1.2	Finite lifetime effect on diffusion	227
B.1.3	Modeling our experimental configuration	228
B.2	Literature review	229
B.2.1	Harvard-group determination of the Diffusivity of ^4He	229

B.2.2	Motivating our measurement	232
B.3	Experimental determination of ${}^6\text{He}$ absorption into kapton	232
B.3.1	Methods	232
B.4	Solution to Barrer’s equation	234
B.5	Particular situations: The finite solid with fixed pressure	235
B.6	Time evolution towards steady state pressure	237
Appendix C: PostgreSQL Database (<code>he6cres_db</code>) Details		241
Appendix D: Particle-level Monte Carlo Simulation Documentation		245
D.1	Purpose	245
D.2	Overview	245
D.3	Trap Field	246
D.3.1	Comparison between Kassiopeia and Python	246
D.3.2	Trap field ρ dependence	246
D.4	Analytic calculations of CRES dynamics	248
D.4.1	Dynamics assuming no ρ dependence in B	248
D.4.2	Dynamics assuming quadratic ρ dependence on B	250
D.5	Comparison to Kassiopeia	254
D.6	Conclusion	255
Appendix E: Deep learning Source Code		261
E.1	UNET model implementation	261
E.2	UNET lightning module	264
Bibliography		268

LIST OF FIGURES

Figure Number	Page
1.1 Fundamental particles of the Standard Model	3
1.2 Quark level β decay Feynman diagram	6
1.3 Example diagrams of chirality-flipping interactions in models beyond the SM	15
1.4 Feynman diagrams corresponding to order- α radiative corrections	20
1.5 Allowed regions for $Re(\epsilon_S)$ and $Re(\epsilon_T)$ given 10^{-3} measurements of b_{Fierz} . .	24
1.6 Allowed regions for $Re(\epsilon_S)$ and $Re(\epsilon_T)$ from LHC vs Low Energy	25
2.1 Schematic of harmonic magnetic trap	32
2.2 Three motions of CRES electron in trap	35
2.3 Anatomy of a CRES event	38
3.1 Energy regions of ^{19}Ne and ^6He sampled in ratio measurement	42
3.2 He6-CRES apparatus schematic	43
3.3 Simple Monte Carlo demonstration of ratio measurement	46
3.4 Layout of Tandem, Cave 1, Cave 2	47
3.5 Vacuum and transport systems	50
3.6 Cryo-II trap profile	51
3.7 Sketch of lithium target for ^6He production	52
3.8 Photo of lithium target for ^6He production	53
3.9 SF_6 flow system for ^{19}Ne production	54
3.10 ^{19}Ne production target	55
3.11 SF_6 cold trap	55
3.12 Decay cell with trap coils	61
3.13 Photo of decay cell and coil form	61
3.14 Failed trap emptying due to eddy currents	63
3.15 Thin walled decay cell	63
3.16 Design drawing of the 3/8" ID Ti radioactivity transport tube	64
3.17 RF components surrounding decay cell	66

3.18	Photo of copper box components	67
3.19	Full He6-CRES apparatus design	69
3.20	Labeled He6-CRES apparatus section	70
3.22	Temperature distribution of apparatus	73
3.23	SNR vs chirp slope vs bitcode	77
3.24	Design of apparatus and magnet rail system	78
3.25	Re-entrant tube for NMR probe	80
3.26	Stability of <i>beta monitor</i>	83
3.27	Stability of <i>beta telescope</i>	84
3.28	Axis definition relative to quarter wave plate dielectric and WR-42 waveguide.	90
3.29	Transmission and reflection vs quarter wave plate configuration	94
3.30	Field Fox measurement set-up	106
3.31	Terminator to LNA RF test configuration	107
3.32	Noise floor, full RF vs terminator only	108
4.1	$^{83\text{m}}\text{Kr}$ lines (7 – 32 keV) detected via CRES	110
4.2	Slope resonances observed with $^{83\text{m}}\text{Kr}$	111
4.3	CRES events from ^6He and ^{19}Ne at 1, 2, 3 T	112
4.4	Density of CRES events with and without trap slewing	113
4.5	Track slopes vs magnetic field	115
4.6	CRES event crossing the $\text{TE}_{2,2}$ ($h = 3$) resonance	117
4.7	Event length vs slope	118
5.1	Simple Monte Carlo of ratio measurement	120
5.2	Sensitivity to main field determination error	126
5.3	Sensitivity to linear β monitor drift	128
5.4	Wall effect visualization	134
5.5	^{19}Ne and ^6He spectral densities observed during ratio experiment	135
5.6	r_{cycl} 's observed during ratio experiment	136
5.7	Wall-effect efficiencies affecting ratio experiment	137
5.8	Wall-effect distortion of ratio	141
5.9	Wall-effect distortion of ^{19}Ne spectra (0)	142
5.10	Wall-effect distortion of ^{19}Ne spectra (1)	143
5.11	^{19}Ne and ^6He spectra with extended bandwidth	144

5.12	Percent of events coming from below vs field vs isotope	145
5.13	Percent change of ratio due to events from below	146
5.14	Particle-level Monte Carlo demonstration	147
5.15	Simulated <i>spec</i> file	148
6.1	Demonstration of the three main event reconstruction steps	154
6.2	Effect of the Katydid Fourier bin SNR threshold	155
6.3	Main field error (ppm)	158
6.4	Autumn 2022 noise floors for ${}^6\text{He}$ and ${}^{19}\text{Ne}$ runs	159
6.5	Autumn 2022 track segment SNRs from ${}^6\text{He}$ and ${}^{19}\text{Ne}$	160
6.6	Track length distributions	163
6.7	Ratio plot from SNR study including events from below	164
6.8	Ratio plot with adjusted uncertainties, including events from below	167
6.9	Ratio plot with adjusted uncertainties, naive 200 MHz cut	170
6.10	Event start frequency vs normalized SNR with and without cuts	172
6.11	Field-wise event start frequency vs normalized SNR	173
6.12	Ratio plot with adjusted uncertainties, naive 200 MHz cut, plus SNR cut	175
6.13	Efficiency-corrected energy domain plot for ${}^{19}\text{Ne}$, $n_{chunks} = 1$	176
6.14	Efficiency-corrected energy domain plot for ${}^{19}\text{Ne}$, $n_{chunks} = 2$	177
7.1	Visualization of the UNET model architecture	186
7.2	UNET Study One: Prediction on test set	205
7.3	UNET Study One: Prediction on test set (zoomed)	206
7.4	UNET Study One: Performance metrics summary	207
7.5	UNET Study Two: Prediction on test set	208
7.6	UNET Study Two: Prediction on test set (zoomed)	209
7.7	UNET Study Two: Performance metrics summary	210
7.8	UNET Study Three: Prediction on test set (1)	211
7.9	UNET Study Three: Prediction on test set (2)	212
7.10	UNET Study Three: Prediction on test set (zoomed)	213
7.11	UNET Study Three: Performance metrics summary	214
7.12	Example of Mask-RCNN applied to CRES data with scattering	215
A.1	Imagine the Track Slope expressed in units of Freq Bin/Slice.	220
A.2	A track traversing an $a \times b$ grid.	222

B.1	Figure 1 of Schowalter et al	230
B.2	Figure 2 of Schowalter et al	231
B.3	Histogram of decay events with overlaid fit	239
B.4	Plot of Equation B.39 for different times	240
D.1	Python vs Kassiopeia field profiles, z	247
D.2	Python vs Kassiopeia field profiles, ρ	248
D.3	Python vs Kassiopeia field profiles, ρ relative difference	249
D.4	Field profiles, ρ dependence inflection	256
D.5	Coordinate system used in analytical expressions for CRES properties	257
D.6	Kassiopeia vs analytical expressions, simulation configuration 0	258
D.7	Kassiopeia vs analytical expressions, simulation configuration 1	259
D.8	Kassiopeia vs analytical expressions, simulation configuration 2	260

LIST OF TABLES

Table Number	Page
1.1 Summary of VAST current properties	14
3.1 Possible contaminants created during ^{19}Ne production	59
3.2 Details on RF components	71
3.3 Temperature distribution of apparatus	73
3.4 Legend for T_{cres} equation	88
3.5 Transmission and reflection vs quarter wave plate configuration	93
3.6 Cascade analysis inputs	96
3.7 Cascade analysis vs network analyzer measurements	98
3.8 Reflection model of RF noise	104
5.1 Sensitivity to b_{Fierz} experimental setups	123
5.2 Sensitivity to b_{Fierz} constants	124
5.3 Sensitivity to main field determination	127
5.4 Sensitivity to linear β monitor drift	128
5.5 b_{Fierz} distortion from wall-effect	138
6.1 Fit results from ratio measurement	168
C.1 he6cres_db: run_log table	242
C.2 he6cres_db: spec_files table	243
C.3 he6cres_db: nmr table	243
C.4 he6cres_db: monitor table	244

ACKNOWLEDGMENTS

Firstly, thanks to Alejandro Garcia, my advisor. We first met in the Summer of 2015 when I was on a hiking trip to the North West. It's hard to believe that was over 8 years ago! Thank you for your patience, support, and belief in me as a researcher. You've devoted countless hours to troubleshooting hardware with me, thinking through analysis issues with me, and patiently explaining theory to me. And what I'm most struck by is how much fun it's been! Beyond being an excellent teacher and caring advisor, you truly love working on difficult problems and it's contagious. I think the single most valuable thing I've gained in graduate school is an appreciation for how fun it is to work on problems for which you can't look up the answers. I have you to thank for that.

Thanks to my undergraduate advisors Fred Wietfeldt and Scott Dewey. Fred seeing something in me during his office hours at Tulane University and helping me get introduced to physics research via the NIST REU was the catalyst that began my career trajectory. Later on, you connected me with Alejandro which has led to six formative and fun years of graduate work. I'm incredibly grateful for that. Scott, I still remember the day you showed me how to use the Linux command-line. You were giddy with excitement. I aspire to cultivate that level of excitement for my work.

Thanks to all of the core He6-CRES collaborators. Heather Harrington, Nick Buzinsky, RJ Taylor, Brent Graner, Winston DeGraw, and David McClain. We've had our highs and lows over the years but for the most part everyone has stayed positive and solutions oriented. That's no small feat given the number of setbacks we've encountered over the years. I'm truly blown away by what we've accomplished with such a small team. Heather, I can't wait to get a tour of your lab and sit in on one of your lectures once you become a professor

somewhere. Your tenacity towards your work and love for physics is epic and inspiring. You can always give me a call to chat about annoying systematics!

Thanks to the CENPA engineers, research staff, and postdocs who have helped me and our group immensely over the past six years. Eric Smith and Brittney Dodson truly went to battle daily to keep the CENPA accelerator alive and I can't thank them enough for that. Thanks to Gary Holman for being an absolute wizard at all things networking. Thanks to David Peterson, Nate Miedema, Ryan Roehnelt, Tom Burritt, Tim Van Wechsel, and John Amsbaugh for their engineering support of our experiment. It's been so fun talking through different problems with you over the years. Thanks for not laughing at my scratch paper "engineering" drawings. Thanks to Charles Hanretty and Grant Leum for helping us keep our very thirsty magnet full of LHe. Thanks to Clint Wiseman for getting me up to speed with CENPA rocks (the computing cluster) and for always dropping everything to help others despite having your own crushing pile of work.

Thanks to Derek for carefully reading through my thesis and providing valuable feedback. You are always (sometimes ruthlessly) boiling arguments (or systems, procedures, etc.) down to their simplest, clearest form. I aspire to do the same.

Thanks to my parents, Erin and Bill, who have always supported me unconditionally. I've never felt pressure from them to be different or better than I am and that's been the foundation from which I've been able to enjoy doing challenging things.

Thanks to my brother, Nick, who (according to some napkin-math) has likely spend over 600 hours talking to me over the phone within the past six years. There might not be anyone else who has been so informed on the day-to-day ups and downs of my PhD. I'm not sure you'll miss me complaining about the CENPA accelerator.

Thanks to the beloved group house, Grauman's Center, and all of it's amazing (past) inhabitants. Grauman's Center might be dead but it'll never be forgotten. Dan, Jesse, Linh, Lindsey, Corwin, Greg, and Jordan; you all will always feel like family. That said, it has

been nice not having to deal with Joe.

Finally, thanks to my partner and wife, Anna, for her immense support and patience throughout my graduate career. Leaving our sunny and comfortable New Orleans life to move to the dreary Pacific North West so that I would be guaranteed less than minimum wage for 6 years was, to say the least, a leap of faith. Thanks for bearing with me for the first year and thanks for all the good times after that.

DEDICATION

To Anna. Anna somehow reacted to a move across the country and 6 years of me being in graduate school with (almost) nothing but love and support. And to Graham. I'm not religious but Graham is unquestionably a miracle.

Chapter 1

THEORY AND MOTIVATION FOR ADVANCES IN PRECISION β SPECTROSCOPY

The study of nuclear β decay has played a crucial role in the development of the Standard Model (SM) of particle physics, which represents one of humanity's greatest scientific achievements. Radioactivity was first discovered in the late 1800's by Henri Becquerel in uranium, and Marie and Pierre Curie in thorium. A few years later Ernest Rutherford systematically organized radioactivity into alpha, beta, and gamma, based on their penetration depth in different materials [97]. Development from Fermi and others over the following decades revealed that the process of β decay was governed by what was deemed the weak interaction, one of the four fundamental forces of nature. In the late 1950's, following theoretical work by Lee and Yang suggesting the possibility, Wu *et. al.* confirmed experimentally that the weak interaction was parity violating by observing an asymmetry in the decay of polarized ^{60}Co [115]. This discovery shed first light on the peculiar chiral properties of the weak interaction. In the 1970's the electroweak theory was developed by Sheldon Glashow, Abdus Salam, and Steven Weinberg, unifying the seemingly distinct electromagnetic and weak interactions, and resulting in the 1979 Nobel Prize in Physics. The electroweak theory predicted the existence of the W and Z gauge bosons, which were both observed experimentally at CERN in 1983. The intimate interaction between the experimental study of the weak interaction, including the study of β decay, and the development of new theory continues to this day in our current search for beyond the Standard Model (BSM) physics.

In Section 1.1 of this chapter we present an overview of the basics of the SM with an emphasis on the charged current weak interaction that governs β decay and discuss how new physics would present itself in the weak Lagrangian. In Section 1.2 we discuss the theoretical

context of the Fierz interference term (b_{Fierz}), the specific low-energy observable we aim to measure. Section 1.3 covers the current experimental limits on the exotic currents we are sensitive to via a high precision measurement of b_{Fierz} . In Section 1.4 we cover the specific form of b_{Fierz} in terms of exotic coupling constants for our isotopes of interest (${}^6\text{He}$ and ${}^{19}\text{Ne}$). Finally, Section 1.5 provides a brief summary of the motivation for novel developments in precision β spectroscopy.

1.1 Weak interaction in the Standard Model

1.1.1 Standard Model overview

There are four known fundamental forces or interactions that govern our universe: gravitational, electromagnetic, weak, and strong¹. Each acts with a different strength and on different length scales. Gravity is at least 36 orders of magnitude weaker than the other three forces though unintuitively dominates the macroscopic behaviour of the universe due to the short range of the strong and weak forces ($\mathcal{O}(10^{-15} \text{ m})$), and the relative neutrality of the electromagnetic charge of macroscopic objects like planets and stars.

The SM of particle physics is built with the mathematical framework of Quantum Field Theory and comprehensively and accurately describes the interactions between all subatomic particles over a wide range of energies. The SM includes the electromagnetic, weak, and strong interaction but does not provide an adequate description of gravity. Many textbooks, such as References [87, 48], provide a detailed account of the SM.

The set of Standard Model fundamental particles, as shown in Figure 1.1, consists of three generations of quarks (and their associated anti-particles), three generations of leptons (and their associated anti-particles), the gauge bosons (photon, gluon, W^\pm , and Z^0), and the relatively recently discovered Higgs boson [32]. Spin- $\frac{1}{2}$ fermions provide the building blocks for matter within the SM and there are four fundamental interactions between them. As was stated above, the SM does not provide a satisfactory treatment of gravity. The three

¹Technically, one could argue, there are only three in that the weak and electromagnetic are unified into the electroweak force at high energies.

	← spin-1/2 particles →			← bosons →	
mass charge spin Quarks Leptons	2.4 MeV/c² +2/3 1/2 up	1.3 GeV/c² +2/3 1/2 charm	173 GeV/c² +2/3 1/2 top	0 0 1 gluon	125 GeV/c² 0 0 Higgs
	2.4 MeV/c² -1/3 1/2 down	0.1 GeV/c² -1/3 1/2 strange	4.2 GeV/c² -1/3 1/2 bottom	0 0 1 photon	
	0.5 MeV/c² -1 1/2 electron	0.1 GeV/c² -1 1/2 muon	1.8 GeV/c² -1 1/2 tau	90 GeV/c² 0 1 Z	
	≈ 1.1 eV/c² 0 1/2 neutrino 1	≈ 1.1 eV/c² 0 1/2 neutrino 2	≈ 1.1 eV/c² 0 1/2 neutrino 3	80 GeV/c² +1/-1 1 W^{+/-}	

Figure 1.1: The fundamental particles of the Standard Model. Figure created by Alejandro Garcia of He6-CRES.

remaining forces (strong, weak, and electromagnetic) described by the SM are generated via the coupling of fermions to gauge fields. The gauge fields are associated with different local gauge symmetries, meaning the gauge fields are responsible for imposing the local symmetries within the Lagrangian. As a concrete example, consider a charged particle coupled to an electromagnetic field. Applying a $U(1)$ gauge transformation ($\psi \rightarrow \psi e^{iQ\lambda(x)}$) and the appropriate transformation to the potentials leaves the Lagrangian invariant ($\mathcal{L} \rightarrow \mathcal{L}$). The gauge bosons of the SM, also called force mediators or carriers, are the quanta of the gauge fields.

The photon, the quanta of the the electromagnetic field, mediates the electromagnetic interaction, which is symmetric under $U(1)$ gauge transformations. Similarly, gluons are the quanta of the strong field and they mediate the strong interaction which is symmetric under $SU(3)$ gauge transformations. The pattern is broken in the case of the weak interaction because the weak force mediators (W^\pm , and Z^0) are massive and within the framework of a local gauge invariance, mediating particles must be massless. In this way a pure left-handed $SU(2)$ gauge theory fails to adequately describe the weak interaction. This apparent conflict was resolved by Sheldon Glashow, Abdus Salam, and Steven Weinberg, when they developed a theory that unifies the electromagnetic and weak interactions into an $SU(2)_L \otimes U(1)$ gauge theory, where the L subscript refers to the fact that the weak interaction only couples left-handed particles. In the electroweak theory there are four fields ($W^i (i = 1, 2, 3)$ and B) associated with the un-broken $SU(2)_L \otimes U(1)$ symmetry which mix with one another after the spontaneous breaking of this symmetry via the Higgs mechanism and result in the photon and the three massive fields (W^\pm , and Z^0) [55]. The Higgs mechanism explains the process by which the symmetric, unified electroweak interaction breaks into the two very distinct interactions we observe today by introducing the Higgs field² The massless fields ($W^i (i = 1, 2, 3)$ and B) become massive (except for the photon) through the interaction with

²Note that the $W^i (i = 1, 2, 3)$ particles have charges $(0, \pm 1)$ and it is only the neutral W^i that mixes with B, the $U(1)$ boson, to generate Z^0 and the photon. This is why the neutral component of the weak interaction is not left handed.

the Higgs field. Fermions also interact with the Higgs field through their weak isospin (the spin associated with the $SU(2)_L \otimes U(1)$ symmetry) which gives them mass. The Higgs field implies an associated boson, the Higgs boson, which until its discovery in 2012 at the Large Hadron Collider (LHC), represented the last major unobserved component of the model [32]. The SM, outlined above, has been successful in describing almost all experimental particle physics observations. Its limitations will be discussed in Subsection 1.1.3.

The proton (uud) and neutron (udd) each consist of three up or down quarks, and constitute the vast majority of the visible mass in the universe. Bound systems of protons and neutrons form nuclei and these nuclei bind to leptons (electrons for the most part) to form atoms. These systems (nuclei and atoms) transform and interact with one another via the strong or electroweak interactions. In the SM, β^- (β^+) decay is described as the transformation of a down (up) quark to an up (down) quark via the emission of a W^- (W^+) boson, which quickly decays into an electron (positron) and anti-neutrino (neutrino).

1.1.2 Charged-current weak interactions in the Standard Model

Precision measurements of β decay observables such as nuclear decay lifetimes, Ft values, spectral distortions, and angular correlations are strong probes for physics beyond the SM since these low-energy observables are readily calculated to high accuracy via perturbation theory. Here we develop the mathematical formalism necessary to understand how these observables are calculated and derive the generic Standard Model differential decay rate for β decays.

Fermi's golden rule states that the decay rate Γ from an initial state i to final state f is given by

$$\Gamma_{if} = \frac{2\pi}{\hbar} \int |\mathcal{M}_{if}|^2 dPS^{(3)}, \quad (1.1)$$

where $dPS^{(3)}$ is the kinematic phase space available to the incoming and outgoing particles, $\mathcal{M}_{if} = \langle f | \mathcal{H}_{int} | i \rangle$ is referred to as the matrix element or amplitude, and \mathcal{H}_{int} is the inter-




Figure 1.2: The quark level Feynman diagram associated with neutron β^- decay.

action Hamiltonian governing the interaction [47] [48]. The phase space term depends on the masses, energies, and momenta of the incoming and outgoing particles and contains the kinematic constraints of the interaction such as conservation of mass-energy. The intuition being that how likely a process is to occur scales with the number of kinematically allowed final states. An extreme example of this is a particle decay where the products outweigh the particle. In such a case there is no available phase space (given the initial particle is at rest) and the decay rate is 0. The matrix element \mathcal{M}_{if} contains the dynamical, model dependent information and we calculate it using the Feynman diagrams and corresponding Feynman rules associated with the interaction. For nuclear β -decay, this term specifies the permitted interactions of the effective weak interaction Hamiltonian \mathcal{H}_{int} .

We use the effective weak interaction Hamiltonian in Equation 1.1 because, despite the weak force acting on quarks at a base level, it is currently challenging to calculate observables in nuclear β decays using quarks as fundamental fields. Over the next few paragraphs we will work from the quark level interaction towards an understanding of the effective weak interaction Hamiltonian. The relevant quark level process for β^- decay is $d \rightarrow ue\bar{\nu}_e$. The quark level Feynman diagram for neutron β^- decay is shown in Figure 1.2.

The propagator associated with the massive W boson is given by

$$\frac{-i(g_{\mu\nu} - (p_\mu p_\nu)/M_W^2)}{p^2 - M_W^2}, \quad (1.2)$$

where $g_{\mu\nu}$ is the metric ($g_{\mu\mu} = -1$, $g_{\mu\nu} = 0$ for $\mu, \nu \in 0,1,2,3$. Except $g_{00} = 1$), and where M_W , the mass of the W boson, is $\approx 80 \text{ GeV}/c^2$. Since the typical momenta of the particles involved in β decays are at the MeV level ($p^2 \ll M_W^2$) we can treat the W propagator as a constant and simply use

$$\frac{ig_{\mu\nu}}{M_W^2}. \quad (1.3)$$

The charged weak vertex factor is given by

$$\frac{-ig_w}{2\sqrt{2}}\gamma^\mu(1 + \gamma^5), \quad (1.4)$$

where g_w is the weak charge current coupling constant ($g_w = 0.78$), and γ^μ is the μ -th ($\mu = 0,1,2,3$) gamma matrix (see Reference [48] or [47] for their explicit definitions), and $\gamma^5 = i\gamma^1\gamma^2\gamma^3\gamma^0$. In this thesis we adopt the Weyl representation of the Dirac matrices. The left-handed projection operator is $(1 + \gamma^5)/2$ and its presence in the charged weak vertex factor (Equation 1.4) means that only left-handed charged currents are involved in weak interactions. Putting together our propagator and our vertex factor, and assuming $p^2 \gg M_W^2$, we get the following interaction Lagrangian,

$$\mathcal{L}_{int} = -\frac{G_F V_{ud}^{CKM}}{\sqrt{2}}(\bar{u}\gamma_\mu(1 + \gamma^5)d)(\bar{e}\gamma^\mu(1 + \gamma^5)\nu_e), \quad (1.5)$$

where V^{CKM} is the Cabibbo-Kobayashi-Maskawa (CKM) quark-mixing matrix which relates the mass and weak eigenstates of quarks, and G_F is the Fermi coupling constant for the *point interaction* Lagrangian which is defined as

$$G_F = \frac{\sqrt{2}g_w^2}{8M_W^2} = 1.166 \times 10^{-5} \text{ GeV}^{-2} \quad (1.6)$$

in natural units. The ‘weak fine structure constant’ is given by

$$\alpha_w = \frac{g_w^2}{4\pi} \approx \frac{1}{30}. \quad (1.7)$$

Comparing the electromagnetic fine structure constant ($\alpha_e = \frac{1}{137}$) to charge weak interaction fine structure constant we see that, against intuition, the weak fine structure constant is larger. Equation 1.6 illustrates that despite the larger intrinsic coupling constant, the weak interaction is weaker (by roughly 3 orders of magnitude) than the electromagnetic interaction due to the large mass of the mediator. It's important to note that while g_w is dimensionless, G_F has units of E^{-2} , and this is indicative of the fact that the point interaction theory of β decay is not renormalizable. For tree-level calculations of β decay, which includes most of what we are concerned with in this thesis, the Lagrangian shown in Equation 1.5 works well.

Even when using the simplified point interaction Lagrangian of Equation 1.5, it is not often practical or possible to calculate observables using the quark-level interactions due to the complications that the strong force within the nucleus introduces. Since quarks are always bound inside a nucleon via the strong interaction and the energies involved in β decay are small compared to these strong force binding energies, we instead use the nucleons (protons and neutrons) as the fundamental fields. Below is the effective field theory Lagrangian:

$$\mathcal{L}_{int}^{Eff} = -\frac{G_F V_{ud}^{CKM}}{\sqrt{2}} (g_V \bar{p} \gamma_\mu n - g_A \bar{p} \gamma_\mu \gamma^5 n) (\bar{e} \gamma^\mu (1 + \gamma^5) \nu_e), \quad (1.8)$$

which we can use to make calculations of observables like the decay rate. In the effective Lagrangian of Equation 1.8 the form factors g_V , and g_A encode the vector and axial-vector current corrections due to the strong interactions in the nucleons. Note that for the quark-level Lagrangian of Equation 1.5, we have $g_V = +1$, and $g_A = -1$, a pure V-A form, but that this is not true in general once we treat the proton and nucleon as fundamental fields. In practice the form factors are determined through a combination of experimental and theoretical efforts. To first order in the up and down quark mass difference squared, g_V is consistent with unity [28]. The fact that the weak vector current is unaffected by the strong interactions in the nucleus is referred to as the Conserved Vector Current (CVC) hypothesis

[106]. Recent lattice QCD calculations have determined g_A for neutron decay to better than 1%, and to this level there is accordance with the most up to date experimental value of -1.2723 ± 23 [10, 82]. The fact that the value of g_A only deviates marginally from -1 , is referred to as the partially conserved axial current (PCAC). In general it is not guaranteed that g_A for other nuclei is the same as the g_A observed for neutron decay, but in the case of light nuclei like ${}^6\text{He}$, where the two-body currents and nuclear structure are simpler, the value of g_A is consistent with the neutron value [68].

Now we can use our Lagrangian from Equation 1.8 and Fermi's Golden Rule (Equation 1.1) to work towards the generic Standard Model differential decay rate for β decays. Note that the interaction Hamiltonian in Equation 1.1 is given by: $\mathcal{H}_{int} = -\mathcal{L}_{int}^{Eff}$. We refactor the matrix element ($\mathcal{M}_{if} = \int \langle f | \mathcal{H}_{int} | i \rangle d^3x$) into leptonic and hadronic currents, as

$$\mathcal{M}_{if} = \frac{G_F V_{ud}^{CKM}}{\sqrt{2}} \int (J_{N\mu}^V - J_{N\mu}^A) J_l^\mu d^3x, \quad (1.9)$$

where $J_{N\mu}^V$ and $J_{N\mu}^A$ are respectively referred to as the vector current and axial current because of the operators they contain (the N subscript indicates nucleon) and J_l^μ is referred to as the leptonic current. Their specific forms are,

$$J_l^\mu = \langle e\bar{\nu} | \bar{e}(x)\gamma^\mu(1 + \gamma^5)\nu_e(x) | 0 \rangle, \quad (1.10)$$

$$J_{N\mu}^V = \langle f | g_V \bar{p}(x)\gamma_\mu n(x) | i \rangle, \quad (1.11)$$

$$J_{N\mu}^A = \langle f | g_A \bar{p}(x)\gamma_\mu\gamma^5\nu_e(x)n(x) | i \rangle. \quad (1.12)$$

The leptonic current (Equation 1.10) represents the transition from vacuum to the outgoing electron and anti-neutrino and looks similar for the case of β^+ decay. The initial and final states of the vector current (Equation 1.11) and axial current (Equation 1.12) are the nuclear states of the parent and daughter nuclei.

The final lepton states are taken to be plane waves such that the state's x dependence is given by $\exp(i(p_e + p_\nu)x)$. Since β decay provides at most MeV scale energy to the leptons, the lepton wavelengths are $\mathcal{O}(100)$ times larger than the size of the nucleus and therefore the approximation $\exp(i(p_e + p_\nu)x) \approx 1$ is appropriate. This implies that the leptonic current can be pulled out of the integral in Equation 1.9.

The integrals of Equation 1.10 and Equation 1.12 can be completed by utilizing the fact that both the parent and daughter nucleon kinetic energy is small compared to it's mass. We quote the final form in the limit that the $\mathcal{O}(v/c)$ terms go to zero here (see Reference [57] for a detailed derivation of these integrals),

$$\lim_{\mathcal{O}(v/c) \rightarrow 0} \int J_{N\mu}^V d^3x = \begin{cases} 2\sqrt{M_i M_r} M_F & \mu = 0 \\ 0 & \mu = 1, 2, 3 \end{cases} \quad (1.13)$$

$$\lim_{\mathcal{O}(v/c) \rightarrow 0} \int J_{N\mu}^A d^3x = \begin{cases} 0 & \mu = 0 \\ 2\sqrt{M_i M_r} M_{GT} & \mu = 1, 2, 3 \end{cases} \quad (1.14)$$

where M_i and M_r are respectively the masses of the parent and daughter nuclei, and M_F and M_{GT} are respectively called the Fermi and Gamow-Teller matrix elements. Nuclear β decays involving purely vector currents are called Fermi decays and those involving purely axial currents are called Gamow-Teller decays. Decays involving both currents, such as for neutrons and ^{19}Ne , are called mixed decays.

These two different types of decays have different selection rules. For Fermi decays $\Delta J = 0$ meaning that the parent nucleus cannot undergo a spin flip. For Gamow-Teller decays $\Delta J = \pm 1, 0$ but not $J = 0 \rightarrow 0$. These selection rules result in the $|e\nu\rangle$ final lepton state either coupling to the singlet $S = 0$ state (Fermi) or the triplet $S = 1$ state (Gamow-Teller). They can mathematically be understood by the fact that in the non-relativistic approximation the vector currents do not contain the Pauli spin operators, while the axial currents associated with the Gamow-Teller decay do contain Pauli spin operators, which allows them to instigate a nuclear spin flip transition.

The matrix elements shown in Equation 1.13 and Equation 1.14 only include the leading order terms and as such are only approximately correct. For meaningful comparison to high precision measurements, higher order corrections (radiative, recoil, etc.) must be accounted for. These corrections are discussed in Subsection 1.2.3.

Combining our above results, we obtain the (leading order) Standard Model differential decay rate,

$$d\Gamma = \frac{1}{2M_i} |\mathcal{M}_{if}|^2 (2\pi)^4 \delta^{(4)}(p_i - p_f - p_e - p_\nu) \frac{d^3 p_e}{(2\pi)^3 2E_e} \frac{d^3 p_\nu}{(2\pi)^3 2E_\nu} \frac{d^3 p_f}{(2\pi)^3 2E_f} \quad (1.15)$$

where E_f and p_f are the energy and 3-momentum of the daughter recoil nucleus. The matrix element is given by

$$\mathcal{M}_{if} = \frac{G_F V_{ud}^{CKM}}{\sqrt{2}} (g_V M_F \bar{e} \gamma^0 (1 + \gamma^5) \nu_e - g_A M_{GT} \bar{e} \gamma^k (1 + \gamma^5) \nu_e), \quad k = 1, 2, 3. \quad (1.16)$$

In Subsection 1.2.1 we will discuss how to derive the shape of the β decay spectrum from the generic decay rate of Equation 1.15.

1.1.3 Charged-current weak interactions beyond the Standard Model

The SM has been extremely successful in describing the vast majority of modern experimental observations but we know it has shortcomings and it is sensible to think it will someday be superseded by a more complete theory of particle physics. The SM does not properly explain gravity, it does not account for dark energy or matter, and it does not explain the origin of the neutrino masses. In addition, there are mounting discrepancies between experimental observations and the Standard Model, for example the Muon g-2 experiment measured a value for the positive muon magnetic anomaly in 2021 that is 4.2σ from the SM prediction [1]. There are also nineteen free parameters in the SM (excluding the neutrino masses and mixing angles) such as the CKM matrix, the mass of the Higgs, and the coupling of the Higgs for

which we rely solely on experiment, and it is hoped that a more advanced theory will explain the particular values of at least some of these parameters [47, 86].

The field of particle physics is currently driven by symbiotic experimental and theoretical efforts. On the one hand theorists are working to propose deeper theories and articulate the experimental observables associated with these new theories. On the other hand, experimentalists are working to detect or measure the observables associated with the best motivated theories. In the absence of a novel observation, these experimental efforts put important constraints on the parameter space of possible theories, thus feeding back into and directing the theoretical efforts.

The experimental effort described in this thesis fits into this framework by searching for signatures of new physics via the precision measurement of β decay spectra. As was stated above, calculations of β decay observables can be calculated accurately with perturbation theory and therefore precise measurements are easily compared to accurate SM predictions. In the following paragraphs we present the theory necessary to understand how physics beyond the SM would influence β decay observables.

We begin with a phenomenological (model independent) approach and write out all possible weak currents. In Subsection 1.1.2 we discussed that the charged weak current in the SM is purely (V-A) in nature, but in general there are other currents that could be present at small levels being introduced by new physics. There are five irreducible bilinear forms allowed by Lorentz invariance, each with an associated current. With these five forms we build the most general form of the weak interaction Hamiltonian (originally written as such by Jackson, Treiman, and Wyld in Reference [61]),

$$\mathcal{H}_{int}^{(V,A)} = \sum_{i=V,A} (\bar{p}O^i n) ((C_i + C'_i)\bar{e}^L O_i \nu_e^L + (C_i - C'_i)\bar{e}^R O_i \nu_e^R) \quad (1.17)$$

$$\mathcal{H}_{int}^{(S,T,P)} = \sum_{i=S,T,P} (\bar{p}O^i n) ((C_i + C'_i)\bar{e}^R O_i \nu_e^L + (C_i - C'_i)\bar{e}^L O_i \nu_e^R) \quad (1.18)$$

where the L and R superscripts indicate the handedness (chirality) of the lepton field, the

C_i s are coupling constants, and the operators O_i s are

$$O_S = 1, \quad O_P = \gamma^5, \quad O_V = \gamma_\mu, \quad O_A = i\gamma_\mu\gamma^5, \quad O_T = -i(\gamma_\mu\gamma_\nu - \gamma_\nu\gamma_\mu)/(2\sqrt{2}). \quad (1.19)$$

The currents ($\bar{p}O^i n$) corresponding to each operator are respectively referred to as scalar, pseudo-scalar, vector, axial (or equivalently pseudo-vector or axial-vector), and tensor currents. See Section III of Reference [57] for a full derivation of the generic interaction Hamiltonian of Equation 1.17 and Equation 1.18 and see Chapter 9.7 in Reference [47] for a more detailed treatment of the chiral properties of the weak interaction.

Within the SM there are well defined constraints on the coupling constants:

- $C_V = C'_V$ and $C_A = C'_A$.
- C_V and C_A are real.
- $C_S = C'_S = C_T = C'_T = C_P = C'_P = 0$.

From the above constraints and Equation 1.17, and Equation 1.18 we see that the SM weak interaction contains only vector and axial currents, has maximal parity violation, and only includes left handed leptons. The values of C_V and C_A relate linearly to the coupling constants g_V and g_A presented in Subsection 1.1.2 (see Reference [20] for the explicit relations).

The chiral components of a dirac spinor are obtained via the application of the left and right projection operators,

$$\psi = P_L\psi + P_R\psi = \psi^L + \psi^R, \quad P_{L/R} = (1 \pm \gamma^5)/2. \quad (1.20)$$

The chiral properties of the different currents of Equation 1.17 and Equation 1.18 originate from the way a projection operator commutes through the current [57]. From Equation 1.17 and Equation 1.18 we see that vector and axial currents couple only like-handed

states whereas scalar, tensor and pseudo-scalar currents exclusively couple states of opposite handedness.

In the same way that the SM nuclear matrix elements were calculated in Subsection 1.1.2 (see Equation 1.13 and Equation 1.14), they can be calculated for the non-SM currents. In the non-relativistic limit, where terms of $\mathcal{O}(v/c)$ are taken to zero, we find that the leading order matrix element for the scalar current is identical to that of the vector current (M_F) and that the leading order matrix element for the tensor current is identical to that of the axial-vector current (M_{GT}). Intuitively, this can be understood as arising from the fact that the axial-vector and tensor bispinor representations in the low energy limit contain spin operators of nucleons (Pauli matrices) and are thus able to cause a nuclear spin flip, while the scalar and vector currents do not contain spin operators of nucleons and so cannot cause a nuclear spin flip. A detailed derivation can be found in the Appendix of Reference [57].

In practice this means that precision measurements of Fermi (Gamow-Teller) channel β decay observables provides sensitivity to vector and scalar (axial-vector and tensor) currents. These relations are summarized in Table 1.1.

	Fermi	Gamow-Teller
relate same handed leptons (SM)	V	A
relate opposite handed leptons (possible new physics)	S	T

Table 1.1: Summary of vector (V), axial-vector (A), scalar (S), and tensor (T) currents in terms of their chiral properties and which channel of β decay is sensitive to them.

The leading order term of the pseudoscalar current is of $\mathcal{O}(v/c)$ so low energy β decay experiments have suppressed sensitivity to such currents. Recent calculations have shown that g_{PS} is surprisingly large (≈ 350) and thus this partially overwhelms the $\mathcal{O}(v/c)$ suppression factor [43]. Despite this concern, there are already strong constraints on pseudoscalar currents imposed by pion decay experiments and therefore we are justified in neglecting their

presence in the majority of this thesis [108].

The generic β decay Hamiltonian presented in Equation 1.17 and Equation 1.18 sheds light on the types of decays that provide sensitivity to specific currents in a model independent way. In the context of the discussion at the beginning of this Subsection, it is worth asking if there is a theoretical basis for searching for new physics in the form of chirality-flipping interactions (scalar and tensor currents) in the weak interaction. Many models proposed to supersede the SM such as supersymmetric models, grand unified theory models, and left-right symmetric models all naturally produce chirality-flipping interactions [8, 56, 73, 92, 90].




Figure 1.3: Example Feynman diagrams of effective chirality-flipping interactions in models beyond the SM. Top: leptoquarks suggested by grand-unified models with electric charge of $+2/3$ on the left and $-1/3$ on the right. Bottom: loop diagrams suggested by supersymmetric models: \tilde{u} , \tilde{d} , \tilde{l} represent squarks and sleptons, while the $\tilde{\chi}$ represent charged and neutral supersymmetric particles. This figure is taken from Reference [90]).

Examples of Feynman diagrams of such hypothetical interactions are shown in Figure 1.3. Recent reviews covering the motivation for the presence of exotic currents in the weak interaction are given in References [21, 107, 45, 103].

1.2 Energy spectrum and b_{Fierz}

1.2.1 Energy spectrum

We now want to work from the most general decay rate expression shown in Equation 1.15 towards the expression for the energy spectrum ($W(E_e) \equiv \frac{d\Gamma}{dE_e}$). We only care about the electron energy spectrum so we begin by integrating over d^3p_r and E_ν , which results in an expression for $W(E_e, \Omega_e, \Omega_\nu)$. See Chapter 9.3 of Reference [47] for the details of the algebra.

$$\begin{aligned} W(E_e, \Omega_e, \Omega_\nu) &= \frac{d\Gamma}{dE_e d\Omega_e d\Omega_\nu} \\ &= |\mathcal{M}_{if}|^2 \frac{1}{(2\pi)^5 16 M_i E_e E_\nu E_r} \frac{E_e p_e E_\nu^2}{1 + (E_\nu + p_e \cos(\theta_{\beta\nu}))/M_r}, \end{aligned} \quad (1.21)$$

where $\theta_{\beta\nu}$ is the angle between the momenta of the outgoing β and neutrino. The neutrino energy E_ν in Equation 1.21 is constrained by conservation of energy-momentum,

$$E_\nu = -(p_e \cos(\theta_{\beta\nu}) + M_r) + \sqrt{(p_e \cos(\theta_{\beta\nu}) + M_r)^2 + 2(M_i - M_r - E_e)M_r - p_e^2}, \quad (1.22)$$

and the energy of the recoil ion is then,

$$E_r = \sqrt{M_r^2 + (p_e + p_\nu)^2}. \quad (1.23)$$

The energy released in β decay is $\mathcal{O}(1 \text{ MeV})$ and the mass-energy of even very light initial and final nuclei is $\mathcal{O}(1 \text{ GeV})$, so we can further simplify the expression of Equation 1.21 by taking $M_i, M_r \rightarrow \infty$ while maintaining $M_i - M_r$ as the energy available to the β decay. This limit effectively reduces the 3-body phase space to a 2-body phase space. In this limit $E_r \approx M_r$ which means that the neutrino energy reduces to

$$E_\nu \approx M_i - M_r - E_e \quad (1.24)$$

and so $W(E_e, \Omega_e, \Omega_\nu)$ is reduced to

$$W(E_e, \Omega_e, \Omega_\nu) = |\mathcal{M}_{if}|^2 \frac{1}{(2\pi)^5 16 M_i E_e E_\nu E_r} \sqrt{E_e^2 - m_e^2} E_e (M_i - M_r - E_e)^2, \quad (1.25)$$

where we have used the relation $p_e = \sqrt{E_e^2 - m_e^2}$.

Next we integrate over all outgoing angles of the leptons (electron and anti-neutrino), then we average over the initial nuclear spin states and sum over the helicities of the outgoing leptons. This results in the β -energy spectrum (without any higher order corrections),

$$W(E_e) \equiv \frac{d\Gamma}{dE_e} = \frac{1}{(2\pi)^3} \sqrt{E_e^2 - m_e^2} E_e (E_0 - E_e)^2 \xi \left(1 + b_{Fierz} \frac{m_e}{E_e}\right), \quad (1.26)$$

where $E_0 = M_i - M_r$, and ξ and b_{Fierz} are functions of the nuclear matrix elements and coupling constants in the nucleon level charged weak Hamiltonian[47, 48, 57]. The exact form of ξ and b_{Fierz} will be discussed in Subsection 1.2.2. The $16M_i E_e E_\nu E_r$ in the denominator of Equation 1.21 is cancelled by the normalization factors of $|\mathcal{M}_{if}|^2$. The form of $W(E_e)$ shown in Equation 1.26 is naive in that it lacks any correction factors. The correction terms necessary to meaningfully compare a β spectrum with 10^{-3} relative uncertainty to the Standard Model are presented in Subsection 1.2.3.

1.2.2 Fierz interference term

The term $b_{Fierz} \frac{m_e}{E_e}$ is referred to as the Fierz interference term because it arises from the cross terms between the SM currents (V,A) and the beyond the Standard Model (BSM) currents (S,T) and therefore in the SM $b_{Fierz} = 0$ (no exotic currents in SM). As will be illustrated, precise measurements of b_{Fierz} are a powerful probe of new physics because b_{Fierz} is linear in the exotic coupling constants (C_S, C_T). This is in contrast to many other β decay observables, such as the $\beta - \nu$ angular correlation ($a_{\beta\nu}$) for example, that are quadratic in the exotic coupling constants and are therefore significantly less sensitive to their presence. Progress towards a 10^{-3} measurement of b_{Fierz} via a novel experimental technique is the main topic of this thesis.

When evaluating $|\mathcal{M}_{if}|^2$ in order to make the step from Equation 1.25 to Equation 1.26, we take the product of the transition matrix element and its hermitian conjugate $\langle f | \mathcal{H}_{int} | i \rangle \langle f | \mathcal{H}_{int}^* | i \rangle$ where \mathcal{H}_{int} is generically given by Equation 1.17 and Equation 1.18 [57]. The self-interacting terms such as $\langle f | \mathcal{H}_{int}^j | i \rangle \langle f | \mathcal{H}_{int}^{j*} | i \rangle$ for $j = V, A, S, T$ give rise to the ξ term in Equation 1.26 where ξ is,

$$\begin{aligned} \xi = & |M_F|^2 (|C_S|^2 + |C_V|^2 + |C'_S|^2 + |C'_V|^2) \\ & + |M_{GT}|^2 (|C_T|^2 + |C_A|^2 + |C'_T|^2 + |C'_A|^2). \end{aligned} \quad (1.27)$$

Cross terms, or interference terms, of the form $\langle f | \mathcal{H}_{int}^{V(A)} | i \rangle \langle f | \mathcal{H}_{int}^{S(T)*} | i \rangle$ give rise to the Fierz interference term in Equation 1.26 [57, 35]. As ${}^6\text{He}$ is experimentally one of our main isotopes of interest, we consider Gamow-Teller decays so that only axial and tensor currents contribute. We set the neutrino mass to zero, meaning that any term that involves neutrinos with both chirality states will vanish because one of the two chiral projections will result in a zero. Substituting the form of the interaction Hamiltonian (Equation 1.17 and Equation 1.18) into the interference term, the resulting non-zero terms are,

$$\begin{aligned} & (C_A + C'_A)(C_T + C'_T) \langle f_l | \bar{e}^L O_A \nu_e^L | i_l \rangle \langle i_l | \bar{\nu}_e^L O_T e^R | f_l \rangle \times (\text{hadronic current}) \\ & + (C_A - C'_A)(C_T - C'_T) \langle f_l | \bar{e}^R O_A \nu_e^R | i_l \rangle \langle i_l | \bar{\nu}_e^R O_T e^L | f_l \rangle \times (\text{hadronic current}) \\ & + h.c. \end{aligned} \quad (1.28)$$

where $\langle i_l |$ and $\langle f_l |$ are the initial and final lepton states respectively³, and *h.c.* stands for the hermitian conjugate. Notice that regardless of the helicity of state $\langle f_l |$ in Equation 1.28, there will be a projection of both e^R and e^L , meaning both the left-handed ($\sqrt{1 - p_e/E_e}$) and right-handed ($\sqrt{1 + p_e/E_e}$) amplitudes will be projected out. This means that regardless of the specifics of the coupling constant algebra, all terms in Equation 1.28 will be proportional to

³For our case of interest (β decay) $\langle i_l | = \langle 0 |$, the vacuum state.

$$\sqrt{1 - p_e/E_e}\sqrt{1 + p_e/E_e} = \sqrt{1 - (p_e/E_e)^2} = \frac{m_e}{E_e}. \quad (1.29)$$

The same $\frac{m_e}{E_e}$ is present in a pure Fermi decay where only vector and scalar currents are allowed. The complete expression for b_{Fierz} is obtained by summing over the final electron helicity, and calculating the leptonic and hadronic matrix elements in the non-relativistic limit. See Section IV of Reference [57] for a more detailed treatment. We find that b_{Fierz} is given by,

$$b_{Fierz}\xi = \pm 2(|M_F|^2(C_S C_V + C'_S C'_V) + |M_{GT}|^2(C_T C_A + C'_T C'_A)) \quad (1.30)$$

where the C_i coefficients are taken to be real (no time reversal violation), M_{GT} is the reduced matrix element common to the M_{GTi} 's ($i = 1, 2, 3$), and the upper (lower) sign applies to β^- (β^+) decay.

1.2.3 Energy spectrum corrections

For a high precision extraction of b_{Fierz} from the β spectrum, several higher order corrections must be applied to the β -energy spectrum $W(E_e)$ shown in Equation 1.26. There are in principle many ways to classify and organize these effects but we follow the conventions established in References [52] and [51]. The three largest effects that must be accounted for in order to reach our current target b_{Fierz} precision of 10^{-3} are the Coulomb, radiative, and recoil effects which are discussed below. See Reference [52, 51] for a complete overview of important spectral corrections.

To correct for the distortion of the wave function of the outgoing β s caused by the Coulomb field of the daughter nucleus to order α (the fine structure constant), the uncorrected energy spectrum is multiplied by the Fermi Function $F(\pm Z, E_e)$ (upper (lower) sign applies to β^- (β^+) decay) [114],

$$F(\pm Z, E_e) = \frac{2(1 + S)}{\Gamma(1 + 2S)^2} (2p_e \rho)^{2S-2} e^{\pi\nu} |\Gamma(S + i\nu)|^2, \quad (1.31)$$




Figure 1.4: Feynman diagrams corresponding to the order- α radiative corrections (single photon or virtual photon). The top three diagrams involve virtual photons and the bottom two involve real photons. This figure is taken from Reference [52].

where $S = \sqrt{1 - \alpha^2 Z^2}$, $\nu = \alpha Z E_e / p_e$, and α is the fine structure constant, Z is the proton number in the daughter nucleus, and ρ is the radius of the nucleus. The Fermi function represents a 1 – 2% energy dependent correction over the beta spectrum. See Section III of Reference [52] for an in-depth description of the Fermi function.

The Coulomb effect, accounted for by the Fermi function of Equation 1.31, is not the only quantum electrodynamics (QED) process at play as the β -particle exits the nucleus. There are also photons, real and virtual, being emitted and exchanged and their effect on the β -energy spectrum are generically referred to as radiative corrections. There is a distinction made between inner radiative corrections, which are fundamental to the weak interaction and are often incorporated into the effective coupling constants, and the outer radiative corrections, which are energy and nucleus dependent. We are primarily concerned with the outer radiative corrections as these are what differ between ${}^6\text{He}$ and ${}^{19}\text{Ne}$, our decays of

interest.

Outer radiative corrections are further differentiated based on how many photons (real or virtual) are involved in their Feynman diagrams. Order- α corrections to the tree level interaction involve a single photon, order- α^2 corrections involve two photons, and so on. Figure 1.4 illustrates all of the order- α Feynman diagrams. The top three diagrams involve virtual photons and the bottom two involve real photons. Typically the real photons are broken up into soft photons, which are very low energy and are hard to distinguish from β s in experiments using standard solid state detectors, and hard photons, which are higher energy and could be vetoed. In a CRES experiment (introduced in Chapter 3), only β s are detected; there is no risk of misidentifying a low energy photon as a β . Therefore we must account for both soft and hard real photons in the β spectrum used in our Monte Carlo calculations. See Section V of Reference [52] for a complete description of the radiative effects.

Recoil order terms account for the finite size and mass of the nucleus. The derivation of the uncorrected β spectrum of Equation 1.26 assumed infinite nuclear mass ($M_i, M_r \rightarrow \infty$) and infinitesimal physical extent, both of which are only approximately true. Attributing a finite mass to the nucleus turns the two-body kinematics to three-body kinematics and effects not only the energy of the leptons but also the Coulomb effect itself now that the charge of the nucleus is not taken as static. Accounting for the finite size of the nucleus necessitates corrections on the S matrix as the nuclear potential is no longer a simple δ -function. See Section IV of Reference [52] for a full treatment of the recoil order effects.

Correctly accounting for the Coulomb, radiative, and recoil effects described above is sufficient for a 10^{-3} measurement. However, Hayen et al. ([52, 51]) provide a python implementation of their fully corrected beta spectra (<https://github.com/leenderthayen/thecobs>) that should account for corrections to the 10^{-4} level. We use these fully corrected spectra (via their GitHub repository) in our Monte Carlo simulations (see Chapter 5) so that as corrections are updated or reconsidered our simulation efforts will be relatively effortlessly up-to-date and as precise as the current nuclear theory allows.

1.3 Experimental limits on exotic currents

Observables such as the decay lifetime, angular correlation coefficients, and the shape of the β -energy spectrum (as we are specifically interested with in this thesis) have been measured experimentally for the neutron and for various nuclei since the late 1950's and these measurements provide constraints on the strength of scalar and tensor couplings in the weak interaction. Reference [20] provides a detailed overview of the best aggregated and single-experiment limits on exotic currents set by low energy measurements and compares these limits to those set by the high energy experiments of the LHC. More in depth reviews of the low energy experiments and their corresponding constraints on exotic currents can be found in References [108, 103]. In order to meaningfully compare the limits set by the high energy LHC experiments to the low energy experiments, a shared renormalization scale must be chosen because the quark-level effective coupling constants are, despite their name, not truly constant but do dependent on the energy scale chosen [20]. A renormalization scale of 2 GeV was used in Reference [20] used when quoting limits on all coupling constants. Equation 2.17 of Reference [20] provides the complete set of linear relationships between the quark-level effective theory coupling constants (ϵ_i for $i \in V, A, S, T$) and the nucleon-level effective theory coupling constants (C_i for $i \in \{V, A, S, T\}$, as seen in Equation 1.17 and Equation 1.18). But for the exotic tensor and scalar couplings we are most interested the relevant relations are,

$$\epsilon_T = \frac{C_T + C'_T}{8C_V g_T} \quad (1.32)$$

$$\epsilon_S = \frac{C_S + C'_S}{2C_V g_S} \quad (1.33)$$

where the difference in the constant factor in the denominator is due to convention. Currently the best constraints on the scalar coupling $\epsilon_S = (1.4 \pm 1.3) \times 10^{-3}$ arises from $0^+ \rightarrow 0^+$ nuclear beta decays [49]. The best experimental bounds on the tensor coupling ϵ_T arises from the study of radiative pion decay $\pi^+ \rightarrow e^+ \nu_e \gamma$. The 90% Confidence Level (CL) experimental

constraint from this study (evolved to 2 GeV as discussed above) is,

$$-1.1 \times 10^{-3} < Re(\epsilon_T) < 1.36 \times 10^{-3} \quad (90\% \text{ CL}). \quad (1.34)$$

We focus the majority of our attention on the tensor currents as our current decays of experimental interest are ${}^6\text{He}$ (Gamow-Teller) and ${}^{19}\text{Ne}$ (mixed) so we are more sensitive to the presence of tensor currents than scalar currents. As noted in Reference [20], though not currently competitive with the best limits, high precision measurements of b_{Fierz} in nuclear β decay could prove to beat the current best low energy limits shown in Equation 1.34. Figure 1.5 shows the effective coupling constant $\epsilon_{S,T}$ limits associated with a 10^{-3} measurement of b_{Fierz} .

The BSM physics probed by low energy precision measurements can also be probed directly at high-energy colliders in a model independent way by assuming that the mediators of the new physics remain point-like at the TeV-scale energies of the Large Hadron Collider (LHC). In Reference [11] bounds for $\epsilon_{S,T}$ are derived by analyzing LHC data in the $pp \rightarrow e\nu + X$ channel at $\sqrt{s} = 7$ TeV and 1 fb^{-1} integrated luminosity. In Reference [22], the analysis has been extended further, including all non-standard couplings with integrated luminosity of 5 fb^{-1} using both $pp \rightarrow e\nu + X$ and $pp \rightarrow e^+e^- + X$ channels.

The $pp \rightarrow e\nu + X$ channel provides direct information about the weak interaction because the parton-level process is $\bar{u}d \rightarrow e\nu$. The transverse mass of the lepton pair is defined as

$$m_T \equiv \sqrt{2E_e^T E_\nu^T (1 - \cos(\Delta\phi_{e\nu}))}, \quad (1.35)$$

where E_e^T and E_ν^T are the transverse momenta of the outgoing electron and anti-neutrino and $\Delta\phi_{e\nu}$ is the angle between the leptons in the transverse plane. At large m_T the SM background goes to zero while the presence of exotic currents would produce events and therefore an excess could be detected or the lack of one can be used to put a limit on the exotic coupling constants.

For massive mediators the propagator is proportional to $1/(q^2 - M^2)$ where q is the




Figure 1.5: Current and projected 90% CL allowed regions for $Re(\epsilon_S)$ and $Re(\epsilon_T)$. The horizontal green band limit is provided by existing bounds on b_F . The red (inner) and blue (outer) bow-tie shaped regions are projected limits from 10^{-3} measurements of b in neutron decay. The vertical orange band is from a (projected) 10^{-3} measurement of b_{GT} in ${}^6\text{He}$. The left panel shows the limits given conservative ranges for g_S and g_T . The right panel shows the constraints when the scalar and tensor charges are taken from recent, more precise, lattice QCD work: $g_s = .8(4)$ and $g_T = 1.05(35)$. The effective couplings $\epsilon_{S,T}$ are defined within a 2 GeV renormalization scheme. This figure is taken from Reference [20].

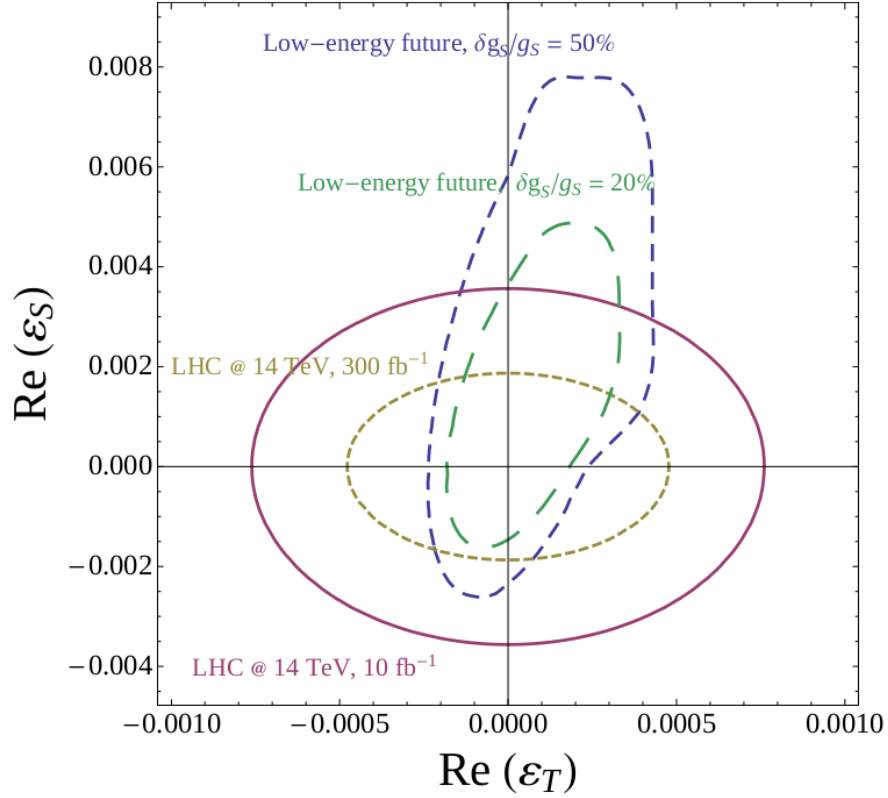


Figure 1.6: Projected 90% CL allowed regions for $Re(\epsilon_S)$ and $Re(\epsilon_T)$ from β decay measurements and the LHC at $\sqrt{s} = 14$ TeV. The low energy constraints correspond to 10^{-3} measurements of b , B in neutron decay and b in ${}^6\text{He}$ under two different scenarios for the lattice QCD uncertainties in $g_{S,T}$. The blue curve is an aggregation of the measurements show in the right side panel of Figure 1.5. See the text for details on how the LHC limits are obtained. The effective couplings $\epsilon_{S,T}$ are defined within a 2 GeV renormalization scheme. This figure is taken from Reference [20].

four-momentum of the propagator and M is the mass of the propagator. In low energy experiments, as was discussed earlier, the propagator is therefore proportional to $1/(M_W^2)$ and at the LHC, because $m_T^2 \sim q^2$ and $\text{TeV} \gg \text{MeV}$, the propagator scales as $1/(m_T^2)$. If we assume that the scale of the propagator associated with the new physics (NP) M_X is much larger than the TeV energy scale of the LHC, then in either the low or high energy case the NP propagator goes as $1/(M_X^2)$. This means that the ratio of the NP propagators to the SM propagators for the two cases is given by,

$$\text{Propagator Ratio (NP/SM): } \beta \text{ decay} \sim \frac{M_W^2}{M_X^2}, \quad \text{LHC} \sim \frac{m_T^2}{M_X^2}. \quad (1.36)$$

The m_T^2 in the numerator of Equation 1.36 explains how the LHC experiments obtain better limits on exotic currents with higher energy collisions. They can also improve their limits by increasing the luminosity of the beam, which increases the statistical power of the measurement in the high m_T region.

The projected constraints on the effective coupling constants $\epsilon_{S,T}$ for the LHC and low energy experiments are compared in Figure 1.6 [65]. This figure illustrates the impact of the uncertainty of the tensor and scalar form factors for nucleons $g_{S,T}$ on the limits set by low-energy experiments. The uncertainty on these form factors will likely decrease further as theoretical calculations improve. The LHC limits shown in Figure 1.6 represent the best possible limits the LHC can set, and low-energy experiments at the 10^{-3} level are still competitive, illustrating the important and complimentary role such low-energy experiments play in the search for new physics in the weak sector.

1.4 Fierz term b_{Fierz} for ${}^6\text{He}$, ${}^{19}\text{Ne}$

The medium-term (*Phase 1*) goal of the He6-CRES experiment is to measure a ratio of ${}^{19}\text{Ne}$ and ${}^6\text{He}$ β decay spectra, as will be discussed in detail in Chapter 3. Many systematics will cancel in the ratio but since ${}^6\text{He}$ is a β^- decay and ${}^{19}\text{Ne}$ is a β^+ decay they have the opposite sign for b_{Fierz} (see Equation 1.30) and so the ratio of spectra is sensitive to a non-zero b_{Fierz} .

However, since ^{19}Ne is a mixed decay and ^6He is a pure Gamow-Teller decay, the form of their respective Fierz interference terms is different.

Here we derive a relationship between $b_{Fierz}(^{19}\text{Ne})$ and $b_{Fierz}(^6\text{He})$ which we will use to derive the limits on exotic current coupling constants set by a ratio of their spectra. We begin with the expression for b_{Fierz} given in Equation 1.27 and Equation 1.30 and we make the SM assumptions that $C_V = C'_V$, $C_A = C'_A$, and that C_V and C_A are real. We also use the relations between the quark-level coupling constants and the nucleon-level coupling constants provided in Equation 1.32 and Equation 1.33 [20]. We also assume the second order tensor and scalar terms in the denominator can safely be ignored. This leaves us with the following form for b_{Fierz} :

$$b_{Fierz} = \frac{\pm 2(|M_F|^2 g_S \epsilon_S C_V^2 + 4|M_{GT}|^2 g_T \epsilon_T C_V C_A)}{|M_F|^2 C_V^2 + |M_{GT}|^2 C_A^2}. \quad (1.37)$$

Now we define two relations to further simplify:

$$\lambda = \frac{C_A}{C_V}. \quad (1.38)$$

$$\rho = \lambda \frac{|M_{GT}|}{|M_F|}, \quad (1.39)$$

Next we divide both the numerator and denominator of Equation 1.37 through by $|M_F|^2 C_V^2$, resulting in the simplified, generic expression [78, 11]:

$$b_{Fierz} = \frac{\pm 2\gamma(Z)(g_S \epsilon_S + 4\rho^2 g_T \epsilon_T / \lambda)}{1 + \rho^2}, \quad (1.40)$$

where we have added the radiative correction $\gamma(Z)$, defined as

$$\gamma(Z) = \sqrt{1 + (\alpha Z)^2}, \quad (1.41)$$

where Z is the proton number of the daughter nucleus and α is the fine structure constant. The longer-term goals of the He6-CRES experiment include measuring b_{Fierz} for ^6He (β^- ,

pure Gamow-Teller), ^{14}O (β^+ , pure Fermi), and ^{19}Ne (β^+ , mixed) individually because there is a simple relationship between their Fierz interference terms. We obtain this simple relation by noticing that in Equation 1.40, $\rho \rightarrow \infty$ for pure Gamow-Teller decays and $\rho = 0$ for pure Fermi decays which implies:

$$\frac{b_{Fierz}(^6\text{He})}{\gamma(Z=3)} = 8\epsilon_T g_T / \lambda \quad (1.42)$$

and

$$\frac{b_{Fierz}(^{14}\text{O})}{\gamma(Z=7)} = -2\epsilon_S g_S. \quad (1.43)$$

Using the above relations and Equation 1.40, we obtain the following relationship between the Fierz interference terms for the three isotopes,

$$\frac{b_{Fierz}(^{19}\text{Ne})}{\gamma(Z=9)} = \frac{\frac{b_{Fierz}(^{14}\text{O})}{\gamma(Z=7)} - \rho(^{19}\text{Ne})^2 \frac{b_{Fierz}(^6\text{He})}{\gamma(Z=3)}}{1 + \rho(^{19}\text{Ne})^2}. \quad (1.44)$$

We use the $\rho(^{19}\text{Ne})$ calculated in Reference [78], $\rho(^{19}\text{Ne}) = 1.5995(45)$. Our collaboration would ultimately like to measure all three of these isotopes, as the relationship of Equation 1.44 provides a powerful systematic check and the potential for strong single-experiment limits on ϵ_T and ϵ_S . However, as was discussed in Section 1.3, there are already limits on ϵ_S at the level of 10^{-3} [49] and the form of Equation 1.44 suppresses the scalar term coming from $b_{Fierz}(^{14}\text{O})$ by $\rho^2 \approx 2.56$. Therefore, for this thesis, we assume the contribution from the $b_{Fierz}(^{14}\text{O})$ term can be neglected which leads us to the following (approximate) relationship:

$$\begin{aligned} b_{Fierz}(^{19}\text{Ne}) &\approx -\frac{\gamma(Z=9)}{\gamma(Z=3)} \frac{\rho(^{19}\text{Ne})^2}{1 + \rho(^{19}\text{Ne})^2} b_{Fierz}(^6\text{He}) \\ b_{Fierz}(^{19}\text{Ne}) &\approx -.7204 b_{Fierz}(^6\text{He}), \end{aligned} \quad (1.45)$$

which holds at the 10^{-3} level. The relation of Equation 1.45 is what is used to relate $b_{Fierz}(^6\text{He})$ and $b_{Fierz}(^{19}\text{Ne})$ in Monte Carlo simulations of the ratio measurement in the

remainder of this thesis and is sufficiently accurate for the first phase of the He6-CRES experiment.

1.5 Summary of theoretical motivation

In summary, precision measurements of low-energy observables probe the existence of beyond the standard model (BSM) interactions at high energy scales through forces mediated by the exchange of virtual particles with large masses. Beta-decay spectroscopy provides high intrinsic sensitivity to chirality-flipping interactions not included in the weak sector of the Standard Model of particle physics [20, 43]. Crucially, distortions to the β spectrum depend linearly on BSM scalar and tensor couplings via a characteristic Fierz interference term. Measurements of the Fierz parameter, b_{Fierz} , at the part-per-thousand level probe new physics at $\mathcal{O}(10 \text{ TeV})$ [11, 33, 19].

The experimental need for precise β -decay and conversion-electron spectroscopy goes beyond searches for scalar and tensor currents: it also yields the shortest path to resolving ongoing questions regarding solar and reactor neutrino spectra [53, 102, 50, 4, 6, 72]. Due to the simplicity of the driving operators, beta spectroscopy can also be used as a powerful tool to explore nuclei far from stability [74] or to help understand stellar nucleosynthesis [25]. High-resolution conversion-electron spectroscopy can also be valuable for identifying shape coexistence or other interesting nuclear-structure phenomena [38]. For a concise overview of precise β -spectroscopy experiments, see Section 1.5 of Reference [59].

As such, substantive improvements in electron spectroscopy have far reaching effects in the search for new physics and compliment the findings of high energy experiments at the LHC for a fraction of the cost. The development and first-ever application of a technology called Cyclotron Radiation Emission Spectroscopy (CRES), towards precise $\mathcal{O}(\text{MeV})$ -endpoint β^\pm decay spectroscopy is the topic of this thesis.

Chapter 2

OVERVIEW OF THE CRES TECHNIQUE

Traditional spectroscopy technology using semiconductor or scintillator detectors [105, 99, 89, 37, 76, 111, 58, 36] that rely on β energy loss in matter have percent-level corrections due to β (back-) scattering and bremsstrahlung losses. Novel technologies such as neutral atom and ion traps have been employed to reduce these effects, some reaching 0.3% precision in parameters of the angular distribution [46, 14, 2, 34]. Alternative technologies such as superconducting spectrometers [71] and new ion traps [83, 104] are being developed and have shown the potential for similar precision capabilities [39, 109, 15].

By contrast, Cyclotron Radiation Emission Spectroscopy (CRES) [75], initially developed by the Project 8 collaboration [5, 30] for the measurement of the absolute neutrino mass scale from β -decay spectroscopy of tritium (β^- endpoint ≈ 18 keV), is a promising avenue for low-background and high-resolution ($\delta E/E \sim 10^{-3}$) beta spectroscopy. CRES also has the unique property of having identical response functions for β^\pm decays, unlike solid state detectors which have very different responses for positrons and electrons. This chapter will provide an overview of the CRES technique.

2.1 Simple description of CRES

CRES determines the kinetic energy (E_k) of a particle with mass m and charge q by measuring the frequency of the cyclotron radiation emitted in an external magnetic field with magnitude B , both of which can be determined to high precision [15].

$$f_c = \frac{|q| B}{2\pi \gamma m} = \frac{|q| B}{2\pi m + E_k/c^2}{}^1. \quad (2.1)$$

¹Note that $\gamma = \frac{E}{mc^2} = 1 + \frac{E_k}{mc^2}$. We use γ as a proxy for the total energy (E) throughout the text.

The cyclotron orbital frequency f_c for relativistic electrons at MeV energy scales in a $\mathcal{O}(1 \text{ T})$ field is $\mathcal{O}(10 \text{ GHz})$ and the radiation can be detected by a radio frequency (RF) receiving antenna. Because the energy is determined via a frequency measurement, which can easily be made with high accuracy with existing hardware, the CRES technique offers a substantial advantage in energy resolution for individual decay electrons. The frequency, and therefore energy, is determined within a millisecond of the electron's appearance, when only a negligible amount of kinetic energy ($\mathcal{O}(1 \text{ eV})$) has been lost to radiation. We utilize gaseous radioactive sources (limited to ${}^6\text{He}$ and ${}^{19}\text{Ne}$ in this work) which are transparent to the emitted microwaves. Gigahertz cyclotron frequencies imply that any electron orbit leading to wall collisions will end within nanoseconds. Because the electron will lose all of its kinetic energy within nanoseconds after wall collisions and CRES signals must persist for hundreds of microseconds to be detected, electrons which collide with the walls will yield no measurable signal or contribution to background. This also implies that external radiation causing reactions on the walls (such as Compton Scattering) don't create background signals.

2.2 Characteristics of CRES electrons

In order to differentiate a CRES electron from the RF noise floor the electron must be observed for several milliseconds. We define the pitch angle θ of the electron to be the angle between the electron momentum p_e and the main magnetic field direction (\hat{z}). Even low energy, $\mathcal{O}(10 \text{ keV})$, electrons with a pitch angle of 89° (would) travel $\mathcal{O}(10^3 \text{ m})$ in 1 ms. This means that without axial trapping, essentially all electrons will collide with the windows of the $\approx 10 \text{ cm}$ long decay cell (which is itself an RF waveguide) before being detected.

In order to achieve the necessary measurement times of several milliseconds, a weak harmonic magnetic trap is produced with a coil of opposite polarity to the main magnetic field. The specific implementation details of the He6-CRES magnetic trap can be found in Section 3.4 and more exact equations of motion of the electron in the trap are discussed in Chapter 5. We begin with a discussion of the dynamics given an ideal harmonic trap. The field strength along the central axis (B_z) can be approximated by,




Figure 2.1: Sketch of axial component of B field versus z for a typical magnetic trap. The axial component of the electron momentum decreases to zero as the magnitude of the field increases.

$$B_z = B_0 \left(1 + \frac{1}{B_0} \frac{d^2 B}{dz^2} \Big|_{z=0} z^2 \right). \quad (2.2)$$

Figure 2.1 shows a sketch of the magnetic field intensity in the axial direction B_z as a function of z for a harmonic magnetic trap. Note that since the field is primarily axial ($\vec{B} \approx B_z \hat{z}$), we use B and B_z interchangeably. The pitch angle $\theta(t)$ is now changing as a function of time as the electron traverses the trap axially. The initial pitch angle is denoted θ_0 .

For adiabatic motion ($f_c \gg f_z$) the magnetic flux through the electrons' cyclotron orbits is conserved [62]:

$$r_c^2 B = \left(\frac{p \sin \theta(t)}{eB} \right)^2 B = \text{constant}, \quad (2.3)$$

with the cyclotron radius given by,

$$r_c = \frac{p \sin \theta(t)}{eB}. \quad (2.4)$$

Equation 2.3 can be simplified to,

$$\frac{\sin^2 \theta(t)}{B} = \text{constant}. \quad (2.5)$$

Due to the axial component of the momentum, electrons travel from their initial position z_0 to a maximal position z_{\max} such that:

$$\frac{1}{B(z_{\max})} = \frac{\sin^2 \theta_0}{B(z_0)}. \quad (2.6)$$

The electrons then have an oscillatory axial motion between $\pm z_{\max}$. In the non-relativistic limit there is an axial force on the electrons given by $\mu \nabla B_z$ (where $\mu \propto r_c^2$ is the orbital magnetic moment of the cyclotron orbit). We can therefore use Equation 2.2 to get the axial oscillation frequency f_z in relation to the cyclotron frequency f_c :

$$\frac{f_z}{f_c} = \sqrt{\frac{r^2 \frac{d^2 B}{dz^2}}{B_0}}. \quad (2.7)$$

For a reasonable harmonic trap gradient of 0.01 T/cm² and $B_0 = 1$ T, the frequency f_z of oscillation is about 100 MHz.

A trap with a finite magnetic depth ΔB cannot confine all the electrons born within it. In order to be trapped, electrons' pitch angles must satisfy the relation:

$$\tan \theta' > \sqrt{\frac{B(z=0)}{\Delta B}} \quad (2.8)$$

where θ' is the pitch angle at the bottom of the trap. For a typical $\Delta B/B_0 \approx 10^{-3}$ center trap field, the fraction of trapped electrons is $\mathcal{O}(1 \%)$. Crucially, this trapping efficiency depends only on the magnetic trap characteristics and is independent of beta energy. For a radioactive source of uniform density in the region of the trap, there will be many more

electrons with low pitch angles because there is more phase space available to them down to the cut-off angle θ_{min} given by,

$$\theta_{min} = \tan^{-1} \left(\sqrt{\frac{B(z=0)}{\Delta B}} \right). \quad (2.9)$$

There are relatively few electrons with $\theta' \approx 90^\circ$ because these electrons can only be born right at the center of the trap ($z = 0$) with initial pitch angle $\theta_0 \approx 90^\circ$.

The magnetic trap also generates grad-B motion induced by variations of the holding field amplitude along the cyclotron trajectory of the electron (similar to the magnetron drift of a penning trap and so often referred to as magnetron motion) [12]. This effect induces a global rotation of the electron trajectory around the axis of the trap magnetic field at a frequency $f_{\nabla B}$. The velocity induced by this effect is perpendicular to both the main field and the field gradient ($\vec{v}_{\nabla B} \propto \vec{B} \times \nabla B$). There is also a different drift that is caused by the curvature of the field lines but the frequency associated with this drift is roughly an order of magnitude smaller than $f_{\nabla B}$ and is therefore not easily visible in our observation window [23]. For the He6-CRES experiment the ratios of cyclotron, axial, and grad-B frequencies are approximately given by Equation 2.17 of Reference [13]:

$$f_c : f_z : f_{\nabla B} \sim 1 : 5 \times 10^{-3} : 10^{-5}. \quad (2.10)$$

Figure 2.2 illustrates how the three motions interact to give rise to the electron trajectory in the magnetic trap.

The axial motion leads to a frequency modulation or Doppler effect that puts power into sidebands. Since the main magnetic field of Equation 2.1 is in principle changing as a function of time as the CRES electron traverses the trap it is the average magnetic field over one axial oscillation that determines the observed cyclotron frequency. These two effects, among others, will be discussed in detail in Chapter 5.




Figure 2.2: Schematic of the three most important motions of a CRES electron in a magnetic trap; cyclotron (f_c), axial (f_z), and grad-B ($f_{\nabla B}$). The main field axis \hat{z} is perpendicular to the plane of the blue and red circles. Note that we use magnetron and grad-B synonymously in this work. This figure is taken from Reference [112].

2.3 Characteristics of CRES signals

A CRES event radiating in free space would be incredibly difficult to observe due to the small ($\mathcal{O}(\text{fW})$) amount of power radiated and the small solid angle of a free space antenna. To overcome this, we use a waveguide as the decay cell such that the CRES signal couples to the modes of the RF waveguide and the radiation propagates exclusively in $\pm\hat{z}$ which can be picked up by an antenna on either end of the waveguide.

In a waveguide, the time dependence of the wave amplitude for a mode is given by [62, 88, 15]:

$$A_{n,m}(t) \propto \int d^3x \vec{J}(\vec{r}, t) \cdot \vec{E}_{n,m}(\vec{r}), \quad (2.11)$$

where $\vec{J}(\vec{r}, t)$ is the current density of the moving β , and $\vec{E}_{n,m}$ represents the mode electric field. The power coupled into a given waveguide mode is $\propto |A_{n,m}(t)|^2$.

The current experimental bandwidth of the He6-CRES experiment (18–19.1 GHz) was selected to measure only in the $TE_{1,1}$ mode, with cutoff at 15.2 GHz. The next mode, $TM_{0,1}$, has a cutoff at 19.9 GHz. Thus the signal power observed directly is exclusively due to the power propagating in the $TE_{1,1}$. Since the electron is losing energy to radiation, the cyclotron frequency is evolving in time. The *slope* (df_c/dt) is given by

$$\frac{df_c}{dt} = \frac{eBc^2}{2\pi} \frac{P}{(mc^2 + E_k)^2} = \frac{1}{2\pi eBc^2} P f_c^2, \quad (2.12)$$

where P is the power lost to radiation. The slope observable is therefore a proxy for the β 's total radiated power. In addition to the $TE_{1,1}$ mode, higher harmonic frequencies ($f_h = hf_c$) of the β current density couple to higher-order waveguide modes, which can contribute significantly to the total radiated power and thus to the observed slope. However, since this power propagates into unobserved modes, we do not detect this power directly.

2.4 Characteristics of CRES spectrogram

Our raw data is, in its base form, a time series voltage wave but we continuously take Fast Fourier Transforms (FFTs) of the digitized signal in hardware (as discussed in detail in Section 3.5). Arranged sequentially, the FFTs produce a spectrogram displaying signal power as a function of time and frequency. Figure 2.3 shows CRES data taken with the ^{83m}Kr source, which produces monoenergetic lines between 7 - 32 keV and is discussed in more detail in Chapter 6. Consecutive and colinear high-power Fourier bins form *tracks*, which are indicative of a radiating charged particle. As explained above, the gradual increase of frequency versus time is due to the charged particle losing kinetic energy to the cyclotron radiation and the slope of the track is given by Equation 2.12.

The sudden jumps are due to scattering off residual gas atoms. Scattering generally causes the electron to lose energy, increasing the frequency. However, since a scatter will often also change the pitch angle θ of the electron, it can change the average field B_{avg} experienced, potentially leading to a decrease in observed frequency as seen in the first

scatter of the event shown in Figure 2.3. After an arbitrary number of scatters, a trapped electron will end up with a pitch angle less than θ_{min} (Equation 3.6) and no longer be trapped. The set of all tracks associated with a single β is referred to as an *event*. In order to precisely determine β energies the event start frequency and the magnetic field must be well known and input into Equation 2.1. To extract the event start frequency the event must be reconstructed accurately; for example the event in Figure 2.3 needs to be recognized as one event and not mistakenly broken into two or three. Event reconstruction will be discussed in Chapter 6. Event start frequencies can generally be identified to within ≈ 20 MHz, so for cyclotron frequencies of ≈ 20 GHz, this represents 10^{-3} energy resolution² given the reasonable assumption that the magnetic field within the decay region is known to 10^{-4} relative uncertainty or better. Notice that the full scale of the spectrogram in Figure 2.3 is ~ 20 MHz, so a start frequency determination to within 20 MHz is (typically) very conservative. Our current and projected precision on the main magnetic field is discussed in Section 3.6 and Section 6.3.

²Note that this is technically the uncertainty in the total relativistic energy (E). The uncertainty in the kinetic energy (K) is effected by an additional factor of $\frac{E}{K} = \gamma/(\gamma - 1)$ which is an $\mathcal{O}(10)$ factor for low energies but is an $\mathcal{O}(1)$ factor for the MeV β -energies we are most concerned with.

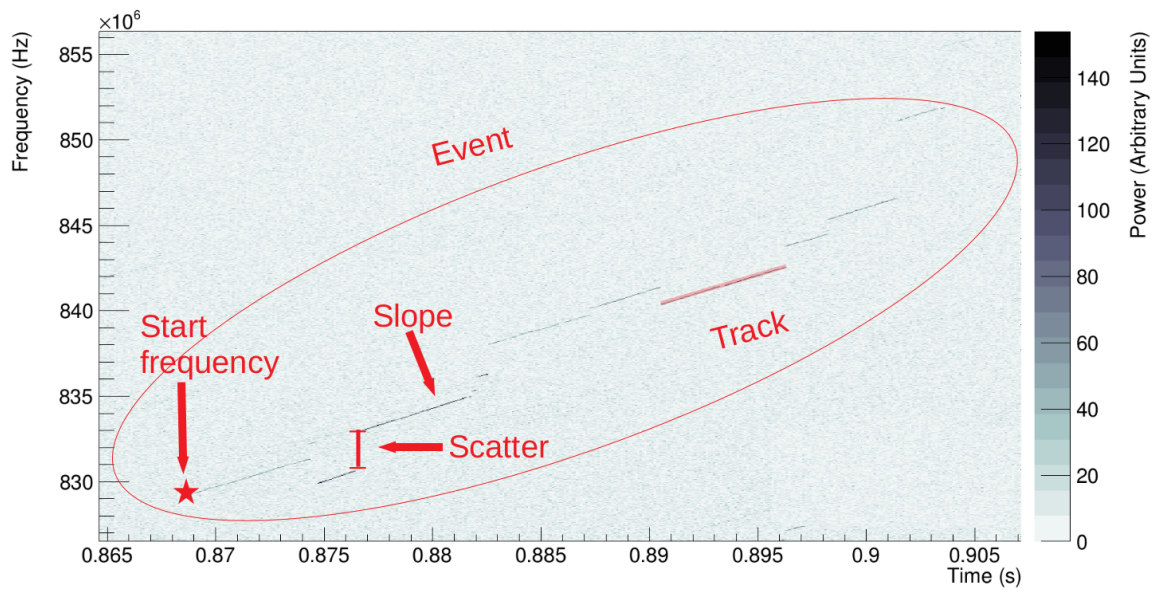


Figure 2.3: Visual illustration of the concepts of track, scatter, start frequency, and event. See the text for explanations. CRES data taken with $^{83\text{m}}\text{Kr}$ source. Figure created by Heather Harrington of He6-CRES.

Chapter 3

OVERVIEW OF THE He6-CRES EXPERIMENT

The He6-CRES collaboration is pursuing $\mathcal{O}(\text{MeV})$ endpoint β spectra measurements via CRES in search of signatures of exotic chirality-flipping interactions. The experiment is being conducted at the Center for Experimental Nuclear Physics and Astrophysics (CENPA) at the University of Washington (UW). This chapter will provide a high level overview of the He6-CRES experiment in Section 3.1, and then discuss the different components of the experiment in detail in Sections 3.2 through 3.9.

3.1 The He6-CRES experiment in a nutshell

Now that we have a high level understanding of the CRES technique (see Chapter 2), we will articulate the experimental goals of *Phase 1* of the He6-CRES experiment and discuss the different experimental components needed to meet these goals. The discrete parts of the experiment will be explored in detail in the remaining Sections of this chapter.

3.1.1 Experimental goals of Phase 1 of the He6-CRES experiment

The overarching goal of *Phase 1* of He6-CRES is to demonstrate that CRES can plausibly be used to achieve a 10^{-3} measurement of b_{Fierz} via the ratio of ${}^6\text{He}$ and ${}^{19}\text{Ne}$ spectra and to discover what experimental design, analysis, or simulation bottlenecks may hinder such a measurement as to direct the efforts of *Phase 2* of the experiment. Specifically, *Phase 1* entails completing the following sub-goals:

Hardware Development:

- Developing the vacuum infrastructure to deliver ${}^6\text{He}$ and ${}^{19}\text{Ne}$ to the He6-CRES apparatus. (Section 3.2)
- Developing high-intensity ${}^6\text{He}$ and ${}^{19}\text{Ne}$ sources. (Section 3.3)
- Building and characterizing the *Phase 1* He6-CRES apparatus. (Section 3.4)
- Developing a high-throughput, broadband RF data acquisition system (DAQ). (Section 3.5)
- Characterizing and precisely manipulating our main magnetic field. (Section 3.6)
- Developing a stable β monitor for spectral normalization. (Section 3.7)

Software Development:

- Developing DAQ software for rapidly writing data to disk, tracking important environmental data (main magnetic field, β monitor, trap current,...), and starting and managing data acquisitions. (Section 3.8)
- Developing adequate Monte Carlo simulation capabilities as to meaningfully compare expectation to observation. (Chapter 5)
- Developing an analysis toolkit to identify CRES events over a large range of energies and event topologies. (Chapter 6)

Experimental Benchmarks:

- Successfully observing CRES signals from ${}^{83\text{m}}\text{Kr}$. (Chapter 4)
- First-ever observations of $\mathcal{O}(\text{MeV})$ β^\pm 's via CRES with ${}^6\text{He}$ and ${}^{19}\text{Ne}$. (Chapter 4)
- Developing and executing a campaign for an $\mathcal{O}(10^4)$ total counts per isotope ratio measurement of ${}^6\text{He}$ and ${}^{19}\text{Ne}$ allowing for a proof-of-principle measurement of b_{Fierz} with $\sim 10^{-1}$ statistical sensitivity. (Chapter 6)

This thesis describes much of the work done to complete *Phase 1* of the He6-CRES experiment.

3.1.2 Phase 1 of the He6-CRES experiment

Here we will step through how the distinct pieces of the He6-CRES experiment fit together with the final *Phase 1* experimental benchmark in mind: an $\mathcal{O}(10^4)$ counts per isotope ratio measurement of ${}^6\text{He}$ and ${}^{19}\text{Ne}$.

In this thesis we demonstrate a frequency-based energy determination of β^\pm 's in the 5 keV to 2.1 MeV (kinetic energy) range, with the capability to reach 5 MeV with the current He6-CRES apparatus [15]. This corresponds to ≈ 100 -times higher energies than previously observed with CRES [91, 30]. The isotopes ${}^6\text{He}$ (β^- endpoint ≈ 3508 keV, $\tau \approx 1$ s) and ${}^{19}\text{Ne}$ (β^+ endpoint ≈ 2216 keV, $\tau \approx 17$ s) are experimentally favorable for precise determinations of b_{Fierz} because they both decay almost exclusively to the ground state of their progeny, have considerably simplified nuclear structure corrections due to an underlying isospin symmetry, and have precise theoretical characterization of their decay properties [79, 69, 24, 77, 66, 41, 15].

A measurement of the ratio of β energy spectra of these isotopes enhances the sensitivity to tensor couplings since the sign of b_{Fierz} is opposite in the two cases [61]. As was discussed in Section 1.4), not only are the b_{Fierz} terms opposite in sign for ${}^6\text{He}$ and ${}^{19}\text{Ne}$, they are also related via an expression that holds at the 10^{-3} level: $b_{Fierz}({}^{19}\text{Ne}) \approx -0.7204 b_{Fierz}({}^6\text{He})$ (see Equation 1.45). This implies that the ratio of spectra can be fit to a single free parameter and simultaneous limits can be set on $b_{Fierz}({}^6\text{He})$ and $b_{Fierz}({}^{19}\text{Ne})$. Crucially, systematic effects arising from common spectral distortions are (mostly) eliminated because the ${}^6\text{He}$ and ${}^{19}\text{Ne}$ experiments are charge-conjugated copies of each other, and, for a given β energy, the trajectories and resulting radiation are identical. The fact that detection efficiencies and line shapes for β^+ and β^- CRES signals are (in principle) identical is unique to CRES. In solid state detectors positrons and electrons have very different signal characteristics due to annihilation-gammas in the case of positrons and other effects, preventing such a cancellation from working. The limits of this powerful cancellation of systematics will be discussed in Chapter 5.




Figure 3.1: Standard model ^{19}Ne (blue) and ^6He (orange) β spectra showing the energy regions sampled by the 11 field settings from 0.75–3.25 T (the highest fields correspond to the highest energies), given the 1.1-GHz RF bandwidth.

From Equation 2.1 one can see that for a given magnetic field strength (B) the energy range of observable β 's is determined from the frequency bandwidth of the apparatus. In the absence of significant signal-processing hardware breakthroughs allowing $\mathcal{O}(30\text{ GHz})$ bandwidths, a limited bandwidth prevents the simultaneous observation of β 's over the entire $\mathcal{O}(\text{MeV})$ -endpoint energy spectrum. To overcome this, the magnetic field must be tuned to observe different regions of the β spectrum.

Figure 3.1 illustrates how the overlapping region of the ^{19}Ne and ^6He β spectra are covered by scanning the magnetic field, given our fixed frequency bandwidth of 18.0 – 19.1 GHz. One can see that due to the B dependence of Equation 2.1, the observation window becomes larger with increasing field. Measurements taken at each field are normalized by a β monitor of the total activity within the decay cell. Labeled “beta monitor” (orange) in the He6-CRES apparatus schematic of Figure 3.2, a 3.8-cm-diameter port with a 1/8-mm thick Cu foil allows β s to be observed by a 5-mm-thick scintillator coupled to 4 SiPM readouts [15].




Figure 3.2: Schematic of the core RF components of the He6-CRES apparatus. The apparatus is presented in detail in Section 3.4.

The density of radioactive atoms in the vacuum side of this port is in dynamical equilibrium with that in the decay cell.

In the simplest approach for constructing the ratio (used throughout this work) the total event count from each observation window for each isotope is divided by the total number of counts registered on the β monitor during the associated CRES acquisition, resulting in a normalized rate for each isotope at each field. Let n_{fields} be the number of discrete fields included in the measurement. Then the ratio is constructed by taking the normalized ^{19}Ne rate over the normalized ^6He rate (field-wise). We situate ^{19}Ne in the numerator to prevent a divergence since ^6He has a larger endpoint. Note that by aggregating all counts for a given field we are sacrificing frequency (and therefore energy) resolution of individual β s¹.

¹We start with the simplest approach and build from it; other approaches that don't sacrifice as much energy resolution are discussed in Chapter 5).

This experimental ratio can then be compared to the SM expectation² which will simply be the field-wise ratio of the integrated spectral density regions presented in Figure 3.1. However, since the efficiency of the β monitor is generically different for positrons (^{19}Ne) and electrons (^6He) due primarily to the presence of 511 keV annihilation photons for β^+ decays, there is an additional multiplicative coefficient (C) needed to compare the expected SM ratio (r_{SM}) to the experimental ratio (r_{exp}). This factor C , along with b_{Fierz} , can be obtained simultaneously via a χ^2 -minimization,

$$\chi^2(b_{Fierz}, C) = \sum_{i=1}^{n_{fields}} \frac{(r_{SM,i}(b_{Fierz}) - Cr_{exp,i})^2}{(C\sigma_{exp,i})^2}. \quad (3.1)$$

Where $r_{exp,i}$ is the experimental ratio at the i^{th} field and is explicitly given by,

$$r_{exp,i} = \frac{N_{ne,cres,i} N_{he,mon,i}}{N_{ne,mon,i} N_{he,cres,i}}, \quad (3.2)$$

where $N_{ne,cres,i}$ ($N_{he,cres,i}$) is the total number of CRES events observed at the i^{th} field for ^{19}Ne (^6He) and $N_{ne,mon,i}$ ($N_{he,mon,i}$) is the total number of monitor counts observed at the i^{th} field for ^{19}Ne (^6He). The uncertainty in the experimental ratio $\sigma_{exp,i}$ is given by,

$$\sigma_{exp,i} = r_{exp,i} \left[\frac{1}{N_{ne,cres,i}} + \frac{1}{N_{ne,mon,i}} + \frac{1}{N_{he,cres,i}} + \frac{1}{N_{he,mon,i}} \right]^{1/2}, \quad (3.3)$$

obtained by propagating a statistical \sqrt{N} uncertainty in the CRES and monitor counts observed for each isotope and field through to the ratio. The best fit values b_{Fierz}^* and C^* for b_{Fierz} and C respectively are given by,

$$b_{Fierz}^*, C^* = \arg \min_{b_{Fierz}, C} \chi^2(b_{Fierz}, C). \quad (3.4)$$

The predicted ratio at the i^{th} field, $r_{SM,i}(b_{Fierz})$, is given by,

²Technically it is incorrect or at least a bit loose to call this the SM expectation since it includes b_{Fierz} (a BSM effect) as a free parameter.

$$\begin{aligned}
r_{SM,i}(b_{Fierz}) &= \frac{\int_{E_{min,i}}^{E_{max,i}} W(E)_{19\text{Ne},SM} (1 + b_{Fierz}({}^{19}\text{Ne}) \frac{m_e}{E}) \epsilon_{e^+}(E) dE}{\int_{E_{min,i}}^{E_{max,i}} W(E)_{6\text{He},SM} (1 + b_{Fierz}({}^6\text{He}) \frac{m_e}{E}) \epsilon_{e^-}(E) dE} \\
&\approx \frac{\int_{E_{min,i}}^{E_{max,i}} W(E)_{19\text{Ne},SM} (1 - 0.7204 \cdot b_{Fierz} \frac{m_e}{E}) dE}{\int_{E_{min,i}}^{E_{max,i}} W(E)_{6\text{He},SM} (1 + b_{Fierz} \frac{m_e}{E}) dE}, \tag{3.5}
\end{aligned}$$

where $\epsilon_{e^{+(-)}}(E)$ is the energy dependent detection efficiency of CRES detection for positrons (electrons) and $W(E)_{SM}$ is the well-known SM differential decay rate with all relevant higher order corrections included (see Subsection 1.2.3 and Equation 1.26). The lower (upper) energy bound of the integral, $E_{min,i}$ ($E_{max,i}$), is obtained by plugging the upper (lower) bound of the frequency bandwidth into Equation 2.1. The efficiency terms pull out of the integral and cancel to first order because the efficiencies are almost constant over the integration windows. Limits to this cancellation and different formulations of the ratio to improve the cancellation are discussed in Chapter 5. Note that in principle there should also be a convolution of the line shape with the beta spectrum included in Equation 3.5 but that since the line width is so small ($\delta E/E \sim 10^{-3}$), we are safe to ignore this effect for our current desired level of precision. Future phases of the experiment will need to address this effect carefully. For simplicity throughout the remainder of this thesis b_{Fierz} is used interchangeably with $b_{Fierz}({}^6\text{He})$ (as is done in the last line of Equation 3.5). One can always use Equation 1.45 to relate a limit on $b_{Fierz}({}^6\text{He})$ to a limit on $b_{Fierz}({}^{19}\text{Ne})$.

Equation 3.5 demonstrates the power of the ratio measurement; it enables the cancellation of energy dependent efficiencies, drastically reducing reliance on complex Monte Carlo to interpret experimental data, and an increased sensitivity to b_{Fierz} by a factor of ≈ 1.7 when compared to a determination of b_{Fierz} from ${}^6\text{He}$ alone.

We are not sensitive to b_{Fierz} via a measurement of the ratio as described here at a single field, as the free parameter C prevents any sensitivity of such a measurement and experimental sensitivity to the Fierz interference is optimized by observing betas over a broad range of kinetic energies, spanning from zero to several times the electron mass [44, 15]. We

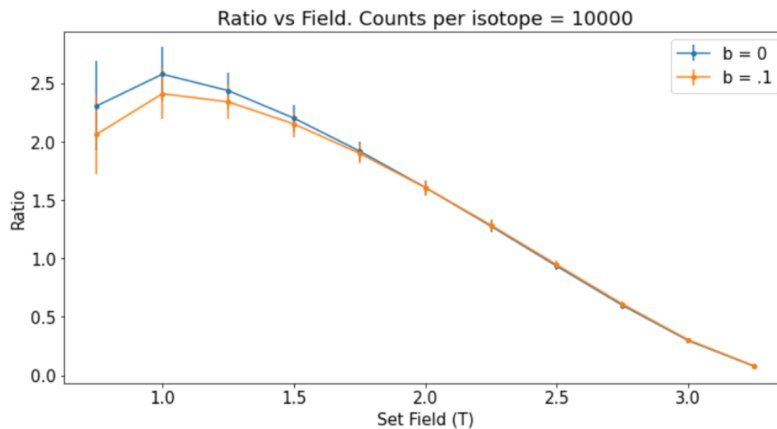


Figure 3.3: Ratio of spectra (Equation 3.5) with $b_{Fierz} = 0$ vs $b_{Fierz} = 0.1$ with 10^4 counts per isotope measured over 11 field settings from 0.75–3.25 T in 0.25 T increments as is illustrated in Figure 3.1.

rely on measurements over a large range in fields (and thus energies). The long lever arm enables sensitivity to b_{Fierz} . Figure 3.3 illustrates how the ratio between an (exaggerated) $b_{Fierz} = .1$ and $b_{Fierz} = 0$ compare with 10^4 counts per isotope measured over 11 field settings from 0.75–3.25 T in 0.25 T increments as is illustrated in Figure 3.1. One can see that the long lever arm enables the fixing of the scaling factor C while still being sensitive to the deviations in the shape of the curve caused by b_{Fierz} . The statistical sensitivity to b_{Fierz} goes (approximately) as,

$$\Delta b_{Fierz} \sim \frac{10}{\sqrt{N}}, \quad (3.6)$$

where N is the number of counts per isotope. Chapter 5 explicitly demonstrates the ratio technique’s sensitivity to b_{Fierz} under different assumptions.

The above provides a description of the measurement technique. Figure 3.4 provides a layout of the rooms at CENPA that are utilized by the experiment. The tandem provides high energy proton and deuteron beams to Cave 2 where the ^{19}Ne and ^6He sources are located. Radioactive isotopes (^{19}Ne or ^6He) are transported to the He6-CRES apparatus in




Figure 3.4: Layout of the Tandem (where proton and deuteron beams are produced), Cave 1 (where He6-CRES apparatus is located), and Cave 2 (where ^{19}Ne and ^6He source are located) within CENPA.

Cave 1 to be detected.

3.1.3 Thesis scope

Now that there is adequate context (via the theoretical motivation described in Chapter 1, a high level description of the ratio measurement, and an articulation of the He6-CRES *Phase 1* goals), we can meaningfully present the scope of the remaining sections of the thesis. The coming sections of this chapter will describe the discrete components of the experiment in detail; Sections 3.2 and 3.3 describe the vacuum system connecting the radioactive sources to the He6-CRES apparatus and the high intensity ^{19}Ne and ^6He sources we've developed at CENPA (publication of this material is in progress), Section 3.4 describes the He6-CRES apparatus, Section 3.5 describes the DAQ, Section 3.6 describes the superconducting solenoid that provides the main magnetic field, Section 3.7 describes the β monitor that provides the relative normalization across fields, Section 3.8 describes the DAQ user interface and SQL database used to manage data files and environmental data, and finally Section 3.9 presents the RF noise characteristics of the *Phase 1* apparatus.

In Chapter 5, the Monte Carlo simulations used to compare expectation to observation and to investigate systematic effects will be discussed. In Chapter 4 initial findings from CRES observations of $^{83\text{m}}\text{Kr}$, ^{19}Ne , and ^6He will be discussed. In Chapter 6, the analysis pipeline used to extract event characteristics from spectrograms will be discussed and a $\mathcal{O}(10^4)$ counts per isotope ratio measurement will be presented, representing the completion of the *Phase 1* experimental benchmarks outlined above. Chapter 7 will present progress towards machine learning based event reconstruction. Finally, Chapter 8 will conclude the thesis and discuss the experimental direction of He6-CRES as we work towards a 10^{-3} measurement of b_{Fierz} .

3.2 Transport, vacuum, and monitoring systems

The ratio technique described above requires high intensity gaseous ^{19}Ne and ^6He sources of very pure radioactivity. Typically, CRES β 's must be observed for many milliseconds

to be detected. As such, pressures below the 10^{-5} Torr level are required for minimizing scattering with residual stable gases to keep track lengths as high as achievable. In addition, undetected radioactive contaminants can lead to unwanted distortions that won't cancel in the ratio. The radioactive sources must also be high intensity as high statistics are crucial to obtaining world leading determinations of b_{Fierz} . Below we present a system to deliver $\sim 10^{10}$ atoms per second of ${}^6\text{He}$ and ${}^{19}\text{Ne}$ with high degree of purity to be used for precision β decay measurements.

In Section 3.3 we describe the experimental setup used for the production and purification of the ${}^{19}\text{Ne}$ and ${}^6\text{He}$ sources but we begin by describing the transport, vacuum, and monitoring systems used.

Figure 3.5 shows a sketch of the most significant parts of the system. The radioactive atoms were injected into a 15-cm-diameter 8-m-long stainless steel pipe via remotely-controlled valves. At the other end a 360 l/s turbo pump (TMP1) was used to compress the activity toward the CRES decay cell. Since all emanation from the pipes contribute detrimentally to the pressure at the CRES decay cell, a cryo pump (Cryo-I) was used to keep the transport pipe clean after baking. When radioactivity was being transported, this cryo pump was isolated via a gate valve. The average transit times are ~ 0.8 s for ${}^6\text{He}$ atoms (about one ${}^6\text{He}$ half-life) and ~ 1.5 s for ${}^{19}\text{Ne}$ atoms (about a tenth of a ${}^{19}\text{Ne}$ half-life.)

We use two beta monitors, a SiPM-based monitor referred to simply as the *beta monitor* and a $\Delta E - E$ monitor referred to as the *beta telescope*, which both connect to the rest of the system as shown in Figure 3.5. These beta monitors serve the dual purpose of enabling the characterization of the radioactive intensity and purity of the sources as well as being used for spectral normalization during ratio measurements. The two beta monitors are discussed in Section 3.7.

Figure 3.6 shows more details on the Cryo-II station, designed to improve purity of radioactivity delivered to the CRES decay cell.




Figure 3.5: Vacuum and transport systems: ^6He and ^{19}Ne are injected in the target room and transported to a low-background-radiation room via a 15-cm-diameter SS pipe. The pipe is connected to a cryo pump (Cryo-I) with a valve that is closed during transport. At the end of the pipe, a 360 l/s turbo pump (TMP1) compresses the gas into a volume with a copper window to monitor beta activity and a toroidal volume kept at 30 K using a second cryo pump (Cryo-II) and a heater with feedback for temperature stabilization. This allows purifying the activity while preventing adsorption of ^{19}Ne or ^6He . Two getter pumps are used to further improve vacuum. A residual gas analyzer is used to monitor gas contamination. The bottom right connections allow calibrations with a $^{83\text{m}}\text{Kr}$ source.




Figure 3.6: Cryo-II trap profile. The gas containing the radioactivity of interest was passed through the rectangular toroid (spool) shown, regulated at 30 K, to improve purity, while avoiding retention of ^{19}Ne .

3.3 Production, intensity, and purity of radioactive sources

The radioactivity is produced by beams from the FN-tandem accelerator at the University of Washington: via $^7\text{Li}(d, ^3\text{He})^6\text{He}$ bombarding a molten lithium target with a 17.8-MeV deuteron beam [70] and via $^{19}\text{F}(p, n)^{19}\text{Ne}$ on an SF_6 gaseous target with a 12-MeV proton beam, similar to the Berkeley-Princeton source [26, 40, 16, 100].

Following are details on the functioning of each source, the measured intensities, and the limits on potential contaminants.

3.3.1 ^6He source

Details of the apparatus for ^6He production have been published before [70, 117], so here we give only a brief description. ^6He atoms are produced by bombarding a molten lithium target with a deuterium beam, as shown in Figure 3.7. The lithium is held by a stainless steel cup and a $7.5\ \mu\text{m}$ thick tantalum foil, pictured in Figure 3.8. Tantalum was chosen for




Figure 3.7: Sketch of the liquid lithium target for ${}^6\text{He}$ production. The 17.6 MeV deuteron beam passes through a thin tantalum window (in blue) which holds the molten lithium inside a stainless steel cup. The lithium is melted by heating the cup through a heater in a copper block (in red) attached to the back. By moving up and down an asymmetric stirrer (in green), a thin layer of Li that wets the Ta foil is produced, yielding optimal production rates.

its high melting point, high thermal conductivity, low activation under the deuteron beam exposure, slow chemical reaction with lithium, and robustness allowing for a thin entrance window. The top of the cup is open to a 6-inch diameter and 8-meter long beam pipe in which the ${}^6\text{He}$ atoms travel via molecular diffusion through a concrete shielding wall towards a 360 l/s turbo-molecular pump, which compresses them into the CRES apparatus.

To minimize the diffusion time of ${}^6\text{He}$ within the target, the cup is heated to 220°C by a cartridge heater to melt the lithium target. The amount of lithium loaded in the cup is such that the liquid level is approximately centered vertically on the window. Ideally, the lithium target should have a minimal thickness ~ 1 mm to maximize the nuclear reaction yield. A paddle (or stirrer) mounted on a translation device that can be lowered into the lithium and rotated, temporarily raises the Li level over the window and allows the formation of the thin layer of Li coating the Ta window. The half-cylinder shape of the paddle enables it to stir the lithium by rotation. A feedback loop controls the temperature of the target by turning the cartridge heater on and off.




Figure 3.8: Photo of the tantalum window from the center of the beamline. The copper block that contains the heating element is seen behind the stainless steel cup. Air cooling lines can be seen on either side of the cup.

With $10\ \mu\text{A}$ of deuterium beam impinging on the target we obtained production of $\sim 10^{10}$ ${}^6\text{He}$ atoms per second, measured at the beta monitor station.

A possible contaminant is ${}^8\text{Li}$, produced via the ${}^7\text{Li}(d,p)$ reaction with a large cross section. Fortunately, due to the high retention of Li within the Li target and the high sticking probability in the bounces in the walls of the transport pipe, no detectable amount of ${}^8\text{Li}$ reached the CRES decay cell. The much higher endpoint of ${}^8\text{Li}$ (16 MeV) with respect to ${}^6\text{He}$ (3.5 MeV), allowed to place upper limits on its presence near the CRES decay cell by using a thick scintillator, located at the beta monitor station, indicated in Figure 3.5. We calculated an upper limit of 2×10^{-4} (68% CL) on the fraction of ${}^8\text{Li}$ to ${}^6\text{He}$.

3.3.2 ${}^{19}\text{Ne}$ source

The ${}^{19}\text{Ne}$ source is based on the scheme of previous sources [26, 40, 16, 100]. The gas flow is illustrated in Figure 3.9 and details of the target and SF_6 traps are shown in Figures. 3.10 and 3.11 respectively.

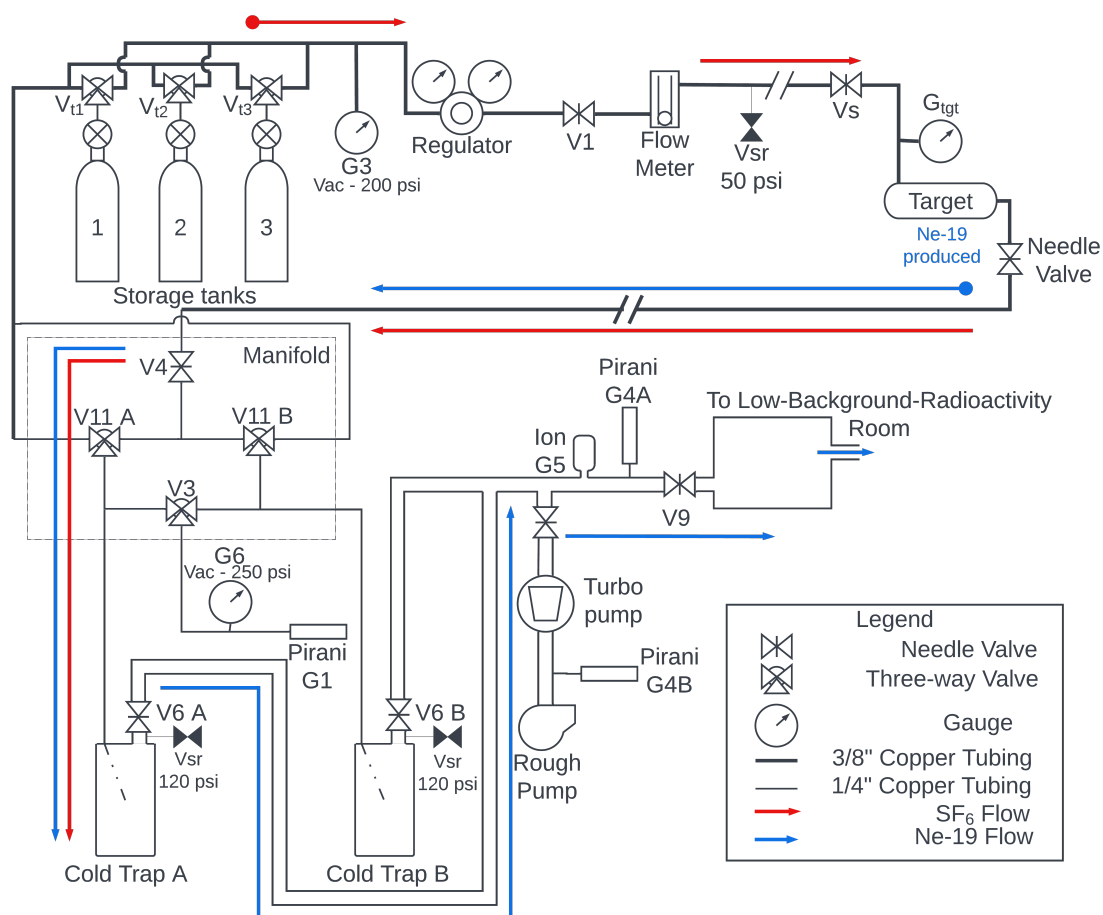


Figure 3.9: The SF_6 flow system for ^{19}Ne production. The colored arrows show direction of flow of SF_6 and ^{19}Ne during the production when Cold Trap A is in use. The flow of SF_6 occurs in the laminar regime, while the ^{19}Ne flow beyond the cold trap is in the molecular regime. Valves enable return of SF_6 to the storage tanks by warming the cold traps. There are two cold traps and the manifold enables using one for ^{19}Ne production while warming the other to return to one of the supply tanks.




Figure 3.10: ^{19}Ne production target. The blue parts are aluminum and the gray ones stainless steel. The beam enters from the left through the Al window (not shown, but sandwiched between the gray and blue pieces on the left) sealed by o-rings on each side. Water held at $\sim 10^\circ\text{C}$ with a refrigerator is circulated through a jacket to keep the target cool during production.




Figure 3.11: SF_6 cold trap. The gas enters through an injection port on the top left via a $1/4''$ tube into a cold trap consisting of a stainless steel chamber with copper baffles soldered to its internal top and bottom walls. When cold, the chamber is kept submerged under liquid nitrogen covering the top wall of the chamber. The injection tube end height was chosen to let SF_6 freeze up to that level while optimizing SF_6 trapping. ^{19}Ne exits through the central tube into the vacuum system.

SF₆ is stored in three 50-l gas bottles at up to 105 psig. At that pressure, a bottle holds about 16 moles of gas. The target pressure is typically set between 20 and 30 psig, so the available gas in each storage tank is up to 11 moles. The gas flows through a pressure regulator to the target cell where bombardment with the 12-MeV proton beam produces ¹⁹Ne. The FN tandem accelerator can produce up to 10 μ A of 12-MeV protons, but to date (March 2023) we have run with up to 3 μ A. The target cell is a 20-cm long cylinder with a water cooling jacket. It is almost twice the range of 12-MeV protons (11.5 MeV after passing through the 0.05-mm Al foil window) in 30-psig SF₆ at room temperature. Although the target is water cooled, the central region of the gas is heated substantially by a 10- μ A beam. This heating has the effect of extending the proton range and was the reason for the conservative length of the target cell.

The Al target window has an unsupported diameter of 6 mm, and is sealed with viton o-rings to the target body. It has been tested to rupture at 11.6 bar and is calculated to be heated to between 70 and 100 K (depending on beam radius) above ambient temperature by a 10- μ A beam. Al was chosen because it is resistant to corrosion by F. Nevertheless, the gas supply line was interlocked to terminate gas flow in case the vacuum fails. The beamline was also interlocked with a fast-acting valve to protect the Tandem in that event. It is possible that free F, produced by ionization of SF₆ could eventually cause the window to break.

The gas flow rate is controlled with the needle valve and monitored with the flow meter. Typical flow rates at 30 psig are up to 2.4 mole/hr (5.4 cc/s). At that rate and pressure, the target volume (57 cm³) is exchanged in about 11 seconds, comparable to the 17-s half life of ¹⁹Ne. To extend the time before the trap fills, we often run at 20 psig with a flow rate of 1.2 mole/hr (3.4 cc/s). Pressure measurements at the ends of the tube between the target and the manifold before the cold trap, combined with known mass flow rates and tube dimensions, indicate a transit time for the gas of only about 2 s for the 30 psig and 2.4 mole/hr example.

The SF₆ with the newly produced ¹⁹Ne passes to the cold trap, where the SF₆ is frozen while the ¹⁹Ne continues into the vacuum system for the experiment. Before using the SF₆

to produce ^{19}Ne , the supply from the commercial bottle was passed through several cycles, freezing it while pumping on the cold trap with the turbo-pump, and then warming the trap to put the SF_6 into the storage tanks. This process should remove air and other volatile contaminants. We found some CF_4 contaminant which was not readily removed by this process. Both the CF_4 and remaining SF_6 (vapor pressure in the 10^{-5} Torr range) have similar partial pressures in our vacuum system upstream of Cryo-II and are removed by it (see Figures. 3.5 and 3.6).

Figure 3.11 shows a profile of one of the SF_6 cold traps, which are immersed in liquid nitrogen. They are 11.4-cm diameter by 23.6-cm length 304 stainless steel cylinders with copper baffles soldered to the inside top and bottom to increase the area for freezing SF_6 . The body tube is 3-mm thick and the end plates 6 mm. The traps have safety relief valves set to 120 psig, even though they have been calculated to withstand 500-psig pressure.

Gas is input through a 1/4" stainless steel tube which enters the top of the cylinder off center. It directs the gas flow downward. ^{19}Ne exits through the 3.8-cm tube on center at the top. The input tube ends half way down in order to leave a substantial volume to accumulate frozen SF_6 . With the tube ending half way down as shown, and with several horizontal baffles above the tube termination to freeze SF_6 before it can get out the exit tube, the trap can freeze about 9 moles of SF_6 before the vapor pressure above the trap increases significantly. Thus we can typically run for 4 to 7 hrs before we need to empty the trap. As it takes about an hour to warm the trap and then cool it down again, we are modifying our system to use two traps. One can be used for ^{19}Ne production with one supply tank while the other is warming and filling another tank.

Calculations based on the cross sections of Jenkin et al. [64], which extend from 5 to 11 MeV proton energy, suggest a 12-MeV proton beam should produce $\sim 10^9$ ^{19}Ne $\text{s}^{-1}\mu\text{A}^{-1}$. This figure may be low, since cross sections measured by Riffle et al. [94] are about 50% larger than those of Jenkin et al. in the 5 - 6 MeV region where the measurements overlap. About half of this activity decays before exiting the target and passing through the cold trap. To determine how long the activity takes to reach the experimental apparatus, we switched

the beam on and off while measuring the time between beam on and arrival of the activity. The activity had a rise time of about 10 s and was delayed 16 to 20 s, for flow rates from 5.4 to 2.7 cc/s. The production rate only decreased by about 15% at the lower flow rate.

Table 3.1 shows a list of potential radioactive contaminants expected from $\text{SF}_6 + p$ at beam energies of $E_p = 12$ MeV. This energy is below the threshold for production of ^{32}Cl , which would be otherwise the most expected contaminant, given the relative abundances of the sulphur isotopes (95%, 0.25%, 4.25% for ^{32}S , ^{33}S , ^{34}S , respectively). Because most produce betas with endpoints higher than the 2.21-MeV value for ^{19}Ne , we assembled the thick-scintillator telescope mentioned in Section 3.2, so as to look for their signatures. It is important to note that both P and Cl are reactive which drastically decreases their chances of traveling all the way from the target to the decay cell (≈ 20 m through SF_6 gas, a cold trap, and a long vacuum chamber). We are also using half-life measurements of ^{19}Ne to put limits on the abundance of ^{18}F . The initial results don't suggest any significant presence of ^{18}F . The work to ensure that there are no contaminants in the decay cell capable of distorting a ^{19}Ne spectrum measurement is still in progress. Contaminants will be an important systematic as we move forward to more precise measurements.

3.4 He6-CRES apparatus

Above we presented the hardware development necessary to produce and transport high purity ^{19}Ne and ^6He to the He6-CRES apparatus. Here we describe the He6-CRES apparatus itself, including details on the decay cell and magnetic trap, the RF system, and the cryo-system and temperature distribution.

3.4.1 Decay cell and magnetic trap

The decay cell is itself a copper 11.56-mm inner diameter, ~ 15 -cm long circular waveguide. The decay cell was chosen to be circular as opposed to rectangular to minimize the probability of wall collisions for a given cross-sectional area (maximizing the event rate). Three coils generate the magnetic trapping field, typically $\mathcal{O}(10^{-3})$ of the main field, used to axially

Product	$t_{1/2}$ (s)	γ_1 branch ^a		β_1 branch ^b	
		E_γ	I_γ	K_{\max}	I_β
		(MeV)	(%)	(MeV)	(%)
²⁹ P	4.1	1.273	1.5	3.9	98
³⁰ P	150	n/a	$\ll 1$	3.2	100
³³ Cl	2.5	0.84	0.5	4.6	99
³⁴ Cl	1.5	n/a	$\ll 1$	4.5	100
^{34,m} Cl	1920	2.127	43	2.5	28
¹⁸ F	6586	n/a	$\ll 1$	0.6	97

^aMain gamma branch energy and intensity;

^bMain beta branch endpoint energy and intensity;

Table 3.1: Potential radioactive contaminants produced by the ¹⁹Ne source and their most prominent decay branches.

confine the betas (see Figure 3.12). The magnetic trap was designed to be quickly slewed (turned off and on) to empty the trap of β s. Once the trap current is at zero, the electrons collide with the decay cell windows and rapidly lose enough energy to be undetectable. So as not to induce an electromotive force (EMF) and associated current in the main magnet during trap slewing and potentially cause a magnet quench, the coils were designed to have no dipole moment (to first order). Therefore there is very little current induced in the superconducting solenoid when emptying the trap via ramping the current to zero. Calculations done by Kris Knutsen found that cycling the deepest traps we intend to use induces a total current of $\approx 1\mu\text{A}$ per cycle in the main magnet at a rate of $\approx 10\text{ mA/s}$ assuming a $100\ \mu\text{s}$ cycling time. The main magnet has a safe-ramping rate of 100 mA/s even at the highest fields ($5 - 7\text{ T}$), so induced currents of this magnitude and rate are a non-issue.

To obtain no dipole moment the center coil of the trap has twice as many windings (88 windings) as the outer two (44 windings) and the three coils are wired in parallel with alternating polarities. The field profile produced is shown in Figure 3.12 and agrees well with simulation.

The trap was constructed using two clam-shell pieces that clamp around the decay cell. The clam-shells are thermally isolated from the decay cell by G10 spacers. There is a layer of Kapton inserted between the two clam-shell pieces in order to cut down on eddy currents created when slewing the trap coils. Each coil was wound using 26-AWG copper-insulated magnet wire as seen in Figure 3.13. After each wound layer GE 7031 varnish was applied to set the wires in place and increase the heat conductivity of the coil. For a 10^{-3} trap in a 6 T main field (our maximum desired field strength) we need 1.45 A in our trap coils. The maximum current rating of the wires is 4 A.

We provide the trap current with a KEPCO BOP 100-4L power supply that is controlled by a SCPI-based python driver. We measured that the rise time for a 0.4 A current in the three coil trap at room temperature ($\sim 300\text{ K}$) was $\sim 100\ \mu\text{s}$, which aligns with the KEPCO data sheet. Using a low inductance pick-up coil and integrating the voltage according to Faraday's Law we measured that the field at the center of the coil had a rise-time constant

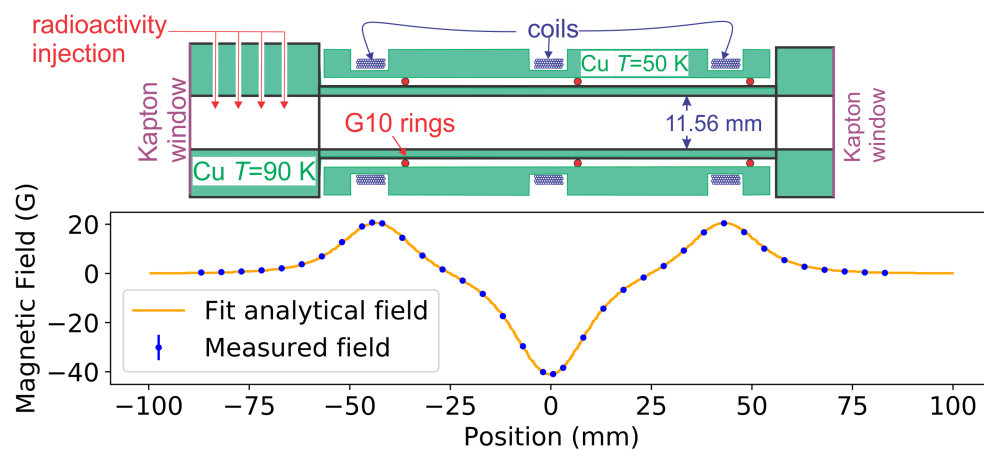


Figure 3.12: Top: decay cell geometry with trap coils. Bottom: axial magnetic field profile from trap coils used for axial confinement of β s.





Figure 3.13: Left: Photo of components of the decay cell and coil form. Kapton was placed between the two sides of the coil form and non-conducting nylon screws were used to prevent circumferential eddy-currents. Right: Decay cell and coil form fully assembled. The flange with a grid of small holes in the foreground lets the radioactivity into the decay cell.

of 1.3 ms (found by fitting the field magnitude to an exponential). According to the theory of field diffusion in a conductor the diffusion time is proportional to the conductivity of the conductor. In-situ the copper decay cell is at about 80 K so the conductivity is 7.8 times larger than at room temperature. Therefore, the expected in-situ 0-95% rise-time is 9.6 ms. This in-situ rise time was verified to within ± 1 ms using the curvature of CRES tracks during trap slewing.

We discovered that a ~ 10 ms rise time would negatively affect CRES measurements because it prevents preferable duty cycles (trap needs to be off for at least ~ 30 ms while eddy currents die down) and it causes tracks to curve as the trap field settles or turns off which confuses our analysis. Figure 3.14 shows the sparse spectrogram of a CRES electron during a 35 ms ON, 5 ms OFF trap slewing duty cycle. One can see that the event frequency is effected (curved) when the trap is turned off (see $t = 0.72, 0.76, 0.80$ s) but that the CRES electron stays trapped due to the long lasting eddy currents preventing the trapping field from adequately vanishing.

Since the clam-shell coil form itself has a radial conductive break, the eddy currents are primarily caused by the copper decay cell which has a 170 mil³ wall thickness. In an effort to address this issue we had a new decay cell with 13 mil walls machined as shown in Figure 3.15. This significantly reduced the rise-time. Based on calculation and verified with room-temperature measurements, we expect the new decay cell to have an in-situ (80 K) 0-95% rise-time of ≈ 2 ms. With this improvement and a protocol of keeping the trap OFF time to ≥ 15 ms, we no longer have significantly curved tracks or any failures to empty the trap.

It's important to note that trap emptying was not an issue for *Phase I* or *Phase II* of Project-8 because the (≈ 18 keV) CRES electrons naturally scattered out of the trap due to collisions with residual gas atoms. At most of the β energies we are interested in (≥ 100 keV) scattering is much more rare, making trap emptying a crucial component of the experiment

³1 mil = 10^{-3} inches.




Figure 3.14: Sparse spectrogram of a CRES electron during a 35 ms ON, 5 ms OFF trap slewing duty cycle. One can see that the track is affected (sidebands collapse toward the main band) when the trap is turned off (see $t \approx 0.72, 0.75, 0.79$ seconds, indicated with red dotted lines) but that the CRES electron stays trapped. This is due to the long lasting eddy currents preventing the trapping field from adequately vanishing.



Figure 3.15: Photo of the thin walled decay cell (13 mil wall thickness) built in order to overcome the eddy current issue. One can see the step down to the thinner walls in the central region of the cell.




Figure 3.16: Design drawing of the Ti transport tube that delivers radioactive isotopes to the decay cell. Originally this tube was designed and built with a 1/4" ID but we found we had better isotope transport and pumping with the larger 3/8" ID.

in order to prevent an abundance of higher-energy CRES events from ‘rising’ into our visible frequency bandwidth. The long tracks we observe will be discussed in detail in Chapter 6.

Our N.C.-State collaborators are designing an electrostatic $\vec{E} \times \vec{B}$ trap-emptying device that will be able to empty the trap in less than 500 μs without any induced eddy currents (or changing the trap current) by providing a brief strong voltage to a section of the decay cell. This should increase the trap’s work to rest ratio and in turn increase our CRES rate within a fixed amount of data taking. It would also drastically improve the issue of CRES events rising from below by enabling shorter trap-on periods.

The radioactivity is brought from the vacuum system portrayed in Figure 3.5 into the decay cell via a 3/8" inner diameter (ID) titanium tube leading to a grid of 25 small (1 mm diameter) holes leading to the inside of the circular waveguide (see Figures 3.16, 3.12, and 3.13). The grid of small holes was shown to not have a measurable effect on the RF properties via bench-top measurements and Ansys HFSS (high-frequency structure simulator) simulations. The radioactivity is kept contained in the decay cell (axially) via 5 mil

Kapton (polyimide) foil windows that are transparent to microwaves. There was worry that the Kapton foil could adsorb the ${}^6\text{He}$ onto its surface to such an extent that the radioactive partial pressure in the decay cell would suffer (hurting our CRES statistics). Calculations and experimental efforts were conducted and showed that this should not be a concern for ${}^6\text{He}$. Given that we found this to be such a small effect these results should generalize to ${}^{19}\text{Ne}$. These calculations and experimental efforts are presented in Appendix B.

3.4.2 RF system

The RF system was designed so that only one waveguide mode is excited within the directly observable bandwidth (17.9 – 19.1 GHz). The mode cutoff frequencies for circular waveguides are given by,

$$\begin{aligned} TE : f_{c_{mn}} &= \frac{p'_{nm}}{2\pi a \sqrt{\mu\epsilon}} \\ TM : f_{c_{mn}} &= \frac{p_{nm}}{2\pi a \sqrt{\mu\epsilon}}, \end{aligned} \quad (3.7)$$

where p'_{nm} denotes the m th zero of J'_n (the derivative of the n^{th} Bessel function of the first kind), p_{nm} denotes the m th zero of J_n (the n^{th} Bessel function of the first kind), and a is the radius of the guide [88]. The mode cutoff frequencies for rectangular waveguides are given by,

$$TE, TM : f_{c_{mn}} = \frac{1}{2\pi \sqrt{\mu\epsilon}} \sqrt{\left(\frac{m\pi}{a}\right)^2 + \left(\frac{n\pi}{b}\right)^2} \quad (3.8)$$

where the TE_{00} , TM_{00} , TM_{10} , and TM_{01} modes are not allowed and a , b are guide side lengths with $a > b$. See Reference [88] for a thorough treatment of waveguide modes and cutoffs. The circularly-polarized cyclotron radiation produced in the circular decay cell couples to the TE_{11} mode (cutoff 15.203 GHz). The next available mode is TM_{01} which opens at 19.890 GHz. After the decay cell there is a $\lambda/4$ phase-retarding polarizer that converts the circularly polarized radiation to a linear polarization state. The polarizer is followed by a




Figure 3.17: Labeled design of the components surrounding the decay cell. Legend: A = WR-42 to Circular adapter, B = $\lambda/4$ polarizer, C = trap coil form surrounding the circular decay cell, D = grid of holes allowing radioactive isotopes into the decay cell, E = Titanium radioactive isotope transport tube, F = WR-42. The kapton windows are compressed between the flanges of B and C.

circular to WR-42⁴ adapter, and WR-42 rectangular waveguide (10.668 mm x 4.318 mm). The signal couples to the TE_{10} mode of the rectangular guide (cutoff 14.051 GHz). The next available mode in the rectangular guide is TE_{02} with cutoff 28.102 GHz. The WR-42 leads to a circulator which directs radiation towards the amplifier chain. The circulator prevents noise sourced from the low noise amplifiers (LNAs) from making its way back into the RF system and also allows for the injection of signals for RF characterization and testing via the 3rd circulator port. The signal-injection input port is composed of a WR-42 to a cryogenic 20 dB coaxial adapter connected to a stainless steel coaxial cable which leads to a vacuum feedthrough. This system transmits and amplifies well for frequencies in the 18 – 24 GHz band, so there is the opportunity to increase our bandwidth without changing any RF hardware. Figure 3.17 shows the RF configuration surrounding the decay cell. Figure 3.18 shows the contents of the RF system surrounding the LNA chain.

⁴WR stands for ‘Waveguide Rectangular’ and the number following it indicates the waveguide dimensions inner width in hundredths of an inch. WR-42 is a standard size used by manufacturers.




Figure 3.18: Photo of internal contents of the copper box (the box itself is labeled F). The cubic component with three WR-42 inputs in the foreground is the circulator (labeled A). The WR-42 waveguide (labeled B) provides radiation transport from the decay cell (out of picture to the right) to the circulator. The first stage LNA (labeled C) takes WR-42 as input and outputs coaxial. The second stage LNA (labeled D) takes coaxial as input and output. The cold-finger from the Cryocooler (labeled E) keeps the LNAs around 5–8 K. The silver material on the left is the aluminumized-mylar super-insulation which helps reduce the amount of heat deposited on the cold components from black-body radiation.

The amplifier chain is composed of a Low Noise Factory cryogenic waveguide-input/coax-output amplifier (1st stage), followed by a cryogenic coax-input/coax-output amplifier (2nd stage). Each set of cryogenic amplifiers has a combined gain of approximately 60 dB with a noise temperature around 8 K when held at a physical temperature of 5 K. The noise characteristics of the RF system will be discussed in detail in Section 3.9. The RF system from decay cell to LNA chain described above is duplicated such that signals propagating out of both sides of the decay cell can be detected. In order to allow for both sets of LNAs to be cooled by a single cryocooler, one piece of WR-42 contains a U, and we refer to this side of the RF system as the U-side. We refer to the other (straight) side as the I-side. Figure 3.19 and Figure 3.20 illustrate how the RF system fits into the scheme of the entire apparatus. Table 3.2 contains details on each of the RF components described above. Note that depending on the β energy there is non-trivial coupling to higher modes of the cyclotron frequency f_c (as will be discussed in Chapter 6) but this power is not directly observed via Fourier bin power, it is only observed indirectly via the track slopes.

A “photonic crystal” was designed to allow for a thermal break in the RF system while maintaining good RF transmission. Four photonic crystal waveguide flange pieces (shown in Figure 3.21) are used in total, two on the U-side and two on the I-side. These flanges have a tiling of square posts roughly $1/4$ wavelength on each side, and are designed to cause maximum destructive interference for waves that would otherwise propagate out of the waveguide (leading to signal degradation) given a simple vacuum gap in the RF system. The use of these photonic crystal flanges allows us to implement a vacuum gap using kapton standoffs from the cryogenic pieces of the waveguide, without appreciable loss of RF power. Vector network analyzer (VNA) measurements of the flanges suggest that millimeter-scale gaps can be introduced without degrading the RF transmission by more than 0.25 dB; substantially better than two flat WR-42 flanges separated by an identical gap.

Once the RF signal exits the LNA chain a stainless steel coaxial cable (chosen for its poor thermal conductivity) brings the amplified signal through a vacuum feedthrough to the ambient (room temperature) receiver where the signal is mixed with a 17.9 GHz reference




Figure 3.19: Full He6-CRES apparatus design. The large grey cylinder is the magnet.

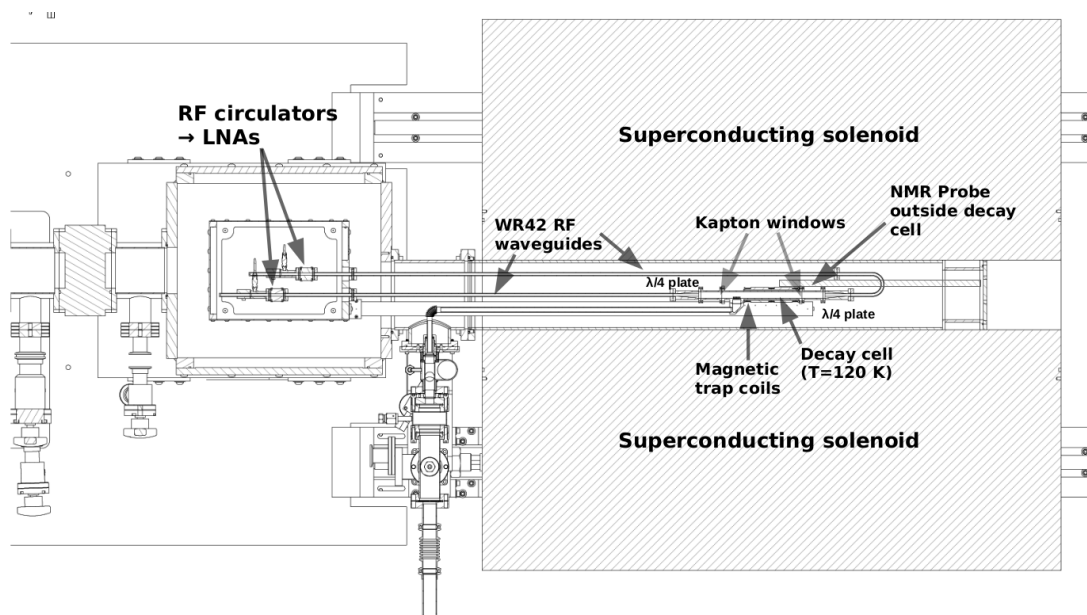


Figure 3.20: Labeled horizontal-plane section of the He6-CRES apparatus. Figure created by Heather Harrington of He6-CRES.




Figure 3.21: Left: Photo of a photonic crystal. Four of these are used as a thermal break within the RF system of the He6-CRES apparatus. Right: Design of photonic crystal.

Component name	Conductor	Size	Company	Part number
Decay cell	copper	ID = 11.56 mm	Custom	N/A
$\lambda/4$ polarizer	gold coated	ID = 11.56 mm	Quin Star	QWL series
Circular to rectangular adapter	gold coated	N/A	Quin Star	QWC series
Rectangular WG	gold coated	WR-42	Sage Millimeter	SWG series
Rectangular U-bend WG	gold coated	WR-42	Quin Star	QAC series
3-port Circulator	gold coated	WR-42	Quin Star	QKY series
Rectangular to coax adapter	gold coated	WR-42 to coax	Quin Star	QWA series
LNA, 1st stage	gold coated	WR-42 to coax	Low Noise Factory	LNC16 28WB
LNA, 2nd stage	N/A	coax to coax	Low Noise Factory	LNC15 29B

Table 3.2: Details on the RF components used in the He6-CRES *Phase 1* apparatus.

oscillator, down-converting signal frequencies from 17.9–19.1 GHz down to 0–1.2 GHz. This is much broader than the 180-MHz total bandwidth used to measure the tritium spectrum with CRES [91]. This lower frequency signal is then fed into the Roach2 Field-programmable gate array (FPGA) board discussed in Section 3.5. Before injection into the ROACH2, a 100 MHz high pass filter (HPF) is applied to kill low frequency aliased signals. Due to this HPF we will often refer to our frequency bandwidth as (effectively) 18.0–19.1 GHz. Currently we only take data out of either the U-side or the I-side.

3.4.3 Cryo-system and temperature distribution

Obtaining the correct thermal distribution throughout the He6-CRES apparatus was a challenge (championed by John Amsbaugh on the engineering side) that took extensive trial and error. A two-stage Gifford-McMahon type cryocooler is used to cool components within the apparatus (Part info: Sumitomo Model SRDK-408D2-F70L cryocooler system, purchased from Janis). The first stage of the cryocooler is 40 K and the second stage is 5 K. If too

much heat is loaded onto the cryocooler, it will fail to keep the two stages at the specified temperatures. The second stage is capable of absorbing 1.5 W of power. The LNAs must be held as close to 5 K as possible in order to minimize their noise temperature. The decay cell needs to be held at ~ 100 K so as not to freeze out the $^{83\text{m}}\text{Kr}$ that we used to calibrate and sanity check the RF system and DAQ. Once we transitioned to taking data with ^6He and ^{19}Ne (and away from $^{83\text{m}}\text{Kr}$), we used a copper braid to reduce the temperature of the decay cell to ~ 80 K in order to further reduce RF noise. Photonic crystals, two on either side of the RF system, (discussed above, shown in Figure 3.21) are used as thermal breaks between the decay cell and the LNAs, as they would otherwise be in direct thermal contact. A copper box (seen in Figure 3.18) connected to the first 40 K stage is used to thermally shield all 5 K components. Copper bars (often referred to as the ‘heat bus’) are mounted to the copper box and bring cooling to the components in the cylindrical nose containing the RF components (see Figure 3.2).

To further reduce the heat load as much as possible we have wrapped the majority of components, within our vacuum system in super-insulation consisting of 6-10 alternating layers of polyester tulle and aluminized Mylar. This successfully cut down on the heat power from black-body radiation dumped onto the cryocooler (see left side of Figure 3.18). We have coupled different components to heat sinks (such as the copper box and copper bars) via copper braid in order to fine tune their temperature. We use silicon diodes and platinum resistance temperature detectors (RTDs, suitable for high field regions) to monitor the temperatures of key components. The temperature distribution of the apparatus is currently satisfactory in that the temperature of the terminators, amplifiers, and waveguides are sufficiently low as to satisfy our RF noise budget (discussed in Section 3.9). The copper clam-shell on which the electron trap coils are wound is cold enough to ensure there is not runaway current in the trap coils. Table 3.3 and Figure 3.3 provide an overview of the temperature distribution profile of the apparatus.




Figure 3.22: Labeled design drawing. Legend and temperature ranges for labeled components found in Table 3.3.

Component name	Temperature (K)	Label in Figure 3.22
LNAs, (both stages)	7	A
Copper box	40	B
WR-42, (U and I-side)	65 – 75	C
Terminator/Attenuator	10–20	D
Copper bars	40 – 50	E
Decay cell	70 – 100	F
Clam-shell coil form	40–50	G

Table 3.3: Temperature ranges for components in He6-CRES apparatus. Corresponding labeled design drawing found in Figure 3.22.

3.5 DAQ

From the ambient receiver the 0.1 – 1.2 GHz signal is brought to a ROACH2 FPGA board (designed by the CASPER radio astronomy collaboration [17]) that digitizes the time-series data at 2.4 GHz in $\approx 3.4 \mu\text{s}$ segments, then converts each segment into the frequency domain by fast Fourier transform (FFT). The magnitude of the FFT is taken (phase information is lost) and adjacent time slices are summed (resulting in $\approx 6.8 \mu\text{s}$ time slices) to save a factor of 4 in data rate relative to the time domain signals. The data is then transported to our Cave 1 DAQ computer via a 10 Gigabit Ethernet (GbE) cable and written to disk at $\approx 600 \text{ MB/s}$. The DAQ computer has $\approx 10 \text{ TB}$ of disk space ($\mathcal{O}(10^4 \text{ s})$ of CRES data), so data is transferred via an `rsync` to a 60 TB RAID 2 drive that is built into the CENPA computing cluster (rocks). Analysis can then be conducted from within rocks.

The ROACH2 FPGA board was programmed as a broadband spectrometer with the Xilinx software (version ISE 14.7) following public tutorials created by the CASPER collaboration. The Xilinx software outputs a ‘bitcode’ which is transmitted to the board CPU on boot-up and contains instructions on how the board should program itself. The ROACH2 uses a rubidium-based 10 MHz reference as its internal clock (Part info: Stanford Research Systems PRS10 10 MHz rubidium oscillator). A 0 – 6.0 GHz signal generator is used to calibrate the ROACH2 ADC cores’ phase, offset, and gain (removing artifacts at integer divisors of the sampling rate).

The choice to not write time domain data to disk was motivated by the desire to maintain a large frequency bandwidth. To write out time domain data we would have had to reduce the active bandwidth by roughly a factor of 4 to have sufficient data throughput given our current hardware constraints. The choice to take FFTs in real time necessitated choosing a fixed number of samples per FFT, N , to hard-code into the FPGA board. Future iterations of the experiment will need to consider if this was the optimal approach. If you increase N for a fixed sampling rate f_s (currently 2.4 GHz) then you increase the time per FFT Δt (slice length) and decrease your frequency bin width Δf while maintaining the same total

bandwidth $f_s/2$ (with $N/2$ positive frequency bins). The relationship to keep in mind is,

$$\Delta f = \frac{1}{\Delta t} = \frac{f_s}{N}. \quad (3.9)$$

The SNR of a chirp of slope m (Hz/s) is highly dependent on the choice of N . There is a balance to be struck between the noise power in a given frequency bin which scales linearly with Δf , and the time a signal is present in a given Fourier bin which scales with Δt . The optimal SNR is obtained when,

$$\Delta f = \sqrt{m} \quad \rightarrow \quad N_{optimal} = \frac{f_s}{\sqrt{m}} \quad (3.10)$$

Appendix A provides a simple proof for why the SNR of a chirp is optimized by the condition given in Equation 3.10. In the energy regime of interest (0 - 2.2 MeV) we primarily observe event slopes between 10^9 Hz/s and $1.5 \cdot 10^{11}$ Hz/s (more discussion on track slope vs energy in Chapter 6). These slopes correspond to $N_{optimal} \in (2^{12.6}, 2^{16.2})$. Given that the ROACH2 must be programmed with N as an even power of 2, we chose to use $N = 2^{13}$, which corresponds to $2^{12} = 4096$ frequency bins. This corresponds to a frequency bin size of 292.97 kHz. Note that for lower energy ^{83m}Kr data (discussed more in Chapter 6), a bitcode with $N = 2^{16}$, which corresponds to $2^{15} = 32768$ frequency bins and a frequency bin size of 36.62 kHz was used.

The SNR in the main Fourier bin for a fixed-power chirp vs the slope for different DAQ number of frequency bins in log base 2 is shown in Figure 3.23. This plot illustrates that regardless of the optimal choice of N , SNRs will roll off for higher slopes. As stated above, we choose to conduct all *Phase 1* measurements with one bitcode (with $N = 2^{13}$) for simplicity. This implies there is a slope-dependent SNR response. This effect is cancelled to first order in the ratio but future versions of the experiment could obtain a significant increase in signal SNR (and therefore event rate) by writing time domain to disk and using match-filtering with different templates for each slope and corresponding field to resolve chirps. Significant hardware development will be necessary to do this without a factor of 4 reduction

in bandwidth.

Each slice is composed of 4096 (8 bit resolution) frequency bins. Data is streamed from the ROACH2 to the DAQ computer over the 10 GbE as User Datagram Protocol (UDP) packets, with one slice per UDP packet. Each packet contains a 32 byte header (8 bits = 1 byte) that includes a sequential packet ID that wraps around to 0 at 2^{20} . Thus each packet is 4128 Bytes, and a 600 MB/s data-rate implies $\approx 1.5 \times 10^5$ packets/s. Initially, given this high packet rate, chunks of dropped packets causing large $\mathcal{O}(10 \text{ ms})$ gaps in spectrograms plagued us. Ultimately, instructions found in Reference [18] describing how to optimize the Ethernet Network Interface Card (NIC) transmit queue and socket receive buffer size in a Linux machine for fast transmission of UDP packets resolved this issue. We now only observe a few dropped packets for every 100 s of CRES data taken (fraction dropped now $\sim 10^{-7}$, down from $\sim 10^{-2}$ before the network improvements).

A receiver script written in C++ using low level socket libraries is responsible for writing data to disk as binary files. We have given our frequency-domain data files the file extension `.spec`. In order to keep file sizes reasonable each `spec` file typically contains one second of CRES data ($\approx 600 \text{ MB}$). Initially, this receiver was designed in a multi-threaded fashion where there were three active threads; one writing, one receiving, and one waiting. This prevented any downtime, meaning data was continuously written to disk. We found that this multi-threading was also contributing to dropped packets so we currently have only one active thread that receives 1 s worth of packets ($\approx 1.5 \times 10^5$ packets), and then writes them to disk, and repeats. The ratio of receiving time to writing time (not receiving packets) is roughly 1:1.

3.6 Magnet

To provide the main magnetic field for CRES measurements, we obtained a refurbished 7-Tesla AMI superconducting magnet (serial number 8946) from our collaborators at Argonne National Laboratory. The magnet's bore is horizontal and designing a movable experimental insert would prove challenging with the complex and rigid vacuum system that supplies




Figure 3.23: SNR vs chirp slope for different choices of bitcode. Choice of bitcode corresponds to log base 2 of the number of frequency bins. For the low energy $^{83\text{m}}\text{Kr}$ data presented in this thesis 2^{15} was used, and for the higher energy ^6He and ^{19}Ne data 2^{12} was used. Figure created by Nick Buzinsky of He6-CRES.




Figure 3.24: Design drawing of He6-CRES apparatus, super-conducting solenoid and the ceramic coated aluminum rails (on the floor) used to extract the magnet from the experimental insert.

radioactive isotopes. We opted to put the magnet on ceramic coated aluminum rails with Frelon-lined linear bearings allowing it to move on-axis in and out of position around a fixed cylindrical experimental insert containing the decay cell and RF waveguides. The use of the rails along with a stopper that aligns the center of the decay cell with the center of the magnet bore allows for repositioning of the magnet relative to the bore to within ~ 1 mil. Figure 3.24 illustrates the positioning of the magnet, rails, and the cylindrical insert assembly.

The magnet is composed of an outer liquid nitrogen (LN2) jacket with an inner liquid helium (LHe) cryostat that the superconducting coils are situated within. Between the LN2 and LHe spaces is a vacuum space and layers of super-insulation to minimize the thermal load on the LHe space. We are fortunate to share a facility with the Axion Dark Matter eXperiment (ADMX), who have a helium recovery and liquefaction system for their own superconducting magnet. We constructed a ~ 150 -ft helium gas recovery line from our superconducting magnet to their system allowing for the capture and re-liquefaction of our boil-off helium. Given that we boil off 3 – 4 % of our 200 liter LHe magnet cryostat (6 – 8

liters) per day and LHe is ≈ 200 USD per liter (in 2023), this greatly mitigates running costs. Over the past five years safe procedures have been developed for all magnet processes; filling the magnet with cryogenics routinely and during an initial cool-down, ramping the magnet up and down in field, retracting the magnet, etc.

With the magnet retracted from the experimental insert, the magnetic field is mapped using a PT2025 NMR Precision Teslameter system from Metrolab. This Teslameter system achieves 5ppm absolute accuracy and $0.1 \mu T$ resolution for measurements and mappings of magnetic fields in the field range of interest (0.5 - 5.0 T). The magnet is quoted to be able to achieve 1 ppm/cm homogeneity at the center of the bore. It was found that without shimming the field homogeneity in the region of interest (within the center of the decay cell) was ≈ 15 ppm/cm for the main field values of interest to this work (.75 T - 3.25 T). With shimming we found we can obtain $\approx 2 - 5$ ppm/cm but in the short term shimming at each field of interest is too time-consuming to be practical. The un-shimmed level of uniformity is more than sufficient for the measurements shown in this thesis, however future iterations of the experiment will need to be more concerned with field determinations. The main field needs to be known to within $\mathcal{O}(10^{-4})$ relative uncertainty for a $\mathcal{O}(10^{-3})$ uncertainty in b_{Fierz} , which is easily attainable with this magnet.

To allow for precise real-time monitoring of the main field during data acquisition a stainless steel tube welded to a circular vacuum flange was built into the end of the cylindrical insert assembly (see Figure 3.25). During data acquisition the NMR tube is slid into the end of the ‘re-entrant’ tube, allowing for the probe to be within ≈ 3 cm of the center of the decay cell. Future versions of the experiment can use a full field map (done previous to data acquisition) to extrapolate from the NMR probe value to the field value at the center of the decay cell. Because this difference in field is of $\mathcal{O}(100 \text{ ppm})$ or less, we are safe to just use the field value at the NMR probe as the main field in this work. For the spectral ratio measurements shown later in the thesis it is more crucial how well matched the field values used for the two different spectral measurements are (these data are taken days or months apart). The repeatability of the field (distinct from the precision of our knowledge of the




Figure 3.25: Design drawing of the ‘re-entrant’ tube (coming in from the left) used to put the NMR probe in close proximity to the decay cell during data acquisition. The three dark bands at center are the magnetic trapping coils.

field) is about 10^{-3} and will be discussed in more detail in Chapter 6.

3.7 β monitor

As was discussed briefly in Section 3.1, the need for a stable β monitor for spectral normalization is critical to our measurement technique. For a 10^{-3} measurement of b_{Fierz} , the β monitor must have a 10^{-4} relative stability over the period of data acquisition, as will be discussed in more detail in Chapter 5. Here we describe the two β monitors we use and present their stability.

To monitor radioactivity during data acquisition and run production tests, a decay volume consisting of a 3.8-cm-diameter 12.5-cm-long SS cylinder (a standard Conflat nipple) with a 0.127-mm-thick copper window, connected to the rest of the system as shown in Figure 3.5, is used. Most of the time a *beta monitor* detector is placed in front of the copper window. This consists of a 2-cm-by-2-cm 5-mm-thick scintillator (Eljen Technologies, EJ-212) read by 4 SiPMs (Hamamatsu S14160-3050HS) with power supplies compensated for temperature variations (Hamamatsu C11204-02) to yield high response stability. A thin single scintillator

was chosen over telescope systems to minimize variations in efficiency due to small gain changes.

In addition, at times, a *beta telescope* system is used instead in order to place upper limits on radioactive contaminants with higher endpoints than the radioactivities of interest (described above). This telescope consists of a 3-mm-thick (ΔE) followed by a 5.1-cm-thick (E) scintillators, both about 5.1 cm in diameter, read by individual phototubes.

During the ratio measurement data campaigns presented in Chapter 6, both the *beta monitor* and the *beta telescope* were used for spectral normalizations. The specific reasons for this will be discussed in Chapter 6. The stability of the monitor over time scales equivalent to the time it takes to measure one isotope is what is relevant experimentally. We conduct stability tests over a 24 hour period, as all data campaigns were more or less contained within 24 hours. Longer stability tests will need to be conducted when longer data campaigns are needed.

We use three metrics to assess the stability of the two monitors; the Allan Deviation, the rolling average rate over time, and the mean bin (a proxy for the gain). The Allan Deviation is given by [95],

$$\sigma_y(\tau) = \sqrt{\langle (\bar{y}_{n+1} - \bar{y}_n)^2 \rangle} \quad (3.11)$$

where \bar{y}_n is the n th fractional frequency average over the observation window τ . The fractional frequency is the normalized difference between adjacent bin rates over time τ . Intuitively, the Allan deviation starts off large for small τ due to large fluctuations over small averaging times. It decreases at longer τ because you are averaging out the noise but at some point it begins to rise again due to systematic effects such as gain drifts, temperature changes, etc. The time scale of τ at which this increase occurs tells you the time scale over which the frequency is stable. The rolling average rate over time provides information about the magnitude of this instability.

Figure 3.26 and Figure 3.27 show the stability metrics for the *beta monitor* and the *beta telescope* respectively. This data was obtained with a ^{90}Sr source that generates β s (via

decay to the shorter lived ^{90}Y) with mean energy ≈ 0.93 MeV. The average rate and mean bin presented in the left plot of Figure 3.26 and Figure 3.27 are normalized to unity for easy stability assessment. In Chapter 5 we will discuss the monitor stability needed to reach a 10^{-3} measurement of b_{Fierz} . To first order an Allan deviation stable down to 10^{-4} is needed, which we haven't yet reached. Development in this area of the experiment is ongoing. The levels of stability presented here are easily sufficient for the $\mathcal{O}(10^{-1})$ proof-of-principle b_{Fierz} determination demonstrated in this thesis.

3.8 DAQ user interface and SQL database

The He6-CRES data acquisition system user interface (DAQ-UI) enables the experimenter to take CRES data and monitor slow controls data (such as magnetic field strength, β monitor rate, vacuum pressures, etc.) remotely. This is necessary given that the DAQ computer in Cave 1 is within a high radiation region during CRES data taking (see Figure 3.4). The user interface is responsible for making the initial connection between the DAQ computer and the ROACH2 FPGA board and providing the instructions for the programming of the board (the bitcode). Once programmed, the ROACH2 continually streams data over the 10 GbE connection towards the DAQ computer, as was discussed in Section 3.5. However, data is only written to disk when the user sends a command over the DAQ-UI. Data is typically acquired in an acquisition of 100 ($\approx 600\text{MB}$) `spec` files, each containing 1 continuous second of CRES data. Each of these blocks of data is assigned a unique `run_id`, each file in the block is assigned a unique `spec_id`, and each `spec` file name contains the date and time the `spec` file was written to disk (`DATE.TIME.spec`).

The user initiates an acquisition (or run) with the method `CLI.take_FD_data(args)` (`CLI` indicates an instance of the DAQ-UI class) and provides (as arguments) the number of files to acquire, the length of each acquisition, the isotope being studied, the main magnet set-field, the RF side being used (U or I-side), notes on the run, and whether or not to delete and rerun files with any dropped packets. When a run is started, all of the arguments described above and all relevant slow controls data are written to the He6-CRES PostgreSQL database




Figure 3.26: Stability of *beta monitor*. Top: Allan deviation over a 24 hour stability run. Bottom: One hour rolling normalized average of the rate and mean bin (a proxy for the monitor gain).




Figure 3.27: Stability of *beta telescope*. Top: Allan deviation over a 24 hour stability run. Bottom: One hour rolling normalized average of the rate and mean bin (a proxy for the monitor gain). Note that “tau” is the observation window for the Allan deviation whereas “Time” refers to the wall-time associated with the test.

(referred to as the `he6cres_db`).

There are four tables within the `he6cres_db` that get written to during data acquisition: `run_log`, `spec_files`, `nmr`, and `monitor`. The `run_log` and `spec_files` tables get written to each time data is acquired whereas the `nmr` and `monitor` tables are written to continuously every 2 – 10 s during a sustained data run and contain the NMR probe value and the β monitor rate respectively. This enables the experimenter to monitor the radioactivity and field stability separate from data acquisition in order to decide when it's appropriate to begin an acquisition. The `nmr` and `monitor` tables are used to assign a field value and monitor value to each 1 s data file, as is discussed in Chapter 6. Details on each column of the four tables (`run_log`, `spec_files`, `nmr`, and `monitor`) is found in Appendix C.

3.9 RF noise and SNR characteristics

The noise temperature of the apparatus, and the resultant signal-to-noise ratio (SNR) are crucial components of the He6-CRES experiment that have a strong impact on the CRES event rate. In this section we cover the most important (both expected and measured) noise and SNR parameters of the *Phase 1* apparatus.

3.9.1 RF noise background equations

Chapter 10 of Reference [88] has excellent explanations regarding the noise characteristics of RF systems. We summarize the equations and concepts needed for this discussion.

A white noise source with gain G and noise power N can be modeled as a resistor of temperature T_e . The equivalent noise temperature of the source is given by [88],

$$T_e = \frac{N}{GkB}. \quad (3.12)$$

where k is Boltzman's constant and B is the frequency bandwidth of the component or system. Microwave components of all kinds are characterized by their equivalent noise temperature T_e which intuitively describes what temperature resistor would create the level of

noise produced by the component, given the bandwidth B . Since bandwidth is generally independent of noise characteristics it is often useful to discuss the noise density of a system or component,

$$T_e = \frac{n_{density}}{Gk}, \quad (3.13)$$

where $n_{density}$ has units of dBm/Hz.

The equivalent noise temperature of a lossy transmission line (for example a piece of waveguide) at physical temperature T is given by [88],

$$T_e = (L - 1)T = \left(\frac{1}{G} - 1\right)T. \quad (3.14)$$

Here $L = \frac{1}{G}$ is defined as the loss. The intuition for this formula is that if the transmission is perfectly lossless, the gain will be 1 and the equivalent noise temperature of the transmission line will be zero because none of its temperature is coupled into the system. For our purposes this means that the better transmission we have the lower the total equivalent noise temperature of our entire RF system and the larger our expected SNR.

If multiple RF components are connected in series the cascaded equivalent noise temperature T_{cas} of the entire RF system is given by [88],

$$T_{cas} = T_{e1} + \frac{T_{e2}}{G_1} + \frac{T_{e3}}{G_1G_2} + \dots. \quad (3.15)$$

Here T_{ei} and G_i are the equivalent noise temperature and gain of the i^{th} component in the series. Given the nature of our apparatus, one would also like to know the expected output signal-to-noise ratio (SNR) for a signal originating at the beginning of a cascade of RF components. If S is the signal power at the source, G is the total linear gain of the multi-component RF system, and T_{cas} is the cascaded noise temperature of the system then the SNR at the output of the cascade will be given by,

$$SNR = \frac{GS}{GkBT_{cas}} = \frac{S}{kBT_{cas}}. \quad (3.16)$$

We see that the total gain actually cancels from the equation and doesn't contribute to the SNR. It is the T_{cas} in the denominator, which is dominated by the physical temperature and gain of the early stages or components of the cascade, that matters.

3.9.2 Simplest working model of the He6-CRES RF system

The key take-away of Eq 3.15 is that it is the early stages of the RF system that determine the noise characteristics. Therefore it is these early stages, such as the terminator, decay cell, WR-42 waveguide components, and the first stage Low Noise Amplifier (LNA) that should be engineered (or purchased) in order to obtain the lowest possible noise floor and highest possible CRES SNR. The cascade analysis discussed below in Subsection 3.9.4 models the entire RF system from the decay cell to the input to the ROACH2 ADCs. This is important to understand the absolute powers input to the ROACH2 ADCs we expect to obtain from CRES events or noise but it isn't necessary to track all of these RF components to pin down what SNR we expect. The equivalent noise temperature is essentially fixed by the time the signal exits the second stage LNA. As a concrete illustration of how little the components after the second stage amplifier matter; for reasonable input parameters the equivalent noise temperature changes from 30.28 K to 30.30 K (a 0.06% change) from the output of the second stage LNA to the input of the ROACH2 ADC.

Below is the simplest working model for obtaining the (technically approximate) cascaded equivalent noise temperature of the current He6-CRES RF System (T_{cres}).

$$T_{cres} = T_{term} + T_1 \left(\frac{1}{G_1} - 1 \right) + \frac{T_{e2}}{G_1} + \frac{T_{e3}}{G_1 G_2} \quad (3.17)$$

Table 3.4 indicates how each part of the apparatus is represented in Equation 3.17. The first term of Equation 3.17 is the terminator's physical temperature. What is relevant is the

Parameter	Denoted in Equation 3.17
Terminator Temperature (K)	T_{term}
Waveguide Temperature (K)	T_1
Stage 1 LNA Equivalent Noise Temp. (K)	T_{e2}
Stage 2 LNA Equivalent Noise Temp. (K)	T_{e3}
Gain of Waveguides (linear)	G_1
Gain of Stage 1 LNA (linear)	G_2
Gain of Stage 2 LNA (linear)	G_3

Table 3.4: Symbolic legend for Equation 3.17.

temperature of the terminator that is connected to the side of the apparatus of interest; it is the U-side terminator temperature that effects the U-side noise level. This is due to RF reflections that are discussed in Section 3.9.3. The second term of Equation 3.17 comes from the (minimally) lossy waveguide components that transport the signal from their source to the circulator and is derived using Equation 3.14. This groups many waveguide components and the decay cell into one term, which is simplistic. But since it is the waveguide components (mainly the quarter wave plate and WR-42 to circular waveguide transition) near the decay cell that have the most loss in this chain and these components are all at approximately a single temperature, it is a reasonable simplification. The third and fourth terms represent the LNA contributions. One can see that the fourth term is largely suppressed by the gain of the first LNA. Its contributions are relatively minuscule for our current RF configuration (on the order of .1% in the total noise temperature of the system). The term is left in place in case future iterations of the experiment were to consider what type or quality of 2nd stage LNA they could install without SNR degradation. For example Equation 3.17 can be used to consider if a room temperature 2nd stage LNA could be viable.

This model can be used to gain an intuition for what parts of our RF system have the largest impact on the noise temperature and then Equation 3.16 can be used to investigate the corresponding SNR of a CRES signal.

3.9.3 Signal injection for RF calibration

The ability to inject CRES-like signals into our apparatus and verify that the entire RF chain is working as expected was crucial to our initial sanity checking of the entire RF chain and to our first CRES detections. This was accomplished by injecting signals into the third port of the circulator situated below the LNA chain and having that signal reflect off of the decay cell assembly and travel back towards the LNA as a CRES signal would (see Figure 3.18). Below we discuss the microwave theory behind why this reflection takes place and also present measurements done to verify our theoretical understanding. The Vaunix LMS-203 Lab Brick Signal Generator (10 - 20 GHz) is what we used to generate these injected signals.

For circularly polarized radiation sourced from a β within our decay cell the quarter wave plates at either end of the decay cell must be aligned (they must share the same slow axis) in order to get signals propagating to both ends of the RF system. The circularly polarized light must be transformed to the one linear polarization state that propagates in the WR-42 (within our bandwidth), corresponding to \vec{E} perpendicular to the long edge of the rectangular guide. Below we will illustrate why a signal injected via the third port of the circulator is reflected off of the decay cell assembly. The RF waveguide (WG) components of interest are, from $-\hat{z}$ to $+\hat{z}$; WR-42 rectangular WG, a rectangular to cylindrical WG converter, a quarter-wave plate, the decay cell (cylindrical WG), another quarter-wave plate, a cylindrical to rectangular WG converter, and another piece of WR-42 rectangular WG (see Figure 3.17 for a visualization of these components). The center five components are collectively referred to as the decay cell assembly.

Whether the dielectric lies in \hat{x} or \hat{y} as labeled in Figure 3.28 is a function of the main magnetic field direction. Here we assume the dielectric lies in the $\hat{y}\hat{z}$ - plane and that the incoming signal from the Vaunix is travelling in the $+\hat{z}$ direction. To analyze the signal

propagation, we will represent the polarization of the fields as a Jones vector and the quarter-wave plate as a matrix operating on said vector [3].

For frequencies in our bandwidth (17.9–18.1 GHz) the TE_{10} WG mode is excited within the rectangular WR-42. Since $E_z = 0$ in the TE modes of propagation we can represent the electric field as,

$$\vec{E}(t) = \begin{bmatrix} E_x(t) \\ E_y(t) \\ 0 \end{bmatrix} = \begin{bmatrix} E_{0x} e^{i(kz-\omega t+\phi_x)} \\ E_{0y} e^{i(kz-\omega t+\phi_y)} \\ 0 \end{bmatrix} = \begin{bmatrix} E_{0x} e^{i\phi_x} \\ E_{0y} e^{i\phi_y} \\ 0 \end{bmatrix} e^{i(kz-\omega t)}. \quad (3.18)$$

Note that the physical electric field is the real part of this vector. The Jones vector representing the above transverse wave is,

$$E = \begin{bmatrix} E_{0x} e^{i\phi_x} \\ E_{0y} e^{i\phi_y} \end{bmatrix}. \quad (3.19)$$

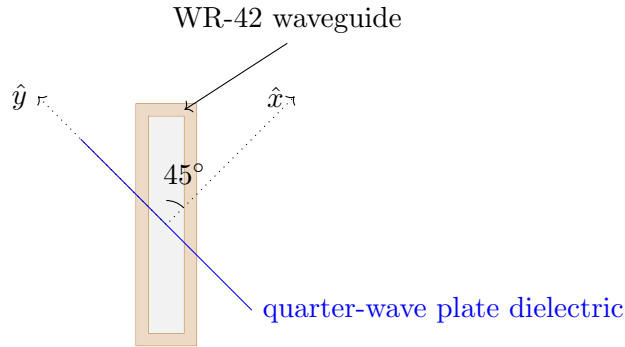


Figure 3.28: Axis definition relative to quarter wave plate dielectric and WR-42 waveguide.

The electric field associated with the Vaunix signal will be linearly polarized perpendicular to the long axis of the WR-42 WG, propagating via the TE_{10} mode. In the coordinates we defined above ($\hat{x} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, $\hat{y} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$),

$$E_{initial} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix}. \quad (3.20)$$

The dielectric plane defines the “slow” axis. Thus the quarter-wave plate’s matrix S operates on the polarization vectors as

$$S = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}. \quad (3.21)$$

Since the Vaunix signal will see two of the same quarter wave plates in series, what is most relevant to our calculation is S^2 ,

$$S^2 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \quad (3.22)$$

So after the second quarter wave plate the polarization of our signal will be,

$$E_{final} = S^2 \cdot E_{initial} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} 1 \\ -1 \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}. \quad (3.23)$$

This results in an electric field polarized parallel to the long axis of the WR-42 WG. The lowest order WR-42 mode that propagates radiation in that orientation is the TE_{01} mode which has a cutoff of: $f_{cut,01} = 34.7 GHz$. Since all of our signals are well below this frequency the radiation is reflected off of the cylindrical to rectangular WG converter back towards the decay cell. Where exactly this reflection takes place is a function of when the relevant dimension becomes too small to allow for the propagation of a $\approx 18 GHz$ signal, which is somewhere in the gradual taper of the converter from circular to rectangular. There is a π phase shift of the entire Jones vector upon reflection but the polarization axis of the radiation remains the same. Since this phase shift has no bearing on the polarization axis we will ignore it. Now the reflected signal travels back through the two quarter wave plates.

The slow axis of the quarter wave plate is unchanged by the reflection of the radiation so we can again apply S^2 :

$$E_{reflected} = S^2 \cdot E_{final} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix} = E_{initial}. \quad (3.24)$$

We see that the resulting polarization lies in the axis perpendicular to the long side of the WR-42 WG. This will propagate back towards the LNAs and the RF source via the WR-42 TE_{10} mode. We can also see that if the slow axis of the quarter wave plate was \hat{x} as opposed to \hat{y} , we would find the same results because $S^4 = I$ (the identity matrix) in either case.

With similar reasoning to above, if the two quarter wave plates are perpendicular to one another, meaning one has the slow axis in the $\hat{y}\hat{z}$ -plane and the other in the $\hat{x}\hat{z}$ -plane, then we would expect transmission through the two quarter wave plates. Note that we wouldn't have the quarter wave plates in this configuration during data taking. A signal originating on the I-side will propagate through the decay cell to the U-side amplifiers and likewise for the inverse. Here we let S' be the action of the quarter wave plate with the dielectric in the $\hat{x}\hat{z}$ -plane on the polarization vectors.

$$S' = \begin{bmatrix} 1 & 0 \\ 0 & -j \end{bmatrix}, \quad (3.25)$$

$$E_{final} = S \cdot S' \cdot E_{initial} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & -i \end{bmatrix} \begin{bmatrix} 1 \\ -1 \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ -1 \end{bmatrix} = E_{initial} \quad (3.26)$$

From Equation 3.26 we see that the initial polarization is unchanged by the action of two orthogonal quarter wave plates in series, regardless of which orientation of dielectric comes first ($S'S = SS' = I$), leading to transmission across the decay cell assembly.

Measurements were taken to validate the above understanding. The Keysight Field-Fox Handheld Microwave Analyzer (0 - 26.5 GHz) was used to measure the reflection and transmission from U to I side and visa-versa with the two different quarter wave plate configurations (parallel and orthogonal) over a 18-20 GHz bandwidth. Note that when this measurement was done we thought 18-20 GHz would be our DAQ's frequency bandwidth.

	QWPs Parallel (dB)		QWPs Orthogonal (dB)	
	U_{out}	I_{out}	U_{out}	I_{out}
U_{in}	-2.5	-25.7	-25.7	-2.5
I_{in}	-25.7	-2.5	-2.5	-25.7

Table 3.5: Summary of the test to investigate the reflection and transmission through the decay cell as a function of the quarter wave plate (QWP) configuration.

The results shown in Table 3.5 are the average over this frequency bandwidth. The setup used to take the measurements is shown in Figure 3.29. The results validate the above understanding in that when we expect reflection (parallel quarter wave plates), there is predominately reflection (more than two orders of magnitude more power reflected than transmitted), and likewise for when we expect transmission (orthogonal quarter wave plates). The symmetry of the results indicate that the U and I-side RF paths, despite being different lengths and one including a U-bend, have approximately the same loss. It is also surprising that the transmission loss is the same as the reflection loss as the transmitted signals travel through a quarter wave plate twice whereas the reflected signals travel through a quarter wave plate four times. This could indicate that the quarter wave plates are not a major source of RF loss.

3.9.4 Comparison of cascade analysis to measurement

3.9.4.1 Cascade analysis: assumptions and limitations

A cascade analysis of the He6-CRES apparatus was initially conducted by Jonathan Tedeschi (PNNL). The goal of the analysis is to estimate the noise and signal characteristics of our RF system by cascading the noise and signal power from the signal origin (decay cell), through all RF components out to the ambient receiver. It works on similar principles to the simple




Figure 3.29: RF setup used to measure the reflection and transmission through the decay cell assembly for both parallel and orthogonal quarter wave plate configurations for signals originating on the I-side. The locations of the coax to WG adapters on the circulators were reversed to measure signals originating on the U-side. We had the terminator in place of the RF twist when the measurement was actually taken, though we found the results didn't depend on the termination of the open circulator port.

RF model presented in Equation 3.17 but it accounts for all components of the RF chain. This enables us to have quantitative expectations of the noise power, signal power, and SNR input to the ROACH2 FPGA board. Whenever possible measured values of quantities were used, however with certain RF components all we had were the RF properties quoted from the manufacturer data-sheet.

The only values that have a significant impact on the SNR for which we need to rely on the data-sheet and not measurement are the LNA equivalent noise temperatures (Stage 1 and 2). The noise temperature for the LNAs is quoted for a 4 K ambient temperature. Since our LNAs are typically held around 5-7 K it's unclear exactly what equivalent noise temperature we should expect. We use 9 K for both the first and second stage LNA here as that is at the upper end of what the data-sheet quotes for a 4 K ambient temperature. This input adds the most to the level of uncertainty to the results of the cascade analysis. The components after the second stage LNA have very little impact on our SNR, as discussed in Subsection 3.9.2, and therefore uncertainties in their quoted equivalent noise temperatures are not very important. Table 3.6 provides a list of the import parameters used as inputs to the cascade analysis.

3.9.4.2 Network analyzer measurements

Measurements of the noise floor, expected SNR, and other important RF parameters were taken with the Keysight FieldFox Handheld Microwave Analyzer (0 - 26.5 GHz). The measurements were conducted by taking advantage of the reflection of the RF signal off of the two quarter wave plates situated on either side of the decay cell as discussed in Section 3.9.3. Signals were injected via the stainless steel SMA cables that connect to the third port of the circulator as discussed in Subsection 3.4.2.

Single frequency tones created by the Vaunix LMS-203 Lab Brick Signal Generator (10 - 20 GHz) were injected via the vacuum feedthroughs. These signals make their way through the stainless steel coaxial cable, attenuator, adapter, and circulator and travel through the

Parameter	Value
CRES Signal Power (dBm)	-120
Stage 1 LNA Equivalent Noise Temp. (K)	9
Stage 2 LNA Equivalent Noise Temp. (K)	9
Terminator Temperature (K)	28
Waveguide Temperature (K)	85
LNA Pair Total Gain (dB)	63
FFT Channel Bandwidth (kHz)	36.62
Waveguide Loss from source to first LNA (dB)	-1

Table 3.6: Important inputs to the cascade analysis.

RF waveguide system towards the decay cell. At the decay cell they are reflected by the two quarter wave plates that are configured orthogonally relative to one another and travel back towards the circulator, and then up into the LNA chain the way a CRES signal originating at the decay cell would. Measurements were conducted to determine how much power a signal lost between it's injection at the aluminum box feedthrough and the point of reflection in the decay cell. This value was then used to determine what power tone would be necessary to correspond to an effective CRES signal of -120 dBm (1 fW) originating at the decay cell. The loss from the point of injection to the decay cell was found to be -24 dB, which includes the 20 dB attenuator. So a signal of $-120 \text{ dBm} + 24 \text{ dB} = -96 \text{ dB}$ was injected at the feedthrough. There is roughly a 1-2 dB level of uncertainty on all of these measurements due to the frequency dependence of this 24 dB loss. The Vaunix Signal Generator can produce signals of minimum power -30 dBm, so a series of high-frequency coaxial attenuators were used to obtain the desired signal power and the power was verified with the FieldFox before injection.

Eleven measurements were taken for the U and I-sides respectively, spanning our current effective frequency bandwidth of 18.0 – 19.1 GHz in 100 MHz steps. For each measurement a tone was injected and a noise floor over the same bandwidth (with the tone off) was measured with the FieldFox. Each FieldFox measurement consists of 400 bins each of bandwidth 750 Hz. The noise was calculated as the average noise power over the 3 MHz total measurement bandwidth with the signal off and was scaled so as to compare to the frequency bin size of 36.6 kHz used in the cascade analysis. Note that this study was done using the bin size used for $^{83\text{m}}\text{Kr}$ CRES data so as to compare to the CRES SNR observed for $^{83\text{m}}\text{Kr}$ (the study was conducted before the observation of higher energy CRES events). The signal power was obtained by taking the maximum power over all bins. The FieldFox was returning the average over 1000 FFT outputs so the strength of these tones were constant in time. We also verified that the tone was always sharp enough to land in one bin and wasn't split into multiple bins. Figure 3.30 presents the signal power, noise power, noise density, and SNR measured for an effective CRES signal of -120 dBm originating at the decay.

3.9.5 Comparison

A comparison of the cascade analysis results and network analyzer measurements for important noise and signal parameters is presented in Table 3.7. In comparing the cascade analysis to measurement, we are attempting to make sense of the magnitude of our observed noise floor and SNR, not to fully characterize and understand their frequency dependence. We compare the results obtained from inputting the Table 3.6 values into the cascade analysis with the mean measured value obtained across our (18.0 - 19.1 GHz) bandwidth (one measurement per 100 MHz, totalling 11 measurements) for both the U and I-side as seen in Figure 3.30.

Given that both the cascade analysis and the FieldFox measurements have uncertainties on the order of 1-2 dB their results are consistent. This comparison gives us no reason to believe our model of the RF noise and gain is missing essential elements. The purpose of this comparison is to validate our cascade analysis model and to validate our understanding of

what RF components contribute most significantly to the total noise power and SNR. Since the results of the cascade analysis and of the simple model presented in Subsection 3.9.2 are effectively identical (to within $\approx .1\%$), this comparison lends validity to the simple model as well.

Parameter	Cascade Analysis	Measured (I-side)	Measured (U-side)
Full Receiver Gain (dB)	79.9	80.0	80.0
ADC Input Signal Power (dBm)	-40.1	-39.7	-39.6
ADC Input Noise Density (dBm/Hz)	-101.1	-99.7	-101.7
Integrated Noise Power to ADC (dBm)	-55.5	-54.1	-56.1
SNR for RMS thermal noise (dB)	15.3	14.4	16.5

Table 3.7: Comparison of the cascade analysis and the network analyzer measurements. The SNR quoted in the last line is for a -120 dBm (1 fW) signal.

3.9.6 Discussion of other relevant RF measurements

3.9.6.1 ^{83m}Kr mean SNR

The maximum power coupled into the TE11 mode of the decay cell expected for a 32 keV ^{83m}Kr conversion electron at 0.70 T is $\approx .84$ fW. This corresponds to a β perfectly centered in the cylindrical waveguide and with a pitch angle of 90° . The maximum SNR we observe for the 32 keV betas is ≈ 22 . Assuming linearity, this means that a -120 dBm (1 fW) signal would have an SNR of $\approx 26.8 = 14.3$ dB. This assumes that we have enough statistics to observe a CRES event with this maximum power condition (centered, pitch angle = 90°) which is reasonable given that this is the max SNR observed over more than 10,000 32 keV CRES events. This calculation is rough and doesn't account for the slope dependence of the SNR, though this should be a small ($\mathcal{O}(1)$) effect at the low slopes we observe with ^{83m}Kr .

Given this caveat, the calculation is still a meaningful sanity check, and is in good agreement with the expected SNR of a -120 dBm signal found from both measurement and the cascade analysis shown in the last row of Table 3.7.

3.9.6.2 Noise floor with terminator only

Lowering the power of our noise floor with all else equal leads to a direct improvement in the SNR of our CRES events. With this in mind a test was conducted where a cold (4.5 K) 50 Ω (impedance matched) terminator was coupled directly to both the I-side and U-side first stage LNAs. The configuration used for the test is pictured in Figure 3.31. This enabled us to locate the lower limit of the noise floor we could accomplish without changing our LNAs.

The I-side noise floor for the entire RF system as configured for data acquisition and the noise floor observed with only terminators coupled to the LNAs (as shown in Figure 3.31) is presented in Figure 3.32. The noise floors shown were constructed by averaging 1000 noise spectra and then doing a 100 point spline fit on that averaged spectrum. The U-side results look qualitatively similar.

The test indicates that the oscillatory behaviour of the noise floor (period ≈ 100 MHz) is a result of the waveguides themselves, not a feature introduced by the LNAs or any component further down the RF chain. This is reasonable given that a 100 MHz frequency corresponds to a 3m wavelength which is on the order of two times the length of the distance from the circulator to the decay cell where signals are reflected and back. Subsection 3.9.8 expands on this simplistic understanding.

The magnitude of the change in the noise power between the full RF system and the system with only a terminator varies from less than 0.5 dB at the troughs to 3 dB at the lowest frequency peaks. Using the terminator and decay cell temperatures relevant to this test as inputs to the cascade analysis one finds that the expected decrease in the noise power would be 0.6 dB (from -57.5 dBm to -58.1 dBm). This aligns well with observation, though doesn't attempt to account for the frequency dependence of the effect which via constructive and destructive interference is complicating the situation. It is important to

note that the decrease in the noise floor predicted is highly dependent on the presumed equivalent noise temperature of the LNAs which are not directly measured but are taken from the manufacturer (Low Noise Factory) data sheet.

One can also compare the absolute power expected in these two cases to the power observed by injecting a synthetic signal of known power directly into the ROACH2 and recording the corresponding ROACH2 power. This provides a rough calibration of the ROACH2 output in absolute units of power. The ROACH2's response does have a frequency response (≈ 1 dB effect) but one finds on average that 1 arb. ROACH2 unit = -63 dBm of input to the ROACH2 (for a 'requant gain' of 20). This can be used to see if the noise floors observed in Figure 3.32 are sensible in absolute units of power. If we assign the full RF system noise floor a value of 6 arb. ROACH2 units and the terminator only noise floor a value of 4 arb. ROACH2 units then we get -56.0 dBm and -57.0 dBm of input power to the ROACH2 compared to a cascade analysis prediction of -57.5 dBm and -58.1 dBm respectively. This can be considered good agreement given that both the absolute power calibration and the cascade analysis are only accurate to within a few dB (given the variability of cable losses and the frequency dependence we ignore). It is important to note that it is the SNR of our signals, which is essentially fixed by RF system up to the second stage LNA, and not the signal power input to the ADCs that matters to us. However, this comparison of absolute power is important in validating that the gain and loss accounting being done in the cascade analysis (from signal source all the way to the input to the Roach) is sensible.

3.9.6.3 Changing the waveguide temperatures

Another test was done to further our understanding of the noise floor where we changed the temperature of the decay cell and looked at how the noise floor was affected. The magnitude of the noise floor, as measured by the Roach, increased by 20 % (.8 dB) when the decay cell temperature increased from 52 K to 133 K. The decay cell was heated with a heater made of phosphor bronze wire wrapped around a small flat piece of copper that was adhered to the top of the decay cell's input gas port. The temperature of the decay cell was measured

by a silicon diode on the bottom of the decay cell's input gas port.

The cascade analysis, using the same inputs as Table 3.6 except for with the decay cell temperatures specified above, predicts a 1.5 dB increase in the noise floor when the decay cell temperature increases from 52 K to 133 K. In linear units this corresponds to a 40 % increase, compared to the 20 % observed increase.

The 0.7 dB discrepancy may be explained by the fact that the heater and silicon diode were on the same side of the cylindrical decay cell (attached to opposite sides of the gas input flange) and therefore the temperature reading would likely be higher than the average temperature over the entire copper body of the cell. This would mean that the 133 K value is likely an over estimate of the actual temperature of the decay cell. The 52 K measurement isn't plagued by this issue because the heater was not used to obtain 52 K; that is just the equilibrium temperature of the decay cell.

Additionally, it is our understanding that it is the temperature of the quarter wave plate and rectangular to circular adapter that matters more than the decay cell itself because that is where most of the loss is occurring. A more appropriate test would involve heating the entire decay cell assembly (including the decay cell, and both quarter wave plates and rectangular to circular adapters) and looking at the resultant change in the noise floor.

According to the cascade analysis the equivalent noise temperature goes from 51.5 K to 72.5 K when the decay cell temperature increases from 52 K to 133 K. According to Equation 3.16 this leads to a 1.5 dB predicted decrease in SNR.

3.9.7 Considerations to optimize SNR

What is the best SNR the He6-CRES experiment can expect to obtain given our current RF hardware? We can use the simple model presented above in Equation 3.17 to investigate this. Given the conservative values of a waveguide loss of 1 dB, a terminator temperature of 25 K, and a decay cell assembly of 80 K, the equivalent noise temperature one would expect for the entire RF system would be around 56.5 K. This would result in an SNR (for a -120 dBm signal) of $15.44 \text{ dB} = 34.99$. In the absolute best case (by cooling everything down

to the extent possible and getting rid of the photonic crystals) we could perhaps obtain a waveguide loss of 0.5 dB, a terminator temperature of 10 K, and a decay cell assembly of 45 K. In this case the equivalent noise temperature one would expect for the entire RF system would be around 25.6 K. This would result in an SNR of $18.88 \text{ dB} = 77.27$. This would entail over a factor of 2 increase in SNR. An improvement larger than this will likely not be possible without current RF hardware.

It's also interesting to consider what we could obtain with fundamentally better RF hardware. If we had no waveguide loss and all other components (terminator and LNAs) had a noise temperature of 4K we would obtain an SNR of $\approx 25 \text{ dB}$ for a -120 dBm signal which would represent roughly an order of magnitude improvement in SNR relative to our current system. This would drastically increase the number of events we see. It would also likely drastically complicate the nature of our detectable signals because so many more sidebands would be above the noise floor.

3.9.8 Reflection model of He6-CRES noise floor

The periodic peaks we observe in our noise floor can be explained by two reflections within the RF system. The phase of the propagating wave is given by

$$\phi(\nu) = z\beta(\nu), \quad (3.27)$$

where β is the frequency dependent phase constant.

There is some frequency ν_0 for which there is constructive interference between two possible paths (direct and reflected) to the amplifiers. The next frequency for which there is also constructive interference is obtained when there is one extra phase cycle on the longer (reflected) path. Therefore,

$$2\pi = \phi(\nu_0 + \Delta\nu) - \phi(\nu_0). \quad (3.28)$$

By Taylor expanding to first order around ν we obtain,

$$\Delta\nu = \frac{1}{\left.\frac{\partial\phi}{\partial\nu}\right|_{\nu_0}} \quad (3.29)$$

And since the length scale of interest is twice the distance between the reflections ($z = 2\ell$), then:

$$\Delta\nu = \frac{\left.\frac{\partial\nu}{\partial\beta}\right|_{\nu_0}}{2\ell} = \frac{c}{2\ell} \sqrt{1 - \left(\frac{f_c}{f}\right)^2}. \quad (3.30)$$

The relevant cutoff frequencies (f_c) for the He6-CRES apparatus are given by Equation 3.7 (circular waveguide) and Equation 3.8 (rectangular waveguide). As was discussed in Subsection 3.9.3, injected signals are relatively efficiently reflected off the decay cell assembly. This feature of the apparatus has the undesired effect of also reflecting any noise excitations generated within the RF system. A reflected signal is also partially reflected by the circulator (≈ 1.3 dB) according to the data sheet and measurements. The length scale of interest (ℓ) is the distance from the circulator to the far side of the decay cell (as was discussed in detail in Subsection 3.9.3), somewhere in the circular to WR-42 transition. The frequency scales of interest ($\Delta\nu$) are the distances between the peaks in the observed noise floors (see the noise floor for the I-side in Figure 3.32).

The RF system is composed of both circular and rectangular waveguide but since the vast majority of the length is rectangular ($\approx 90\%$) we do the calculations assuming propagation only in the TE_{10} mode of the rectangular WR-42 guide (cutoff 14.051 GHz).

Table 3.8 summarizes the results from modeling the oscillations in our RF noise floor as resulting from these two reflections within the RF system. There is a few cm uncertainty in ℓ caused by the ambiguity in exactly where the wave reflects off the waveguide assembly and there is a few MHz uncertainty in $\Delta\nu$ (expressed in the ranges given in Table 3.8) which arise from how you fit the noise floor peaks, which aren't perfectly periodic. Given these uncertainties, there is good agreement between the predicted length scales (ℓ) based on our simple two reflection model (Equation 3.30) and the physical length scales of the two sides of the RF system.

RF side	RF noise floor $\Delta\nu$ (MHz)	Physical length ℓ (m)	Reflection Model ℓ (m)
I-side	97 - 100	1.034	.975 - 1.005
U-side	64 - 71	1.441	1.373 - 1.524

Table 3.8: Comparison of the physical lengths between reflections to predicted lengths based on the $\Delta\nu$'s observed in the I and U-side noise floors. Predictions calculated using Equation 3.30. The range of $\Delta\nu$ values comes from ambiguity in how to determine the (not perfectly periodic) spacing between noise peaks. The physical lengths quoted are from the circulator to the far end of the decay cell assembly (the end of the circular to rectangular adapter). There are inherent $\mathcal{O}(1 \text{ cm})$ uncertainties in where the reflections are taking place within the adapter.

The large oscillations in the noise floor and therefore the SNR make analysis difficult in that long events are coming in and out of being detectable.⁵ Flattening the noise floor would greatly improve our event reconstruction so the above understanding of the origin of the oscillations in our noise floor is crucial to inform experimental redesigns that reduce RF reflections (see Chapter 8).

3.9.9 Conclusion of RF noise and SNR characteristics

The above serves to summarize our current understanding of the RF noise and SNR of the He6-CRES apparatus. The cascade analysis and the simple model expressed in Equation 3.17 agree in magnitude with measurement via the FieldFox Network Analyzer. Other tests, such as the terminator-only noise floor and the heated decay cell test can be made sense of with the cascade analysis. The reflection model sheds light on the oscillatory nature of our noise floor. A more thorough study of the RF noise and gain using HFSS or some other sophisticated

⁵Note that oscillations in the noise floor don't necessarily imply oscillations in the SNR but in the case of the He6-CRES apparatus the RF reflections that affect the noise floor also affect signals (and therefore SNR).

RF simulation software will be necessary to have a full handle on the frequency dependence of the SNR and noise floor.

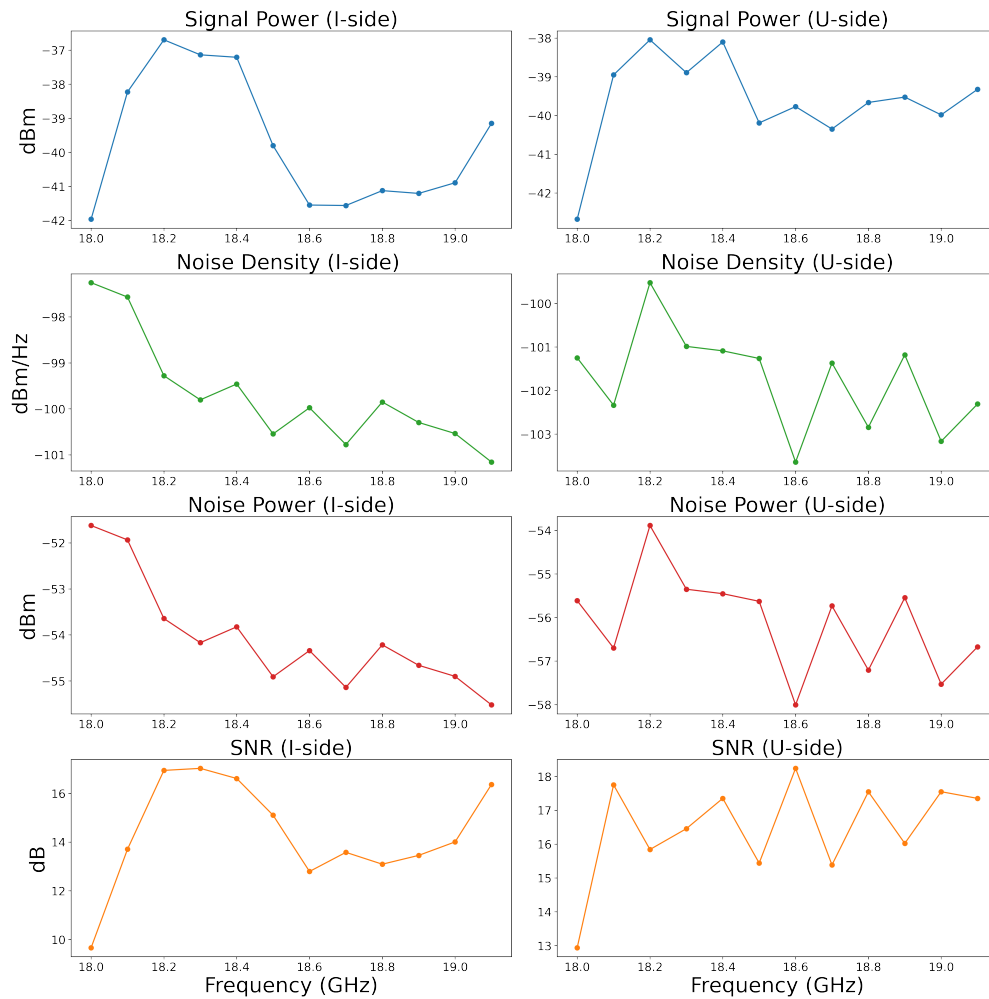


Figure 3.30: Results from the network analyzer measurements. Eleven measurements were taken over our current RF bandwidth, 18.0 - 19.1 GHz in 100 MHz steps.




Figure 3.31: RF configuration used to measure the RF noise floor without the entire RF waveguide assembly hooked up. This was to give us a limit on how much the noise floor could plausibly be reduced by making the waveguides colder and less lossy.




Figure 3.32: I-side noise floor for the full RF-system compared to noise floor with only terminators (at 4.5 K) coupled directly to the Low Noise Amplifiers (LNAs). Requart gain = 20 for both measurements.

Chapter 4

INITIAL OBSERVATIONS OF CRES EVENTS FROM $^{83\text{m}}\text{Kr}$, ^6He , AND ^{19}Ne

In this Chapter we present the initial detections of CRES events originating from $^{83\text{m}}\text{Kr}$, ^6He , and ^{19}Ne with the He6-CRES experimental apparatus. These initial observations represent the realization of the first two experimental benchmarks of *Phase 1* of the He6-CRES experiment articulated in Subsection 3.1.1. In Section 4.1 we present observations of CRES events originating from $^{83\text{m}}\text{Kr}$ and in Section 4.2 we present initial observations of CRES events originating from ^6He and ^{19}Ne . Given that initial observations and the development of our CRES event reconstruction algorithm and analysis procedure went hand-in-hand, we present initial findings here without a description of the analysis and present the analysis procedure in Chapter 6.

4.1 Initial observations of CRES events from $^{83\text{m}}\text{Kr}$

In March 2021 the He6-CRES experiment successfully detected its first CRES events becoming the second collaboration after Project 8 to build a functional CRES apparatus. These first events came from a bottled source of $^{83\text{m}}\text{Kr}$ which can be released into the decay cell region of the apparatus. The $^{83}\text{Rb} - ^{83\text{m}}\text{Kr}$ source was supplied by the DOE by the Isotope Program in the Office of Nuclear Physics and provides monoenergetic electron internal conversion lines of varying intensity in the 7 – 32 keV range [110].

Figure 4.1 illustrates simultaneous observation of the 7–32 keV lines via CRES, demonstrating our 1.1 GHz frequency bandwidth [15]. Figure 4.2 shows the effect of RF resonances on event track slope. As was discussed in Section 2.3, these resonances are caused by higher harmonic frequencies ($f_h = hf_c$) of the β current density coupling to higher-order waveguide




Figure 4.1: $^{83\text{m}}\text{Kr}$ data taken with $B = 0.68$ T. Center: full bandwidth *sparse* spectrogram showing dots for every point with $\text{SNR} > 6$, demonstrating the bandwidth (1.1 GHz) capability. Left: zoomed-in section exemplifying an event composed of multiple tracks. Right: histogram obtained after track identification and event clustering, showing simultaneous observation of the 7–32 keV lines.

modes. These higher harmonics can contribute significantly to the total radiated power and thus to the observed slope (see Equation 2.11).

4.2 Initial observations of CRES events from ^{19}Ne and ^6He

The He6-CRES collaboration first observed CRES events from both ^6He and ^{19}Ne in the Spring of 2022. These observations represent the first detection of high-energy β s (up to 2.1 MeV, corresponding to magnetic fields up to 3.25 T, surpassing the 31.9 keV M-line $^{83\text{m}}\text{Kr}$ conversion electrons observed by both the Project 8 [5] and He6-CRES collaborations) and the first detection of positrons via CRES¹.

Through these initial detections it was discovered that event features vary significantly with β energy. Most prominently, high energy β s result in long duration and highly sloped tracks [15]. Figure 4.3 shows ^6He and ^{19}Ne CRES events for fields 1.0, 2.0, and 3.0 T.

¹Note that ^{19}Ne is a β^+ decay.




Figure 4.2: Track slope vs track start frequency of $^{83\text{m}}\text{Kr}$ events demonstrating the RF resonances present in the detector. Note that since the low energy events associated with $^{83\text{m}}\text{Kr}$ have a very small frequency extent (due to having small slopes and being short in time) their track slopes are very sensitive to narrow slope resonances.

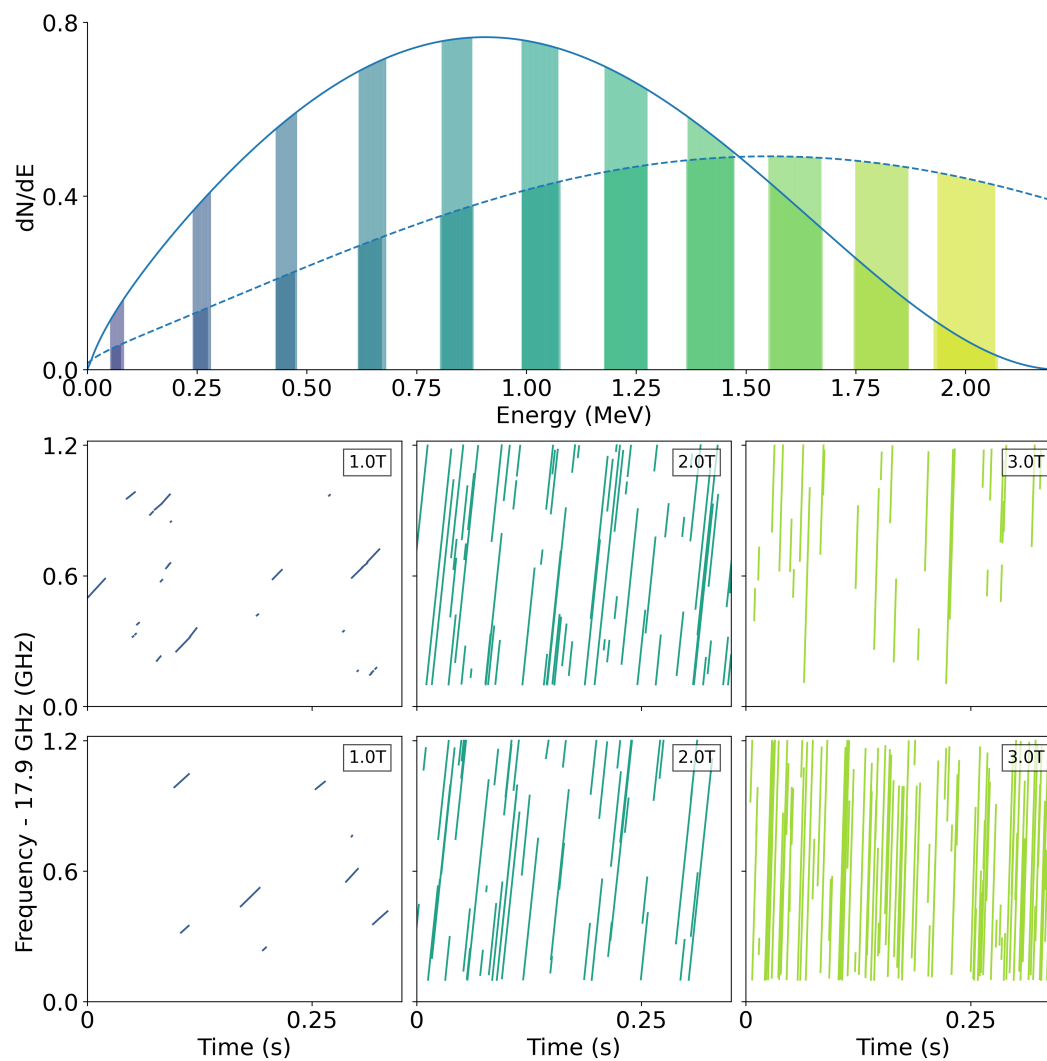


Figure 4.3: Top row: Standard model ^{19}Ne and ^6He β spectra showing the energy regions sampled by the 11 field settings from 0.75–3.25 T, given the 18–19.1 GHz RF bandwidth. Bottom two rows: Grid of identified tracks from ^{19}Ne (top) and ^6He (bottom) at 1, 2, 3 T. Each panel is an overlay of 0.35 seconds of data from 5 separate acquisitions. Track colors correspond to the magnetic field settings from the top row of the plot.




Figure 4.4: Density of CRES events with a constant trap field (left) and with trap slewing with $t_{trap} = 35$ ms (right) illustrating the need to slew the trap in order to accurately detect event start frequencies.

Track durations are limited by the time interval between scatters with the residual gas. Due to the decrease in the scattering cross section for β s on residual gas atoms with increasing β energy the mean track length grows by over a factor of 20 over the observed energy range from 10 keV – 2.1 MeV. Longer tracks make identifying the beginning of an event easier, since scattering can lead to the misidentification of the initial frequency. The low scattering probability also leads to accumulation of β s in the trap, as was discussed in Section 3.4. This results in a high density of long tracks crossing the spectrogram, interfering with the correct reconstruction of events born within our bandwidth. To remedy this, the trapping coils are toggled on and off with a trap-on period (t_{trap}) of 35 ms and trap off period of 20 ms, effectively emptying the trap. Figure 4.4 demonstrates the importance of continuously emptying the trap.

An additional difference observed between the tracks for high-energy β s compared to those from ^{83m}Kr is that the slopes are up to three orders of magnitude larger. The event slopes increasing with energy (and field) can be seen in the bottom two rows of Figure 4.3. The observed slopes vs field are aggregated in Figure 4.5.

In a waveguide, the power radiated into a mode λ at frequency f is given by [62, 88, 30]:

$$P_\lambda(f) \propto \left| \int q\vec{v}(t) \cdot \vec{E}_\lambda(\vec{r}(t))e^{-2\pi ift} dt \right|^2 \quad (4.1)$$

where $\vec{r}(t), \vec{v}(t)$ are the position and velocity vectors of the moving β , and $\vec{E}_\lambda(\vec{r})$ represents the mode electric field. The bandwidth (18–19.1 GHz) was selected to measure only in the TE_{1,1} mode, with cutoff at 15.2 GHz. The next mode, TM_{0,1}, has a cutoff at 19.9 GHz. Thus the track power observed is exclusively due to the power propagating in the TE_{1,1} mode. The track slope, df_c/dt , is a measure of the β 's total radiated power, as was discussed in Section 2.3 [15].

In addition to the TE_{1,1} mode, the β velocity vector couples to higher-order waveguide modes at higher harmonic frequencies ($f = nf_c$), which can contribute significantly to the total radiated power and thus to the observed slope (see Figure 4.5). However, since this power propagates into unobserved modes, we do not detect this power directly[15].

It was expected that some power would be present in higher harmonics of the cyclotron frequency but we were surprised to find that experimentally observed slopes scale in accordance with the expectation for radiation in free space, given by the relativistic Larmor formula [62]. This means that for high energy β s the amount of power in the 18 - 19.1 GHz range (observed directly as the track power and propagating via the TE_{1,1} mode) can be 10^{-3} or less of the total radiated power (observed indirectly as the slope). There may be the opportunity for future experiments to detect the significant power present at these higher harmonics in order to obtain better SNR and higher detection efficiencies.

While in general, the radiated power emitted in a waveguide can vary significantly from the free-space expectation, this result is predicted analytically for orbits centered around the waveguide axis [27] and is confirmed numerically using Equation 4.1, calculated up to 800 harmonics ($h = 800$), as shown in Figure 4.5.

When a mode (TE _{n,m}) cutoff frequency is crossed by a harmonic of the cyclotron frequency there is a momentary increase in radiated power that can be observed as a temporary

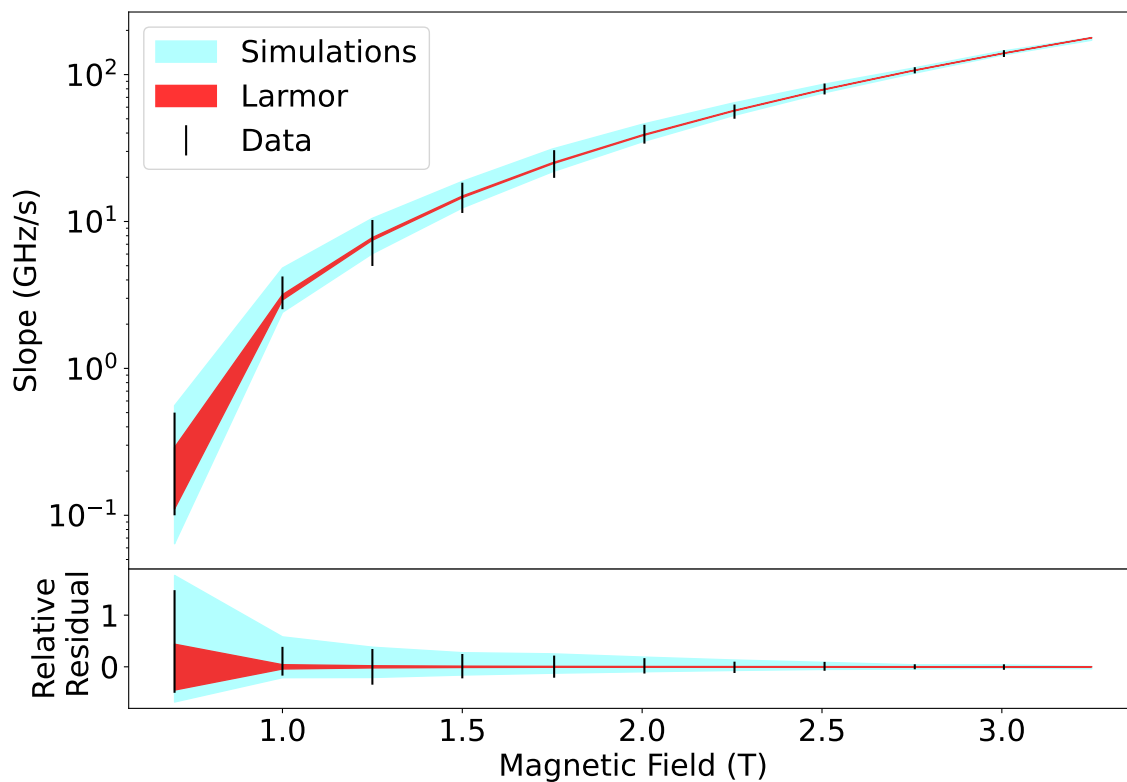


Figure 4.5: Observed track slopes, df_c/dt , in comparison to the expectation from free-space (red) and from waveguide simulations in a circular waveguide (blue) as a function of experimental magnetic field. Data bars and simulation bands correspond to maximum/minimum slopes observed in data and simulations, respectively, due to variations in the underlying beta energy and position. Fractional residuals are shown with respect to a 18.5 GHz beta, in free-space. Figure created by Nick Buzinsky of He6-CRES.

increase in the slope. For example, Figure 4.6 shows a spectrogram of a CRES event as it crosses the $\text{TE}_{2,2}$ mode cutoff for $h = 3$ at 18.45 GHz (the $\text{TE}_{2,2}$ cutoff frequency is $18.45 \times 3 = 55.35$ GHz for our cylindrical waveguide geometry). While there are technically a countably infinite number of resonances (n, m, h) within the 18–19.1 GHz RF bandwidth, the $(2, 2, 3)$ mode has a large (visible) effect on the slope due to its low order and therefore large coupling². The $^{83\text{m}}\text{Kr}$ slope resonances seen in Figure 4.2 are also explained by this effect. A paper describing the numerical simulations and math behind why the radiated power in a waveguide converges to the radiation in free space at high β energies (given by the relativistic Larmor formula [62]) is currently being written by Nick Buzinsky and RJ Taylor of He6-CRES.

A direct side effect of long, steep tracks is that there is a field-dependent maximum time-duration of events, t_{max} , which is given by,

$$t_{max} = \frac{f_{BW}}{df_c/dt}. \quad (4.2)$$

For *Phase 1* of the He6-CRES experiment the active frequency bandwidth, f_{BW} , was 1.1 GHz. The maximum event length t_{max} observed in Figure 4.7 (dotted line) is a function of Equation 4.2. Note that in Figure 4.7 we observe some event lengths that exceed 35 ms (t_{trap}). This is due to eddy currents allowing the trap to persist a few ms beyond 35 ms (see Chapter 6 for more discussion of this effect).

In summary, two primary features characterize CRES events from high-energy β s (a novel energy regime for CRES); long durations and steep slopes. The event reconstruction described below in Chapter 6 was designed to identify start frequencies across the ^{19}Ne β spectrum (up to 3.25 T or, equivalently, 2.1 MeV), and thus for a broad range of track slopes, durations, and signal-to-noise ratios.

²The indices n and m specify the RF mode and h specifies the harmonic of the cyclotron frequency.




Figure 4.6: CRES event crossing the $TE_{2,2}$ ($h = 3$) mode cutoff frequency in data.




Figure 4.7: Event length vs event slope from 11 fields: 0.75 - 3.25 T in 0.25 T steps. The maximum event length observed at a given slope is given by Equation 4.2 (dotted line).

Chapter 5

MONTE CARLO SIMULATION

In this Chapter we begin by leveraging the simplest simulation of a full ratio experiment in Section 5.1 in order to gain intuition for the ratio’s sensitivity to b_{Fierz} and to explore a few important experimental systematics. We also demonstrate an alternate way to present a ratio measurement in the energy domain. In Section 5.2 we look into the effect of events originating below the visible frequency bandwidth (18.0 – 19.1 GHz). Finally, in Section 5.3 we introduce our efforts to build a complete particle-level Monte Carlo that precisely simulates all components of the experiment (RF system, DAQ, and analysis pipeline), a prerequisite for pushing to high precision determinations of b_{Fierz} with CRES.

5.1 Simple Monte Carlo of ratio experiment

This relatively simple Monte Carlo takes as its arguments the different main field settings being used in the experiment, the frequency bandwidth of the apparatus, the assumed value for b_{Fierz} , the number of CRES counts obtained per isotope (N_{CRES}), the total number of counts per field obtained by the β monitor (see Section 3.7), and the normalization constant (C). We begin by calculating the integrated spectral density for each isotope, for each field (see Equation 3.5), using the fully corrected SM β decay spectrum presented in Subsection 1.2.3 [52, 51]. We then scale the integrated densities such that each isotope has N_{CRES} total counts and Poisson vary the obtained counts to simulate statistical fluctuations. Now we have an array of counts equal in length to the number of set fields being used in the experiment and summing to N_{CRES} (in expectation) for each isotope. We also Poisson vary the monitor counts obtained for each field and apply the ratio of the β monitor efficiencies (C) to the ^{19}Ne β monitor counts. Now that we have a simple simulated experiment we




Figure 5.1: Simple Monte Carlo experiment with b_{Fierz} set to zero for 10^4 counts per isotope measured over 11 field settings from 0.75–3.25 T in 0.25 T increments. Plot also shows $b_{Fierz} = \pm 1$ for comparison. Fit results: $b_{Fierz} = -0.065 \pm 0.060$.

follow the ratio analysis procedure explained in Section 3.1 (see Equation 3.4), obtaining a best fit to b_{Fierz} and C with a associated uncertainties.

This simple Monte Carlo lacks many details but it is sufficient to begin to understand what experimental factors most impact sensitivity to b_{Fierz} given our measurement technique. Figure 5.1 presents a simulated experiment.

We utilize the simple Monte Carlo to explore the sensitivity of the ratio measurement to b_{Fierz} in the absence of systematics in Subsection 5.1.1). In Subsection 5.1.2 we explore the ratio’s sensitivity to main magnetic field nonuniformities. In Subsection 5.1.3 we explore the ratio’s sensitivity to the β -monitor stability. In Section 5.1.4 we present an alternate way

of presenting a ratio experiment where we use one isotope to calculate efficiencies and then use this to correct the spectrum of the other isotope. In Subsection 5.1.5 we explore the way that finer binning in frequency further increases the robustness of the ratio to energy-dependent systematics. In Subsection 5.1.6 we apply the formalism developed in the previous two sections to the particular case of the wall-effect systematic.

5.1.1 Ratio measurement sensitivity to b_{Fierz}

An arbitrary experiment's sensitivity to a nuclear β decay parameter such as b_{Fierz} generically goes as $\frac{X}{\sqrt{N}}$ where N is the number of counts observed and the constant X is dependent on the specifics of the experimental approach, as described in [42]. The He6-CRES experiment has two observables with statistical fluctuations that impact sensitivity to b_{Fierz} ; the total number of CRES counts per isotope (N_{CRES}) and the total number of monitor counts per isotope (N_{mon}). We assume here that there is an equal amount of data taken at each field, meaning the relative number of CRES events (in expectation) observed at each field scales with the integrated spectral density of the isotope at that field¹ and frequency bandwidth and that the total number of monitor counts observed per isotope is (on average) constant for each field and (for simplicity) that the β monitor is equally efficient for ^{19}Ne and ^6He ($C = 1$).

Note that future iterations of the experiment could increase sensitivity to b_{Fierz} (for a fixed amount of acquired data) by strategically taking more data in the regions of the spectrum with the most impact on experimental sensitivity to b_{Fierz} . In this study we simulate a fixed amount of data acquired at each field as that was the strategy taken in the proof-of-principle measurements presented in this thesis (see Chapter 6). One may also naively think that since the Fierz interference term ($\propto m/E$) distorts the spectrum most at low energies (and therefore low fields) we should only take data at these low fields. This isn't correct, as the long lever arm of the measurement is what fixes the proportionality constant

¹This is technically incorrect as it doesn't account for field dependent CRES detection efficiencies but for the purposes of this study it is a sufficient assumption.

C (Equation 3.4) and enables sensitivity to the spectral distortion at low fields.

Since both the CRES detector and the β monitor detector are seeing the same density of radioactive gas, the total number of CRES events observed and the total number of monitor counts are related to one another by the efficiencies of these two detectors. We will address the impact of reasonable/physical detector efficiencies on our overall sensitivity to b_{Fierz} later in this section, but we begin with an exploration of our b_{Fierz} sensitivity without reference to the specifics of the detector CRES and β monitor efficiencies.

We define the sensitivity constant X_{CRES} (assuming that we are not statistics limited in the monitor, or $N_{mon} \rightarrow \infty$) as,

$$\sigma_{b_{Fierz}} = \frac{X_{CRES}}{\sqrt{N_{CRES}}}, \quad (5.1)$$

and we define the sensitivity constant X_{mon} (assuming that we are not statistics limited in CRES events, or $N_{CRES} \rightarrow \infty$) as,

$$\sigma_{b_{Fierz}} = \frac{X_{mon}}{\sqrt{N_{mon}}}. \quad (5.2)$$

The values of X_{CRES} and X_{mon} are dependent on the specifics of the experimental setup defined by the set fields used and the frequency bandwidth of the measurement. In this study we investigate two experimental setups described in Table 5.1. Experimental Setup 1 represents the best case set-up, where we measure over the entire energy spectrum of ^{19}Ne and are able to utilize the entire available frequency bandwidth. Experimental Setup 2 represents a worst-case where, due to significant scattering off of residual gas at low fields and β energies, we are only able to take high quality data starting at 1.75 T, and where a 300 MHz window at the bottom of the bandwidth must be used as a veto window for events rising from below. These two effects (scattering and events rising in from below the visible bandwidth) will be discussed in more detail in Chapter 6.

The value for X_{CRES} (X_{mon}) is obtained by simulating 100 trials (Monte Carlo experiments as described in the introduction to Section 5.1) with b_{Fierz} set to zero, for N_{CRES}

	Main Field Values	Frequency Bandwidth
Experimental Setup 1	11 fields, 0.75–3.25 in 0.25 T steps	18.0–19.1 GHz
Experimental Setup 2	7 fields, 1.75–3.25 in 0.25 T steps	18.3–19.1 GHz

Table 5.1: Description of the two experimental setups used to determine X_{CRES} and X_{mon} . Experimental Setup 1 represents the best case and Experimental Setup 2 represents a conservative worst-case setup. For both experimental setups, it is assumed that the same amount of data is taken at each field, as discussed in the text.

(N_{mon}) in the range 10^4 to 10^{11} in factor-of-ten steps, resulting in 800 simulated experiments. For each experiment N_{mon} (N_{CRES}) is set to infinity (effectively), to prevent any contribution in the uncertainty from that component. Each experiment returns a fit value along with an estimated uncertainty. An estimate of X_{CRES} (X_{mon}) is obtained for each simulated experiment by taking the fit uncertainty and multiplying it by $\sqrt{N_{CRES}}$ ($\sqrt{N_{mon}}$), following from Equation 5.1 (Equation 5.2). The results of the study are shown in Table 5.2.

Two sanity checks were conducted to ensure the sensibility of the study results. The first was verifying that for a given value of N_{CRES} (N_{mon}) the standard deviation of the obtained b_{Fierz} fit values matches the mean estimated uncertainty, indicating that the uncertainty quoted by the fit is unbiased. The second was verifying that the mean estimate for X_{CRES} (X_{mon}) was constant over the different values of N_{CRES} (N_{mon}), validating the claim that the sensitivity generically goes as $\frac{X}{\sqrt{N}}$ [42].

Now that we have these generic sensitivity constants X_{CRES} and X_{mon} , which make no physical assumptions about the differences between the CRES detector and the monitor detector we can consider the fact that the detection efficiency for the β monitor for our current experiment is much greater than the CRES detection efficiency. This is also likely to be true for any future iteration of our experiment or similar experiment. For the first attempt at the ratio measurement (Autumn 2022) we used the β monitor for ${}^6\text{He}$ data and

	X_{CRES}	X_{mon}
Experimental Setup 1	7.67 ± 0.10	4.03 ± 0.07
Experimental Setup 2	15.66 ± 0.48	13.35 ± 0.34

Table 5.2: Sensitivity constants X_{CRES} and X_{mon} for two different experimental setups. The two experimental setups used are described by Table 5.1.

found $\frac{N_{mon}}{N_{CRES}} \approx 15,500$. We used the β telescope for ^{19}Ne data and $\frac{N_{mon}}{N_{CRES}} \approx 4,500$ (further details found in Chapter 6). This means that the uncertainty in the experimental ratio σ_{exp} is overwhelmingly dominated by N_{CRES} regardless of which isotope or specific monitor is used, so in practice it is Equation 5.1 that can be used as a rule of thumb scaling of the ratio measurement's statistical sensitivity to b_{Fierz} .

This scaling is important to keep in consideration as He6-CRES works towards high precision determinations of b_{Fierz} . If we assume the best-case (Experimental Setup 1) we will need $N_{CRES} = (\frac{X_{CRES,1}}{10^{-3}})^2 \approx 5.9 * 10^7$ counts per isotope to obtain a $\sigma_{b_{Fierz}} \approx 10^{-3}$ statistical error. If we assume the worst-case (Experimental Setup 2) we will need $N_{CRES} = (\frac{X_{CRES,2}}{10^{-3}})^2 \approx 25 * 10^7$ counts per isotope to obtain a $\sigma_{b_{Fierz}} \approx 10^{-3}$ statistical error. The Monte Carlo study presented here can be used to understand the return on investment of crucial experimental design decisions such as attempting to increase the experimental bandwidth or whether we can get away with dropping low fields due to scattering effects. In addition, one could use this technique to determine the data taking procedures (how much data at each field) would optimize sensitivity to b_{Fierz} (minimize X_{CRES}).

5.1.2 Sensitivity to main magnetic field determination

Now we use the same simple Monte Carlo used above to investigate the impact of our main magnetic field determination/uniformity on our sensitivity to b_{Fierz} . In a ratio experiment it is crucial that the ^{19}Ne and ^6He data be taken at an identical set of main field values

so that the energy dependent efficiencies cancel and so that the predicted ratio (r_{SM} , see Equation 3.5) be calculated at the correct fields. In practice there are uncertainties in the determination of the main field values at the center of the decay cell as well as slight mismatches between the fields used for ^{19}Ne and ^6He . As was discussed in Section 3.6, for our current level of precision these effects are sufficiently small to be safely ignored but as we push to higher precision determinations of b_{Fierz} it will be crucial to understand how well the field needs to be known and how well matched the ^{19}Ne and ^6He fields need to be matched given the target precision of the measurement. As a reminder, the main field is monitored during data acquisition via an NMR probe that is ≈ 3 cm of the center of the decay cell (see Section 3.6 for details).

Two Monte Carlo studies were conducted to understand the field determination and uniformity effects. In the first, we investigate the effect of the true field being Gaussian distributed around the assumed field with a fixed relative standard deviation. In the second, we investigate the effect of the true field for one isotope being linearly offset from the assumed field, while for the other isotope the true and assumed field are the same. This is not the most physical set of assumptions but it sheds light on the limits of the field uncertainty and drift effects. In both studies we use Experimental Setup 1 as defined in Table 5.1 and use (effectively) infinite values for N_{CRES} and N_{mon} so that the effect of the field on b_{Fierz} determinations is isolated. In both studies we are interested in how fitting the simulated data using the assumed fields effects the determination of b_{Fierz} .

In the first study we take the assumed fields (11 fields, 0.75–3.25 in 0.25 T steps) and vary the i^{th} field B_i according to a Gaussian with mean B_i and standard deviation $B_i * \sigma_B$ so that the relative uncertainty at each field is given by σ_B . We simulate 40 trials each for 10 values of σ_B exponentially spaced from $\sigma_B = 10^{-5}$ to $\sigma_B = 10^{-2}$ for a total of 400 simulated experiments. The results of the study are summarized in Figure 5.2. As expected, the spread in B causes a spread in the fit values for b_{Fierz} centered around $b_{Fierz} = 0$. The size of the effect can be characterized by a single parameter $\frac{d\sigma_b}{d\sigma_B}$ obtained by dividing the standard deviation in b_{Fierz} fit values (σ_b) by the relative uncertainty in the field σ_B and




Figure 5.2: Sensitivity to main field determination error. Plot includes results from all three Monte Carlo studies done.

then averaging over all values of σ_B^2 . The value obtained for $\frac{d\sigma_b}{d\sigma_B}$ can be found in Table 5.3.

In the second study we take the assumed fields (11 fields, 0.75–3.25 in 0.25 T steps) and vary the i^{th} field B_i for ^{19}Ne only by a linear relative offset $B_i(1+\Delta B)$ (^6He fields unchanged). We simulate 10 values of ΔB exponentially spaced from $\Delta B = 10^{-5}$ to $\Delta B = 10^{-2}$ for a total of 40 simulated experiments³. The results of the study are summarized in Figure 5.2. The effect of this linear offset is to systematically affect the fit to b_{Fierz} and this effect can be characterized by a single parameter $\frac{db_{Fierz}}{d\Delta B}$ obtained by dividing the b_{Fierz} fit values (σ_b) by the linear offset in the field ΔB and then averaging. The value obtained for $\frac{db_{Fierz}}{d\Delta B}$ can be found in Table 5.3.

²This assumes the systematic is linear which holds true at least for these reasonably small values of σ_B .

³Since $N_{CRES}, N_{mon} \rightarrow \infty$ we don't need to run multiple trials here; there is no randomness.

	Brief description	Parameterization of systematic effect
Study 1	Gaussian spread	$\frac{d\sigma_b}{d\sigma_B} = 7.29 \pm 0.65$
Study 2	Linear offset for Neon	$\frac{db_{Fierz}}{d\Delta B} = -17.08 \pm 0.09$
Study 3	Linear offset for Helium	$\frac{db_{Fierz}}{d\Delta B} = 4.318 \pm 0.001$

Table 5.3: Parameterization of systematic effect associated with three different types of incorrect main field determinations.

The values $\frac{d\sigma_b}{d\sigma_B}$ and $\frac{db_{Fierz}}{d\Delta B}$ can be used to build a systematic error budget for a high precision measurement. In practice a real experiment may exhibit some linear combination of the two effects explored here (a Gaussian spread and a linear offset). The key point from study one is that the physical field values must be consistent with the field values assumed in the analysis to within 10^{-5} for a corresponding systematic uncertainty of $\approx 6.6 * 10^{-5}$ in b_{Fierz} . The key point from study two is that the field offset between the two fields must be to within 10^{-5} for a corresponding shift of $-12.75 * 10^{-5}$ in the fit relative to the true b_{Fierz} . As was discussed in Section 3.6, the field value is already known to within $\sim 10^{-4}$ (without field shimming) and similarly, the two fields match each other to within $\sim 10^{-4}$ (see Chapter 6). This means that with our current apparatus (without field shimming), we already have the main field determination systematics under control for a 10^{-3} measurement of b_{Fierz} .

5.1.3 Sensitivity to monitor instability

We discussed how the statistical uncertainty of the (assumed to be perfect) β monitor effects sensitivity to b_{Fierz} in Section 5.1.1. But what if the monitor's gain drifts, as is prone to happen in physical detectors, over the course of the data acquisition. In practice this drift is likely to be a random walk but in this study, for simplicity, we look at the effect of a linear drift parameterized by Δ_{mon} over the course of the acquisition (for both isotopes). We again use Experimental Setup 1 as defined in Table 5.1. We are assuming we are not statistics




Figure 5.3: Sensitivity of b_{Fierz} determination via a ratio experiment to linear β monitor drift.

Brief study description	Parameterization of systematic effect
Linear drift in β monitor	$\frac{db_{Fierz}}{d\Delta_{mon}} = -1.01 \pm 0.03$

Table 5.4: Parameterization of the sensitivity of a b_{Fierz} determination via a ratio experiment to linear β monitor drift.

limited but that the monitor rate drifts linearly from N_0 to $N_0(1 + \Delta_{mon})$ over the course of data acquisition for both isotopes which is assumed to start at 3.25 T and end at 0.75 T (as we do when taking data).

Results from the study are presented in Figure 5.3 and Table 5.4. This is an important systematic for He6-CRES going forward. In Section 3.7, the β monitor is shown to currently be stable to within 2 parts-per-thousand over 24 hours. For a 10^{-3} determination of b_{Fierz} one would need to limit the β monitor to drift 10^{-4} over the course of data taking to have the monitor drift systematic be well below the statistical error. Given that data taking for such an experiment may take over a week per isotope, ensuring the β monitor is sufficiently

stable for the desired level of sensitivity to b_{Fierz} presents a significant (but not impossible) experimental challenge.

5.1.4 An alternate way to present a ratio experiment

In this Subsection we present an alternate way to present the data obtained from a ratio experiment. In Section 3.1 the ratio is presented as a way to cancel energy-dependent systematics shared by both isotopes. An alternate (though ultimately equivalent) approach is to use one isotope to calculate the energy dependent efficiencies and then apply those efficiencies to the other isotope in order to obtain a “corrected” energy spectrum for one isotope. Below we present the math assuming $b_{Fierz} = 0$ so as to make the general applicability of this method apparent; this method can be used for precise β -energy spectrum determinations with any two isotopes in the energy regime in which they share non-zero spectral densities. At a facility capable of producing many isotopes of interest (like FRIB) one isotope with a high endpoint and well-known spectral shape (such as ${}^6\text{He}$) could be used to calculate efficiencies that could be applied to the precise determination of β^\pm spectra for many other isotopes [7].

The normalized rate for ${}^{19}\text{Ne}$ and ${}^6\text{He}$ observed at the i^{th} field is given by

$$A_{ne} \frac{N_{ne,cres,i}}{N_{ne,mon,i}} = \int_{E_{min,i}}^{E_{max,i}} W(E) {}_{19}\text{Ne} \epsilon_e(E) dE \approx \epsilon_{e,i} \int_{E_{min,i}}^{E_{max,i}} W(E) {}_{19}\text{Ne} dE \quad (5.3)$$

$$A_{he} \frac{N_{he,cres,i}}{N_{he,mon,i}} = \int_{E_{min,i}}^{E_{max,i}} W(E) {}_6\text{He} \epsilon_e(E) dE \approx \epsilon_{e,i} \int_{E_{min,i}}^{E_{max,i}} W(E) {}_{19}\text{Ne} dE \quad (5.4)$$

where we assume $\epsilon_{e^+}(E) = \epsilon_{e^-}(E) = \epsilon_e(E)$ ⁴ and A_{ne} and A_{he} are field-independent (in principle unknown) constants relating β -monitor rate to the density of radioactivity in the decay cell. Note that in general $A_{ne} \neq A_{he}$ due to the fact that the β -monitor has a different detection efficiency for positrons and electrons. By making the approximation that the CRES efficiency $\epsilon_e(E)$ is constant over the energy domain associated with the i^{th} field and frequency domain, we can pull this efficiency out of the integral ($\epsilon_e(E) \rightarrow \epsilon_{e,i}$). The limits

⁴Note that these efficiencies are energy, and therefore field dependent.

of this approximation are explored in Subsection 5.1.5. Now we can solve for the efficiency $\epsilon_{e,i}$ in Equation 5.4 and substitute it into Equation 5.3,

$$\begin{aligned} \int_{E_{min,i}}^{E_{max,i}} W(E)_{19Ne} dE &\approx \frac{A_{ne}}{\epsilon_{e,i}} \frac{N_{ne,cres,i}}{N_{ne,mon,i}} \\ &\approx \frac{A_{ne}}{A_{he}} \frac{N_{ne,cres,i}}{N_{ne,mon,i}} \frac{N_{he,mon,i}}{N_{ne,cres,i}} \int_{E_{min,i}}^{E_{max,i}} W(E)_{6He} dE. \end{aligned} \quad (5.5)$$

We notice that $\frac{N_{ne,cres,i}}{N_{ne,mon,i}} \frac{N_{he,mon,i}}{N_{ne,cres,i}}$ is the experimental ratio $r_{exp,i}$ at the i^{th} field defined in Equation 3.2 and that the ratio of the above normalization constants ($\frac{A_{ne}}{A_{he}}$) is the field-independent normalization constant C defined in Section 3.1. Thus, we can reduce the above to,

$$\int_{E_{min,i}}^{E_{max,i}} W(E)_{19Ne} dE \approx C r_{exp,i} \int_{E_{min,i}}^{E_{max,i}} W(E)_{6He} dE. \quad (5.6)$$

The above is essentially just a different presentation of the ratio algebra, but it makes clear that assuming the SM β spectrum for one isotope (${}^6\text{He}$) allows for producing an efficiency-corrected energy spectrum for the other isotope (${}^{19}\text{Ne}$).

5.1.5 Finer frequency binning to improve efficiency cancellation

The quality of the efficiency substitution in Equation 5.5 (or equivalently the cancellation of the efficiencies in Equation 3.5) is a function of the exact shape of that efficiency $\epsilon_e(E)$ in the interval $E_{min,i}$ to $E_{max,i}$. If the efficiency is varying wildly from $E_{min,i}$ to $E_{max,i}$, then the cancellation will be imperfect. If the efficiency is constant from $E_{min,i}$ to $E_{max,i}$, then the cancellation will be perfect. In principle a correction could be applied to the ratio making up for the imperfect cancellation but this partially defeats the point of the ratio which is to never need to have efficiencies very well characterized to obtain high precision determinations of β spectra. Below we present a way of greatly mitigating the imperfect cancellations of efficiencies; to utilize finer frequency binning in the construction of the ratio.

As was discussed in Section 3.1, the relevant integration window at a given field ($E_{min,i}$, $E_{max,i}$) is determined by using Equation 2.1 for main field B_i and the experimental frequency bandwidth. By breaking the full bandwidth up into n_{chunks} (presumably equal sized) pieces and conducting n_{chunks} distinct ratio analyses we greatly reduce the size of the integration windows (E_{min} , E_{max}) and therefore decrease the size of systematics associated with pulling the efficiency term ($\epsilon_{e,i}$) out of the spectral integrals. Each of the n_{chunks} ratio measurements obtained will be lower in statistics than the baseline case ($n_{chunks} = 1$) but due to the smaller integration windows there will be a more complete cancellation of efficiencies. Then one can take each of the best fits for b_{Fierz} (see Equation 3.4) and take the statistical uncertainty weighted-average of them to obtain a final measurement of b_{Fierz} .

Applying the finer frequency bins to the corrected energy spectrum formulation presented in Equation 5.7, we have for the i^{th} field and j^{th} frequency chunk,

$$\int_{E_{min,i,j}}^{E_{max,i,j}} W(E)_{19Ne} dE \approx C_j r_{exp,i,j} \int_{E_{min,i,j}}^{E_{max,i,j}} W(E)_{6He} dE. \quad (5.7)$$

So long as the frequency determination of events (δf) is small compared to the size of the frequency bandwidth chunks, we won't need to worry about line shape effects. The finer frequency bins drastically decreases the effect of imperfect efficiency cancellations, further eliminating the need to have efficiencies well-characterized in Monte Carlo. We will explore this claim quantitatively with a concrete example in the next Subsection.

5.1.6 Wall-effect efficiency

Here we will explore the formulations presented in Subsections 5.1.4 and 5.1.5 via the “wall-effect”. The wall-effect denotes the fact that the cyclotron radius (see Equation 2.4) is energy and field dependent and so the fiducial volume available to CRES electrons in the decay cell is energy and field dependent.

In principle the energy dependent efficiency $\epsilon_e(E)$ is the total efficiency, including all

systematic effects such as the frequency dependent SNR, the field dependent SNR, etc (full list covered at the end of Chapter 6). Here we investigate only the efficiency associated with the wall-effect so as to make things concrete and because it is a particularly dangerous systematic effect for He6-CRES because its effect is very similar in shape to that of b_{Fierz} .

The energy and field dependent wall-effect efficiency is obtained by considering that the center of the electron orbit can be at most $\rho_{max} = r_{WG} - r_{cycl}$ from the central axis of the decay cell (where ρ denotes the distance from the central axis of the decay cell to the guiding center of the cyclotron orbit), r_{WG} denotes the decay cell radius (≈ 5.78 mm), and r_{cycl} is the energy and main field dependent cyclotron radius given by Equation 2.4. If ρ exceeds ρ_{max} the electron will collide with the decay cell wall and lose all of its energy in $\mathcal{O}(1$ ns) and therefore be undetectable. Figure 5.4 provides a visualization of the effect. CRES electrons with smaller r_{cycl} have more fiducial volume available to them and will therefore be observed more frequently with all else equal.

Note that so far the generic efficiency has been presented as only having energy dependence ($\epsilon_e(E)_{we}$) but it can also depend explicitly on the main field B or the frequency f . These are omitted in the presentation of the ratio technique to reduce the busyness of the expressions. Mathematically, the wall-effect efficiency is given by,

$$\epsilon_e(E, B)_{we} = \left(1 - \frac{r_{cycl}(E, B)}{r_{WG}}\right)^2. \quad (5.8)$$

We consider a simulated ratio experiment consisting of main fields in the range 0.75–3.25 in 0.25 T steps, and a frequency bandwidth of 18.1–19.1 GHz (allowing for a realistic 100 MHz window for which to veto events from below), with $b_{Fierz} = 0$ assumed. This is identical to Experimental Setup 1 as defined by Table 5.1 except for the veto window between 18.0–18.1 GHz. The spectral densities, r_{cycl} values, and associated efficiencies involved in such a ratio measurement are shown in Figures 5.5, 5.6 and 5.7, respectively. From Figure 5.7 we can see that the efficiency changes drastically over the energies observed within a single main field value and between main fields. We also notice that the shape of the effect is roughly proportional to $\frac{1}{\gamma}$, as is our physics signal of interest: Fierz interference, parameterized by

b_{Fierz} .

In the absence of the ratio technique one would have a very hard time differentiating between b_{Fierz} and this effect. We would also be forced to know r_{WG} to high precision to accurately pin down this effect as the efficiency depends on it quadratically. The way the wall-effect alters the ^{19}Ne spectrum observed is almost identical to the way the Fierz interference term with $b_{Fierz} = +1$ alters the ^{19}Ne spectrum. Similarly, the way the wall-effect alters the ^6He spectrum is almost identical to the way the Fierz interference term with $b_{Fierz} = -1$ alters the ^6He spectrum. This means precision determinations of b_{Fierz} without the ratio technique would need to apply $\mathcal{O}(1)$ corrections to b_{Fierz} while in search of a $\mathcal{O}(10^{-3})$ (or smaller) effect. However, since the wall-effect is identical for the two isotopes and the effect of b_{Fierz} is opposite, this systematic effect (that would be devastating if one were attempting to pursue a b_{Fierz} measurement via CRES with a single isotope) is largely mitigated by the ratio technique.

Figure 5.8 illustrates the absolute value of the distortion to the ratio caused by the wall effect. Figure 5.8 (top) shows the distortion with the baseline analysis; $n_{chunks} = 1$ (see formalism introduced in Subsection 5.1.5). This illustrates the real power of the ratio measurement. The wall-effect efficiency varies by a factor of ≈ 2 (see Figure 5.7) and yet the distortion to the ratio is less than 1% over the energy range of interest.

Figure 5.8 (bottom) shows the distortion(s) with finer frequency binning; $n_{chunks} = 5$. The distortions with finer binning is dependent on the specific frequency chunk but is universally almost 2 orders of magnitude less than with the distortion of the baseline analysis of using a single frequency bin. This concretely illustrates the power of the finer binning. It allows for the energy dependent efficiencies to cancel much better than in the baseline case due to the smaller integration windows (E_{min}, E_{max}) as was discussed in Subsection 5.1.5.

Figures 5.9 and 5.10 demonstrate the application of the alternate presentation of a ratio experiment described in Subsection 5.1.4 in the context of the wall-effect. In this alternate formulation or presentation we still fit to C and b_{Fierz} using Equation 3.6 but then we present a corrected ^{19}Ne spectrum using Equation 5.7 instead of the ratio itself. The Monte Carlo




Figure 5.4: Wall effect visualization. The wall-effect efficiency presented in Equation 5.8 is the ratio of the area available to allowed decays (shown in blue) to the total available area defined by the waveguide inner walls.

simulated data is created via the prescription presented at the beginning of this Section but now the spectral integrals include the wall-effect (Equation 5.8). In simulation b_{Fierz} is taken to be zero. To show meaningful residuals we assume $N_{CRES} = 10^6$ counts per isotope but we don't Poisson vary the simulated counts for ease of interpretability.

Figure 5.9 (top) shows the uncorrected ^{19}Ne spectra obtained from a simulated experiment. We see that the wall-effect causes a huge distortion to the spectrum relative to the true probability density function (pdf). As mentioned above a naive fit to b_{Fierz} here results in $b_{Fierz} \approx 1$. Figure 5.9 (bottom) shows the corrected spectrum with the baseline analysis ($n_{chunks} = 1$). Figure 5.10 shows corrected spectra with finer frequency binning: $n_{chunks} = 2, 5$ in the top and bottom sub-figures, respectively.

Table 5.5 shows the size of the wall-effect systematic (δb_{Fierz}) as a function of measure-




Figure 5.5: ^{19}Ne and ^6He spectral densities observed given a simulated ratio experiment consisting of main fields in the range 0.75–3.25 in 0.25 T steps, and a frequency bandwidth of 18.1–19.1 GHz.

ment technique. With 100 MHz frequency bins ($n_{chunks} = 10$) the size of the distortion to b_{Fierz} is well below 10^{-4} . This means that (in principle) a measurement using this technique would not need to worry about this systematic until measurements at the 10^{-4} level. Here we only address the wall-effect in detail but the cancellation of many systematics should follow this general trend, as most systematics don't vary as much as this one over a given field. Note that utilizing 100 MHz frequency bins is only possible if your frequency resolution is small compared to this bin size. The limits of this assumption are discussed in Chapter 6.

Also note that applying corrections to b_{Fierz} isn't that simple as many of our systematics are intermingled. For example, in reality, the wall-effect interacts with the SNR dependent detection efficiency: electrons with smaller ρ couple more strongly to the TE_{11} circular waveguide mode and therefore produce higher power, more detectable events. To know exactly how this folds in you need an absolute power calibration and to simulate the detection efficiency through our entire analysis pipeline. The power of the binned ratio technique




Figure 5.6: Field and energy dependent r_{cycl} values observed given a simulated ratio experiment consisting of main fields in the range 0.75–3.25 in 0.25 T steps, and a frequency bandwidth of 18.1–19.1 GHz.

is that we can (potentially) make very precise determinations of b_{Fierz} without precisely characterizing every systematic because of the powerful cancellations illustrated here.

5.2 Monte Carlo including CRES events born below the detection window

An initially unanticipated side effect of high energy CRES events having very large slopes is that during the time the trap is on ($t_{trap} = 35$ ms for our *Phase 1* preliminary ratio measurement presented in Chapter 6) events with initial cyclotron frequencies from below the active detection window (18.0 – 19.1 GHz in *Phase 1* of the He6-CRES experiment) rise into being detectable. At 3.25 T event slopes are greater than 100 GHz/s. This means that at 3.25 T the 1.1 GHz experimental bandwidth is traversed in ≈ 11 ms. It also implies that for trap-on duration (t_{trap}) of 35 ms CRES electrons born with as low a cyclotron frequency as ≈ 14.5 GHz can rise into the visible bandwidth during the trap-on period.

Determining b_{Fierz} via a ratio measurement as described by Equation 3.6 relies on only




Figure 5.7: Wall-effect efficiencies at play during a ratio measurement given a simulated ratio experiment consisting of main fields in the range 0.75–3.25 in 0.25 T steps, and a frequency bandwidth of 18.1–19.1 GHz.

counting CRES events that originate exclusively in the active detection window (18.0 – 19.1 GHz). This means successfully vetoing all CRES events that arise from below. Due to the oscillations in the SNR response vetoing all events from below is difficult (see Chapter 6 for more details). The oscillating SNR response is an incredibly important systematic effect for the He6-CRES experiment because if CRES events are broken and mistakenly identified as two events, the argument that systematics cancel perfectly in the ratio due to the narrow line-shape of the detector is broken. Here we present a simple Monte Carlo simulation that provides knowledge of the density of events coming up from below and sheds light on the effect of imperfect vetoing of these events.

Figure 5.11 illustrates the “extended bandwidth” for ^{19}Ne and ^6He . The energy-width of the 18.0 – 19.1 GHz frequency bandwidth for magnetic field values .75 T - 3.25 T in 0.5 T steps is shown via a shaded bar plot⁵. The energy-width of the extended bandwidth

⁵We use 0.5 T steps instead of 0.25 T steps in order to reduce the busyness of the plot.

Description of Measurement Technique	n_{chunks}	$\delta b_{Fierz} * 10^3$
No efficiency cancellation using ratio	N/A	1000
Ratio measurement	1	3.2 ± 1.0
Ratio measurement	2	1.19 ± 0.38
Ratio measurement	5	0.374 ± 0.096
Ratio measurement	10	0.170 ± 0.034

Table 5.5: Distortion to measured b_{Fierz} caused by the wall-effect with different measurement approaches given a simulated experiment consisting of main fields in the range 0.75–3.25 in 0.25 T steps, and a frequency bandwidth of 18.1–19.1 GHz.

(red outline) illustrates the energies in which CRES events could possibly rise into being visible for $t_{trap} = 35$ ms. The key issue is that the area under the curve within this extended bandwidth for a given field is drastically different for ^{19}Ne and ^6He . At 3.25 T (the farthest right portion of Figure 5.11, there will be drastically more events coming from below for ^6He as for ^{19}Ne . Concretely, this means that imperfect vetoing is not a systematic that will cancel in the ratio because it will not effect the two isotopes in equal proportion.

Figure 5.11 also demonstrates that the density of events coming from below is marginal at low fields due to the relatively small CRES event slopes at low fields. As has been demonstrated above, it is the low-field regions that are most sensitive to b_{Fierz} . However, it is the high-field regions that fix the normalization constant C in the fit (see Equation 3.6) so systematics that primarily effect high-field/high-energy regions of the ratio are still crucial to understand and characterize.

The extended bandwidths presented in Figure 5.11 provide the farthest reaches in energy in which an event could possibly end up being detected in the visible bandwidth. For example for ^6He an event born at the endpoint could end up in the visible bandwidth at 3.25 T. However, this event would need to be born right as the trap is turned on to have the time

to traverse such a large distance in frequency⁶. We generate β 's within the extended energy regions presented in Figure 5.11 according to the SM spectral density and assign them an event start time uniformly within the trap on period. We then propagate the events to their frequency location at the end of the trap on period (t_{trap}).

We use the Monte Carlo described above to calculate the percent of all detectable events that originated below 18.0 GHz for $t_{trap} = 35$ ms (top) and $t_{trap} = 10$ ms (bottom) (see Figure 5.12). By detectable we mean an event is either born above 18.0 GHz or that it rises above 18.0 GHz during its lifetime (ending when the trap turns off). We see that for high fields and a trap on period, t_{trap} , of 35 ms up to 70% of detectable ${}^6\text{He}$ events are coming from below. This implies that one needs to successfully veto 2 out of 3 ${}^6\text{He}$ events at these fields. In principle, it sounds simple to veto such events: simply cut all events with a start frequency of 18.0 GHz. However, as is discussed in detail in Chapter 6, due to SNR oscillations which cause broken tracks this is not a simple task. The $t_{trap} = 10$ ms plot (bottom of Figure 5.12) is included to illustrate that the issue is improved but not at all fully mitigated with a decrease in the trap-on time⁷.

Figure 5.13 shows the relative difference (%) between the naive ratio (see Equation 3.6) and the ratio obtained when including all detectable events (including those from below that rise into being detectable) for $t_{trap} = 35$ ms and $t_{trap} = 10$ ms. This implies that vetoing 90% of all events from below will result in a $\approx 7\%$ deviation from the naive (expected) ratio at 3.25 T for $t_{trap} = 35$ ms. Intuitively, this is a function of many more events arising from below for ${}^6\text{He}$ than for ${}^{19}\text{Ne}$ due to the larger density of higher energy events for ${}^6\text{He}$. The Monte Carlo presented here will be used extensively in Chapter 6 to make sense of *Phase 1* ratio data.

⁶Technically the slope is a function of both main field and frequency so it is changing as the CRES event loses energy to radiation. This is accounted for in the Monte Carlo.

⁷Note that the trap on period cannot be made arbitrarily small as it does take finite time ($\sim 1 - 5$ ms, but it is field/slope dependent) to clearly differentiate CRES events from RF noise fluctuations.

5.3 Particle-level Monte Carlo

The Monte Carlo simulations of Section 5.1 and Section 5.2 are very useful for gaining initial insights into the b_{Fierz} sensitivity and the most important systematics of a ratio measurement. However, as the He6-CRES collaboration works towards a 10^{-3} determination of b_{Fierz} a more robust Monte Carlo simulation will be necessary to develop a comprehensive and defensible systematic error budget. As was demonstrated above, the ratio method is powerful in its ability to cancel certain types of spectral distortions, such as the wall-effect. But there are systematics that do not cancel well in the ratio such as differing scattering environments between ${}^6\text{He}$ and ${}^{19}\text{Ne}$ data acquisitions and events arising from below.

We have created a python-based particle-level Monte Carlo simulation package (*he6-cres-spec-sims*), which simulates the CRES dynamics of individual CRES electrons in our magnetic trap geometry, calculates the corresponding time domain signals, simulates the action of the DAQ, and outputs `spec` files that can be input into our data analysis pipeline (described in Chapter 6). The analytic formulas used to calculate the CRES features of the β such as cyclotron frequency, average magnetic field experienced, maximum z position, axial frequency, and modulation index from initial conditions (energy, position in decay cell, pitch angle) have been validated against Kassiopeia (<https://github.com/KATRIN-Experiment/Kassiopeia>). The derivation of the analytic formulas used and comparison to Kassiopeia are presented in Appendix D. This Monte Carlo is still being calibrated and validated and so hasn't yet been used to analyze any systematic effects in detail. Figure 5.14 and Figure 5.15 illustrate some of the current capabilities of the simulation package.





(a) Wall-effect distortion of ratio for $n_{chunks} = 1$.(b) Wall-effect distortion of ratio for $n_{chunks} = 5$.

Figure 5.8: Distortion to ratio vs field for different binning schemes ($n_{chunks} = 1, 5$). Note that the cusps are caused by the absolute value crossing zero and that in order to present a smooth line a 0.25 T window was swept continuously over the fields of interest. Note that in the y-axis label ($abs\left(\frac{R_{we}}{R_0} - 1\right)$) R_{we} denotes the ratio including the wall effect efficiency and R_0 denotes the ratio obtained without any efficiency included.




(a) Simulated wall-effect distortion of ^{19}Ne spectrum. No ratio measurement correction done here.




(b) Simulated wall-effect distortion of ^{19}Ne ratio-corrected spectrum for $n_{\text{chunks}} = 1$.

Figure 5.9: The top plot shows how the wall-effect effects drastically effects the spectral shape of ^{19}Ne in the absence of the ratio technique. The bottom plot shows how the ratio-corrected does have some minor distortion relative to the true pdf.



(a) Simulated wall-effect distortion of ^{19}Ne ratio-corrected spectrum for $n_{chunks} = 2$.



(b) Simulated wall-effect distortion of ^{19}Ne ratio-corrected spectrum for $n_{chunks} = 5$.

Figure 5.10: Simulated wall-effect distortion of ^{19}Ne ratio-corrected spectrum for $n_{chunks} = 2, 5$ demonstrating the power of finer binning for cancelling efficiencies.

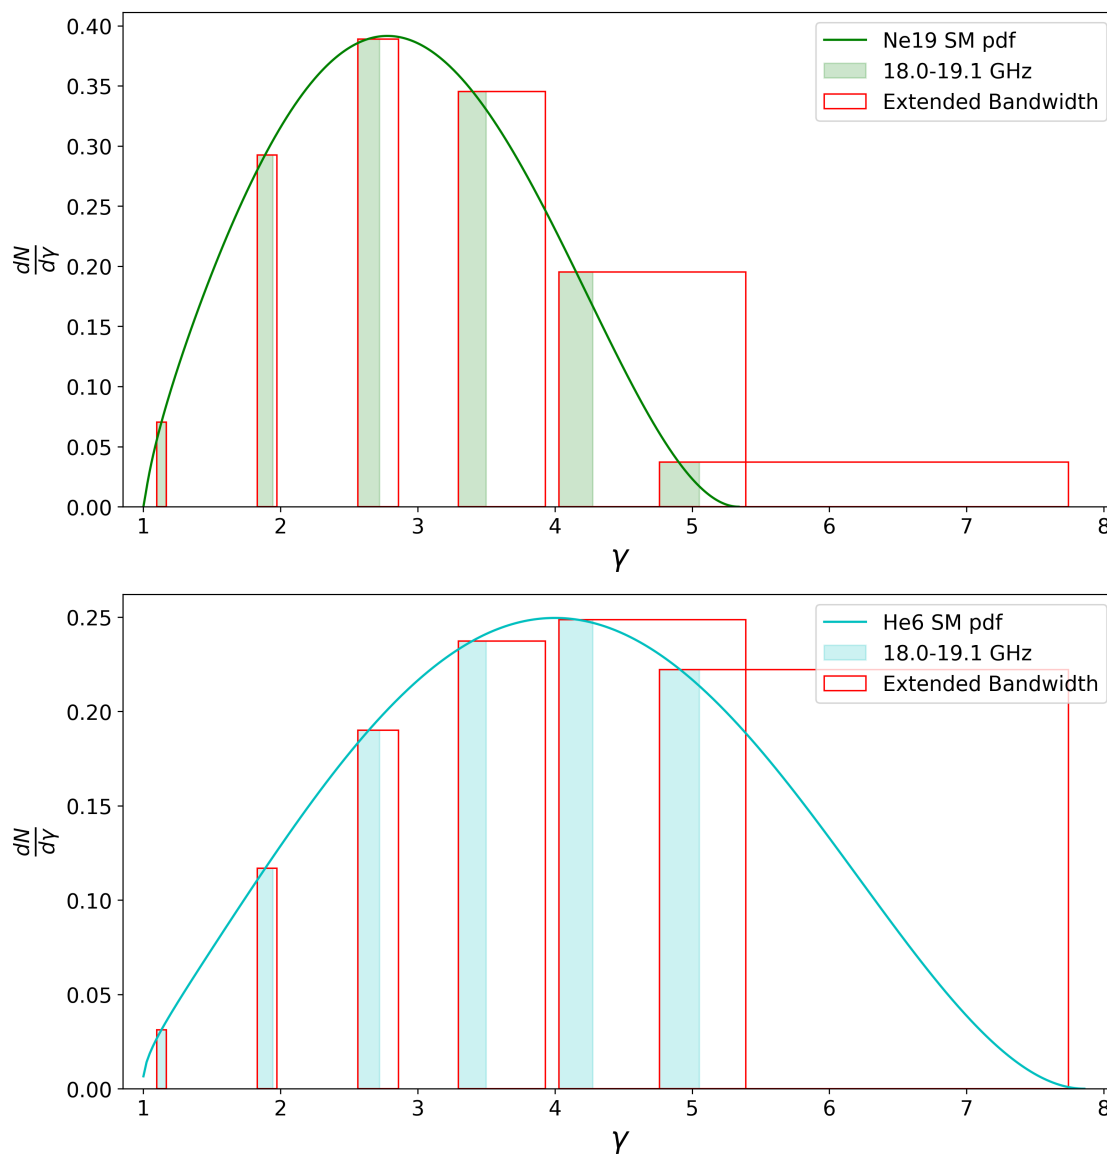


Figure 5.11: ^{19}Ne (top) and ^6He (bottom) SM spectral density (solid line) with the energy-width of the 18.0 – 19.1 GHz frequency bandwidth for magnetic field values .75 T - 3.25 T in 0.5 T steps (shaded). The energy-width of the extended bandwidth (red outline) illustrates the energies in which CRES events could possibly rise into being visible for $t_{\text{trap}} = 35$ ms.




Figure 5.12: Percent of all detectable events (an event either born above 18.0 GHz or that rises above 18.0 GHz during its lifetime) that originated below 18.0 GHz for $t_{trap} = 35$ ms (top) and $t_{trap} = 10$ ms (bottom).




Figure 5.13: The relative difference (%) between the naive ratio (see Equation 3.6) and the ratio obtained when including all detectable events (including those from below that rise into being detectable) for $t_{trap} = 35$ ms and $t_{trap} = 10$ ms.




Figure 5.14: Example output of the *he6-cres-spec-sims* package (for demonstration). Here we see that first order sidebands ($\text{band_num} = \pm 1$) gain power as the maximum z position (and therefore the modulation index) of an event increases.




Figure 5.15: Sparse spectrogram (SNR threshold of 3) from a simulated *spec* file. The simulated SNR response is oscillating with frequency and the simulated event lasts from 1 to 8 ms, has a 125 MHz axial frequency, and three visible tracks associated with it.

Chapter 6

DATA ANALYSIS AND PRELIMINARY RATIO MEASUREMENT

In this Chapter we begin with a description of the event reconstruction procedure in Section 6.1. Event reconstruction refers to the process of going from a raw spectrogram (`spec` file) to a set of distinct events and associated event properties (such as event SNR, start/stop time/frequency, etc.). In Section 6.2 we describe the generic procedure of taking a set of reconstructed events originating from both ^{19}Ne and ^6He , building the ratio from that, and comparing that ratio to expectation. In Section 6.3 we present an analysis of a full (preliminary) ratio measurement taken at CENPA in the Fall of 2022. Finally, in Section 6.4 we discuss the most important systematics to address as He6-CRES works towards a high precision measurement of b_{Fierz} .

6.1 *Event reconstruction*

To robustly reconstruct events our data analysis involves (1) applying a signal-to-noise-ratio (SNR) cut to the raw spectrogram, (2) clustering above-threshold points into track segments, and (3) clustering track segments into events. `Katydid`, a modular C++-based analysis framework developed by the Project 8 collaboration, has been adapted for the specific event reconstruction needs of He6-CRES [81]. A python-based post-processing does one final clustering of colinear tracks. The quality of event reconstruction was verified using both visual inspection and comparison of expected and observed distributions of reconstructed event properties.

6.1.1 *Katydid*

Katydid utilizes processors that are connected via a `yaml` configuration file that specifies the signal/slot relationships between the processors. Each processor receives data inputs (signals), does some data processing, and then passes data outputs to the next processor (slots). He6-CRES forked the original Project 8 Katydid repository (<https://github.com/project8/katydid>) in order to adapt it to the specific needs of our experiment.

Project 8 wrote time-domain data to disk whereas He6-CRES writes frequency domain data to disk (as `spec` files) so a processor (`spec-processor`) that takes frequency domain data as an input was constructed. The Katydid configuration file that we used for the ratio measurement presented below took two `spec` files as input: one noise file, and one data file containing CRES events to be reconstructed. A `gain-variation` processor averages the frequency spectra of the noise file and fits the resulting mean colored noise curve to a spline. This spline function is then used by the `discriminator` processor to threshold the data. This is done according to a pre-configured threshold SNR value which is multiplied by the frequency-dependent mean noise to obtain the threshold: typically set to around 9, and referred to as the ‘Katydid Fourier bin SNR threshold’. This results in a (relatively) uniform density of above-threshold points. The above-threshold points are then grouped into track segments by the `sequential-track-finder`. The `sequential-track-finder` processes the `spec` file slice by slice, taking each above-threshold point and seeing if another above-threshold point sits in line with it (to within some pre-configured time/frequency error) according to the pre-configured initial slope and continues in this way (recalculating the slope with each additional point) to build up track segment objects. The `sequential-track-finder` processor is configured to only find track segments between 3 and 5 points long. This is to overcome the fact that tracks are curved. Then these track segments are clustered via a `dbscan` processor that utilizes a C++ implementation of the DBSCAN clustering algorithm [31]¹. The clustering radius (in frequency and time) is pre-configured and is scaled according to the

¹The Katydid DBSCAN implementation was written by Noah Oblath of He6-CRES and Project 8.

field-dependent slope of the CRES events such that the clustering area is kept constant. The output of this final processor are long tracks (not yet considered events)².

Note that the analysis procedure described here is optimized for long tracks with small jumps due to scattering with residual gas or curvatures due to resonances. This was to optimize our analysis's performance over the largest swath of CRES energies without needing to implement a unique analysis strategies in each energy regime. Future iterations of the experiment will likely want to implement a Katydid analysis similar to the low-energy, high scattering analysis utilized by Project 8 for the low-field regions.

As was described in Section 3.8, when an acquisition of CRES data is initiated by the DAQ user, a specified number of `spec` files with the exact same configuration (main field, trap current, etc.) are acquired. The group of `spec` files (often 100 per acquisition) are assigned a single `run_id` and unique `file_ids`. The slow controls data (main field, trap current, etc.) associated with each `run_id` are written to the `he6cres_db` (housed on `rocks`, the CENPA computing cluster). Every night `spec` files are `rsynced` from the DAQ computer to `rocks`. Katydid is run on a large set of nodes within `rocks` and the field dependent configuration parameters such as the initial slope in the `sequential-track-finder` processor and the DBSCAN radius in the `dbscan` processor are calculated automatically by looking up the field associated with the `run_id` in the `he6cres_db` and using that field to calculate the (approximate) track slope via the via the free-space Larmor formula [62].

Katydid on `rocks` outputs a set of `root` files (one per `spec` file) organized by `run_id`. Each of these `root` files contains important track-level properties such as SNR, track start time/frequency, track end time/frequency, slope, etc.

6.1.2 Python post-processing

A python-based post-processing procedure is then responsible for aggregating the results from an entire experiment (defined loosely as a set of `run_ids`, typically consisting of a single

²Heather Harrington is the main Katydid developer within He6-CRES.

isotope covering a range of different fields), and applying a final clustering and cleaning.

The Katydid `dbscan` processor clusters based purely on the proximity of tracks in frequency and time which leads to track segments that are not colinear being mistakenly clustered into a track object. These ‘barnacle’ track segments are cut by dropping any track segment whose time intercept is more than 2 sigma away from the mean time intercept of the set of track segments that make up the track object (an example of a barnacle cut from an event is seen in the top row of Figure 6.2).

After this cleaning procedure, a final clustering is done on the time intercept of the track objects. The SNR response of the detector is oscillatory: causing low SNR events to come in and out of detection. This final layer of clustering serves to connect colinear tracks that are far from one another in frequency and time due to these SNR oscillations³. The clustering radius (‘eps’) is chosen based on histograms of the time deltas between track time intercepts. There is a steep rise in counts observed close to zero caused by the broken and colinear tracks and the field-dependent clustering radius is selected as the knee of this histogram.

Figure 6.1 demonstrates the main steps of the analysis procedure applied to ¹⁹Ne CRES events observed at 1.75 T. Figure 6.2 which features a main band and two first order sidebands illustrates the important role the Katydid Fourier bin SNR threshold plays in the resultant set of detected events⁴. The top row of Figure 6.2 also demonstrates the important role the layer of DBSCAN clustering of colinear tracks plays, in that one can see a 200 MHz gap between track segments in the event labeled 11. Without this final clustering layer this event would be mistakenly identified as two independent events.

Figure 6.2 also illustrates the difficulty of determining whether an event originated below the visible bandwidth. Because the SNR response peaks around 600 MHz (see Figure 6.5)

³Since the SNR oscillations have a period of ≈ 100 MHz a clustering based on proximity in frequency and time would lead to severe over-clustering.

⁴Note that our analysis procedure isn’t currently capable of grouping the three bands of Figure 6.2 into one event. We treat each band as a unique event object, despite all three bands originating from one CRES electron. Since the number of detectable sidebands scales linearly with the event density, the effect should cancel in the ratio.


a low SNR event may not be detectable until towards the middle of the bandwidth. This effect is discussed at length in Section 6.3.

Finally, the python post-processing aggregates the events and all calculated features such as event SNR, start/stop time/frequency, etc. across an entire experiment (typically consisting of many files across many different fields for a single isotope) and writes them to one `events.csv` file on `rocks` so that further analysis can easily be conducted on local machines. Beyond the event-wise features the monitor rates and field values associated with each `spec` file in the experiment are written to disk (`files.csv`). The field value and β monitor rate are written to the `he6cres_db` every 1-3 seconds during an experiment. The post-processing uses the time the `spec` file was created to find the main field value and β monitor rate closest in time to that acquisition (usually within 2 seconds). Thus each row of the `files.csv` contains the crucial environmental data associated with each second of acquired data. This environmental data is crucial to the building of the ratio discussed below in Section 6.2.


6.2 Building the ratio

A ratio measurement loosely consists of taking both ^{19}Ne and ^6He CRES data at a variety of different fields. For *Phase 1* we have decided to use the main field values of 0.75 – 3.25 T in 0.25 T steps, for a total of 11 fields. Once data is taken at these 11 fields for both isotopes and one has run the `Katydid` and python post-processing of the data on `rocks` one constructs the experimental ratio as was described in Section 3.1. Here we outline how this is done explicitly for completeness.


Once one has the `events.csv`, containing the event-wise features, and the `files.csv`, containing the file-wise environmental data, one simply sums up the events observed at each field and sum up the monitor rate observed during each second of data at each field. One then takes the ratio of these to obtain a source-intensity normalized CRES event rate per field per isotope. The field-wise ratio of the normalized ^{19}Ne rate and the normalized ^6He rate is the experimental ratio ($r_{exp,i}$, see Equation 3.2). The statistical uncertainty is calculated



(a) Sparse spectrogram with Katydid Fourier bin SNR threshold = 9.



(b) Track segments (3–5 points in length) identified by the `sequential-track-finder`.



(c) Fully reconstructed events after the two layers of DBSCAN clustering (see text).

Figure 6.1: Demonstration of the three main event reconstruction steps for a set of ^{19}Ne CRES events observed at 1.75 T in October 2022 (`run_id = 532`, `run_id = 2`).

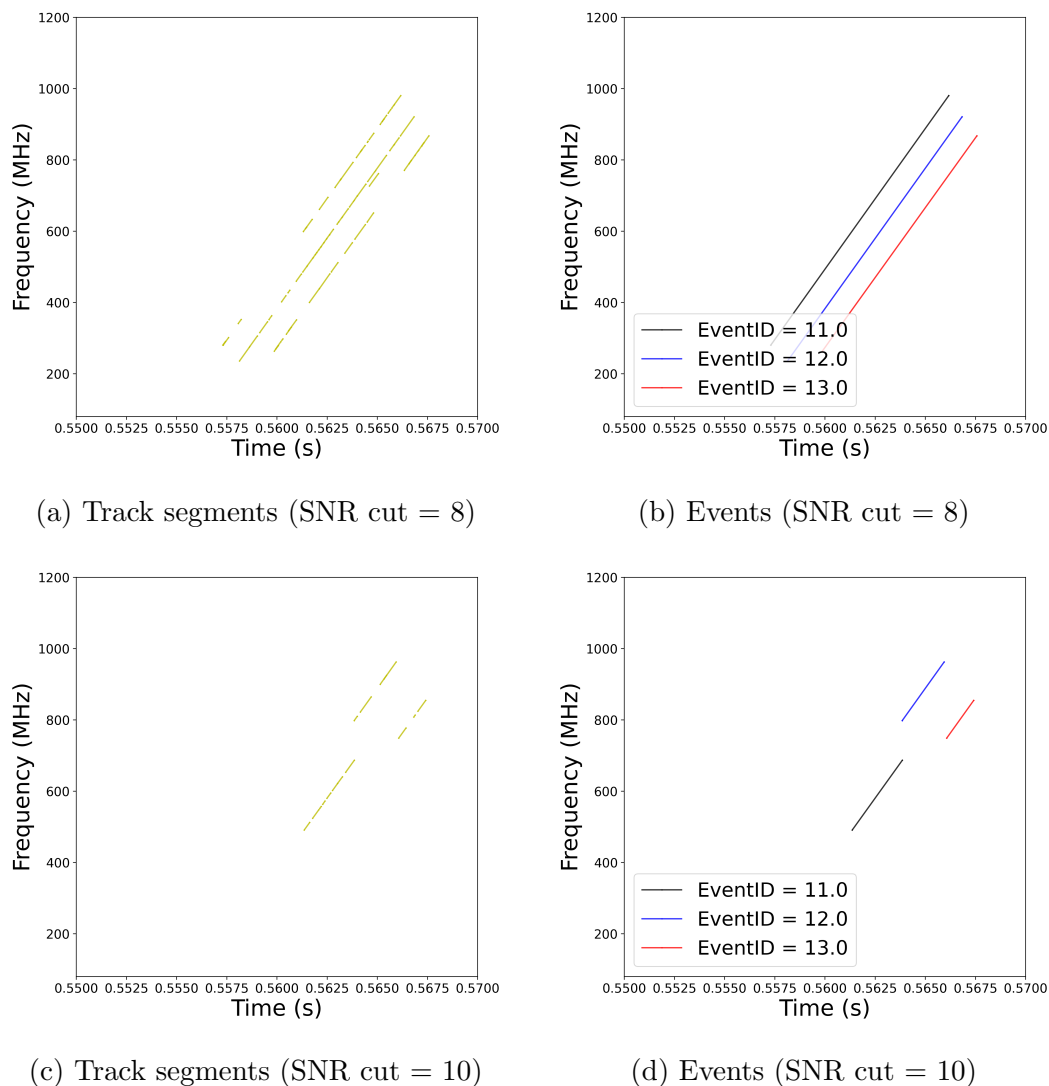


Figure 6.2: Demonstration of the effect of the Katydid Fourier bin SNR threshold (SNR cut) on event reconstruction. The spectrogram contains a main band and two first order sidebands originating from ${}^6\text{He}$ observed at 2.5 T in August 2022 (`run_id = 393`, `run_id = 0`). One can see that the choice of the SNR threshold has a large effect on the resulting set of reconstructed events.

by propagating a \sqrt{N} uncertainty in the four observables (CRES rate and monitor rate for ^{19}Ne and ^6He) through to the ratio (see Equation 3.3).

Then once the experimental ratio ($r_{exp,i}$) and its uncertainty ($\sigma_{exp,i}$) are calculated one can compare to expectation (or obtain a fit to b_{Fierz}) via Equation 3.4. There is some additional complexity to the comparison between observation and expectation that arises from whether or not one attempts to fully veto events originating from below the visible bandwidth. This will be discussed in detail in Section 6.3 below.

6.3 Preliminary ratio measurement

In this Section we present the preliminary ratio measurement taken in Autumn 2022. Subsection 6.3.1 will present the details of the experimental process and Subsection 6.3.2 will present the analysis of the ratio measurement.

6.3.1 Ratio measurement experimental process

In Autumn 2022 500 seconds per field of CRES data for both ^6He and ^{19}Ne were taken at main field values of 0.75 – 3.25 T in 0.25 T steps⁵ with a fixed relative trap depth of 2.55 parts-per-thousand. This trap depth was chosen such that at 3.25 T, we were applying 1.800 A to the trapping coil, the maximum safe current recommended for the 26 AWG gauge wire of the trap coil (this corresponds to 0.415 A at 0.75 T).

The data written to disk sums to ≈ 90 minutes of CRES data per isotope and the data for each isotope was taken in continuous ~ 36 hour campaigns separated by approximately six weeks. The campaigns were long relative to the amount of data written to disk for two primary reasons. First, it took approximately 40 minutes for the main field to come to equilibrium when switching the field value despite leaving the superconducting solenoid in driven mode. Second, there were vacuum issues that needed to be addressed continuously throughout data acquisition. These vacuum issues were particularly tricky with the ^{19}Ne run

⁵We started at high fields and worked down towards 0.75 T because changing fields was faster when ramping the field down in magnitude.

due to the fact that the ^{19}Ne gas system (see Subsection 3.3.2) contains significant amounts of N_2 that gradually saturates the Cryo-II and eventually (after 1-2 hours) overpowers the decay cell with N_2 . In order to prevent this build-up of residual gas we had to turn on the Cryo-II heater between each field in order to burn off and pump the N_2 that had accumulated.

A key challenge for a high precision ratio measurement is ensuring that the experimental conditions are identical (to within the desired uncertainty) between the two isotopes. Figure 6.3 shows the main field error in parts per million (ppm)⁶. The error bars illustrate the standard deviation of field values observed during the time period data was acquired at the given field. Note that the differences between the set field and true field are primarily due to uncertainties in the relationship between the current and resulting field in the superconducting solenoid. Agreement could be improved by iteratively using the observed field to adjust the set current in the magnet but this wasn't necessary for the desired level of precision. The current field matching between ^{19}Ne and ^6He field values would result in a $\approx 2 \times 10^{-3}$ systematic error on a determination of b_{Fierz} (see the field uniformity Monte Carlo study presented in Section 5.1.2). This is more than enough field uniformity for the current proof of principal. Another order of magnitude in field uniformity would be necessary to get to a 10^{-3} determination of b_{Fierz} which would be obtainable with field shimming. However, shimming would make data taking even more time consuming. Another option would be to incorporate the true field values as measured into the Monte Carlo expectation of the ratio.

Due to very high background levels of β^\pm annihilation gamma radiation present during the ^{19}Ne data run we had to use the β telescope as opposed to the SIPM-based β monitor (as was used for the ^6He run) for source-intensity normalizations. The telescope utilizes a coincidence between two different detector regions which cuts down drastically on background counts. As was shown in Section 3.7, the β telescope is stable to the percent level over 24 hours, which is sufficient for this proof-of-principle demonstration of the ratio technique. Background suppression measures are being implemented to ensure that higher precision

⁶Here 0 means the true field is exactly equal to the set field




Figure 6.3: Main Field vs Main Field Error (ppm). Error bars are the standard deviation of field values (in ppm) observed during the time period data was acquired at the given field.

measurements will be able to utilize the β monitor for source-intensity normalizations. The current β monitor is already stable to 2 parts-per-thousand over 24 hours and work is being done to ensure it is even more stable over a longer time period in the future.

An unfortunate consequence of waiting 6 weeks between isotope runs was that the cryocooler which cools the He6-CRES apparatus (see Section 3.4) was shut off and brought up to room temperature for routine vacuum maintenance. When the apparatus was cooled back down it was discovered that the noise floor was slightly different, as is illustrated in Figure 6.4. This caused the SNR response of the detector to be slightly different between the ${}^6\text{He}$ and ${}^{19}\text{Ne}$ data sets. Figure 6.5 shows the mean SNR of the track segments detected for both ${}^6\text{He}$ and ${}^{19}\text{Ne}$ CRES data during the Autumn 2022 data campaign. Despite the distinct noise floors, the SNR response is very similar in shape, with the peaks in ${}^6\text{He}$ SNRs slightly (5-7%) higher than those for ${}^{19}\text{Ne}$. This slight difference in SNR response could effect the detection efficiency and therefore systematically effect the ratio measurement. Simulating




Figure 6.4: Autumn 2022 noise floors for ${}^6\text{He}$ and ${}^{19}\text{Ne}$ runs. The noise floors presented are the spline fits of the mean noise power obtained by the `gain-variation` Katydid processor.

the effect of these slight differences in SNR is beyond the scope of this work given that future iterations of the experiment will be able to trivially mitigate this effect by not turning off the cryocooler between isotope runs. We attempt to get a handle on this potential systematic via running our analysis with a wide range of Katydid Fourier bin SNR thresholds (see Subsection 6.3.2 below). Note that a constant difference in detection efficiency between ${}^{19}\text{Ne}$ and ${}^6\text{He}$ (meaning at each field there was 10% higher detection efficiency for ${}^{19}\text{Ne}$ than ${}^6\text{He}$ for example) will not effect the viability of the ratio measurement as it will be accounted for by the normalization constant C (see Equation 3.6). It is only a field-wise difference in detection efficiency that would cause a systematic effect as these differences would not be dealt with by the normalization constant C .




Figure 6.5: Track segment SNRs from Autumn 2022 ${}^6\text{He}$ and ${}^{19}\text{Ne}$ data runs. The transparency for the ${}^6\text{He}$ points is turned down so that one can compare the shapes of the responses.

6.3.2 *Comparing observed and predicted ratios*

In the following sections we present two different approaches to analyzing the results of the ratio experiment. In Subsection 6.3.2.1 we present an analysis that involves the least number of cuts to data. We compare the total CRES rate observed, including events originating from below the visible frequency bandwidth to the expected rate as found via Monte Carlo simulations that include events originating below the visible frequency bandwidth (see Section 5.2). By “from below” we mean they are born with a cyclotron frequency less than 18 GHz and lose energy to radiation, rising into being detectable before the trap cycle ends.

In Subsection 6.3.2.2 we present an analysis that involves (attempting to) cut all events coming from below the visible bandwidth. This type of analysis proves more difficult but has the advantage that one can easily interpret the ratio results in the energy domain as was discussed in Subsection 5.1.4.

6.3.2.1 *Including events from below*

The simplest approach to analyzing a ratio experiment is to make no cuts to the events output by the python post-processing procedure described above in Section 6.1. This means that all events, even those that originate below the visible frequency bandwidth (below 18.0 GHz) and rise into being detectable before the trap is emptied, are included in the analysis. This approach has two major advantages. First, the number of events born below the visible bandwidth outnumber events born within the visible bandwidth by a factor of up to 2.5 (it’s field and isotope dependent, see Figure 5.13 of Section 5.2). This means that vetoing all events from below results in a very large reduction in statistics. Second, vetoing all events from below is difficult due to the occasionally broken nature of events. It is not easy (even by eye) to determine if the event illustrated in Figure 6.2 originated from below or not. Since low SNR events are more likely to be broken and therefore misclassified as originating in the visible bandwidth, one needs to cut low SNR events to successfully veto all events from below, even further reducing event counts and the statistical power of the measurement. An

attempt at an analysis that cuts all events from below is shown in Subsection 6.3.2.2.

Ratio data taken at 0.75 T and 1.0 T was excluded from the analysis in this Subsection as well as the analysis presented in Subsection 6.3.2.2. The reason for this is that scattering systematics are not cancelled in the spectral ratio at these fields. The residual gas environments, as monitored by a residual gas analyzer (Stanford Research Systems RGA-100), differed between the ^{19}Ne and ^6He data runs. The scattering cross section with ^2H and N_2 , as well as other prominent residual gases, rolls off with energy, meaning that low-field (and therefore low-energy) data is most effected by the differing residual gas environments. The mean event lengths differed greatly between ^{19}Ne and ^6He data at 0.75 T and 1.0 T. At 0.75 T the mean event length for ^{19}Ne was 3.0 ms whereas for ^6He it was 5.7 ms. At 1.0 T the mean event length for ^{19}Ne was 7.9 ms whereas for ^6He it was 12.5 ms. These observations are consistent with higher residual gas pressures measured during ^{19}Ne data acquisition. At fields higher than 1.0 T (1.25 - 3.25 T in .25 T steps), the mean event lengths are within ≈ 1 ms of one another. Figure 6.6 demonstrates that the field-wise track length distributions for ^{19}Ne and ^6He are notably different at 0.75 T and 1.0 T.

As was discussed in Section 6.1 the Katydid analysis conducted on this data was not optimized to deal with low-field CRES data in which tracks are short and events are often comprised of many tracks due to scattering-induced jumps in frequency. In principal, one can address scattering systematics at low fields by either having a robust enough analysis that event reconstruction is immune to the presence of scattering or by ensuring experimentally that the residual gas pressures are identical. Fortunately, at higher fields the scattering systematics are greatly suppressed due to the decrease in scattering cross sections. For the proof-of-principal analysis presented here, we opt to simply exclude these two lowest fields where the scattering environments are resulting in different event length distributions. Future phases of the He6-CRES experiment will need to pin down this systematic more quantitatively through both experimental and Monte Carlo simulation efforts.

The number of events detected per isotope at a given field is a function of the Katydid threshold used in the analysis (see Figure 6.2). The effect of this threshold should, in

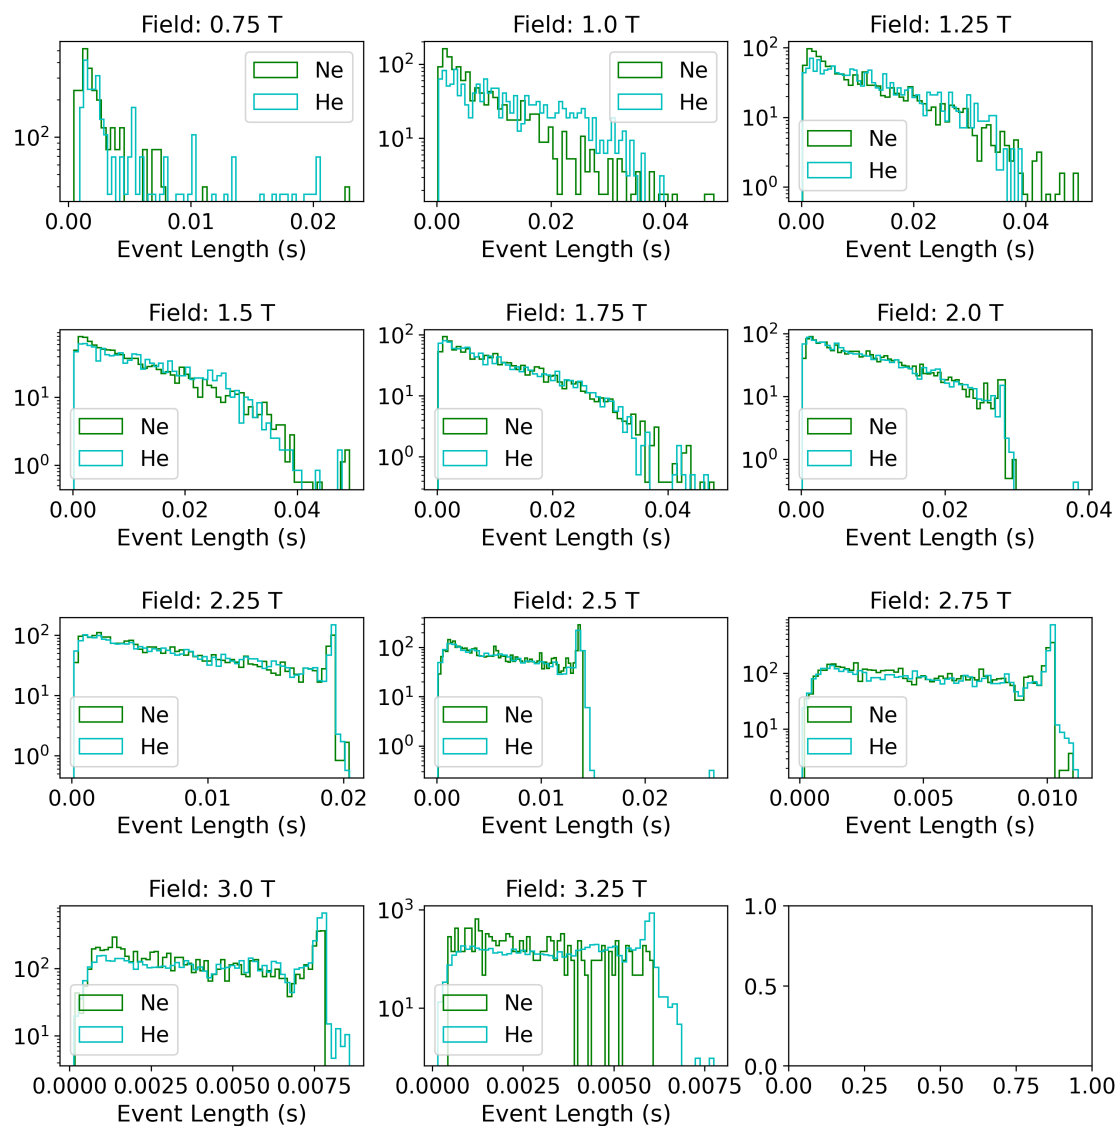


Figure 6.6: Comparison of ^{19}Ne and ^6He event length distributions by field. A Katydid Fourier bin SNR threshold of 9 was used. Note that the peak at the end of the distribution seen at 2.0 T and above is due to events travelling through the entire 1.1 GHz frequency bandwidth.




Figure 6.7: Spectral ratio plot of $^{19}\text{Ne}/^6\text{He}$ observed with CRES including residuals with respect to Monte Carlo (lower inset). Different scattering environments for the two isotopes prevent the cancellation of efficiencies at low fields (not shown).

principal, be identical between the ^{19}Ne and ^6He data and therefore cancel in the ratio. However, as discussed above in Subsection 6.3.1, the noise floors and therefore SNR responses of the detector were slightly different in the ^{19}Ne and ^6He runs. To estimate the size of this effect experimentally we ran our analysis with a range of different Katydid Fourier bin SNR thresholds. The analysis only works within a specific range of SNR thresholds, so we ran the analysis with SNR thresholds that span that functioning range (8, 9, 10, 11)⁷.

The comparison of the experimental ratios obtained to Monte Carlo is shown in Figure 6.7. The comparison is done according to the procedure outlined in Section 6.2.

⁷The reason the analysis only works within a specific range of Katydid Fourier bin SNR thresholds is that at very high thresholds there aren't enough above-threshold points for tracks to be identified by the `sequential-track-finder` processor and with low thresholds there are too many above-threshold points and the `sequential-track-finder` processor is overwhelmed by noise and can't function properly.

We can see that the spread in ratio values obtained for different Katydid SNR thresholds exceeds what one would expect for purely statistical fluctuations, particularly at the highest fields. We assume this indicates a systematic effect that is not cancelling perfectly in the ratio, likely related to the slightly different SNR responses of the ^{19}Ne and ^6He runs. The spread in ratio values obtained from this study cannot (admittedly) be taken to be representative of the full systematic error of this measurement; extensive systematic studies would be necessary to precisely quote a systematic error which is outside of the scope of this work. With the above caveat, a single ratio was obtained with a field-wise average of the ratios obtained in the study above and systematic uncertainties were estimated by calculating half the full range of ratios obtained from different Katydid Fourier bin SNR thresholds (8, 9, 10, 11) and added in quadrature with the statistical uncertainties.

With a more representative field-wise error, we can refit the data. There are three free parameters in the fit; the normalization constant (C), the amount of time the trap is on (t_{trap}), and b_{Fierz} . A minimization of the χ^2 (given by Equation 3.6, except for now the expected ratio is obtained with the ‘from below’ Monte Carlo presented in Section 5.2) was conducted with the python package `lmfit` [80]. As is recommended in the package documentation, first a brute-force grid search over the parameter space is conducted to find rough estimates of the absolute minimum of the objective function. Then, using the best brute-force values to initialize the parameters, the Levenberg–Marquardt algorithm is used to find the best fit values more precisely. The uncertainties assigned to the fit values are obtained by inverting the Hessian matrix of the objective function (χ^2).

Figure 6.8 shows the fit results compared with predicted ratios obtained from Monte Carlo for $b_{\text{Fierz}} = 0, \pm 1$. The fit results for all three free parameters are shown in Table 6.1. Note that the trap-on slew time is left as a free parameter because we don’t have precise knowledge of how long the trap is capable of trapping electrons. We know from both ex-situ measurements of the trapping field with a sensitive pick-up coil and in-situ measurements using the curvature of $^{83\text{m}}\text{Kr}$ events that the 95% - 0% trapping field fall time is roughly ≈ 5 ms. Since the power supply is on for 35 ms, a fit time of ≈ 39 ms is sensible. One can also

see that the maximum event lengths observed in Figure 6.6 and Figure 4.7 are around 40 ms. Future iterations of the experiment would benefit from running the entire experiment with different trap-on times (t_{trap}) in order to get a better handle on these eddy current effects. We are working on an $E \times B$ trap-emptying mechanism which will result in ≈ 0.1 ms trap fall time, largely negating this effect. It is also of note that each of the four ratios shown in Figure 6.7 (for the four different SNR thresholds used), when fit to the ratio predicted by Monte Carlo, result in values for t_{trap} that are consistent within uncertainties with the final value quoted in Table 6.1.

The exact number of CRES events detected per isotope is dependent on the exact Katydid SNR threshold used, with lower SNR cuts resulting in more detected events. For ^{19}Ne we obtained between 2.03×10^4 and 2.58×10^4 total detected CRES events and for ^6He we obtained between 3.81×10^4 and 4.70×10^4 total detected CRES events during the 2022 data run.

The results presented in Figure 6.8 and Table 6.1 represent the realization of the final experimental benchmark of *Phase 1* of the He6-CRES experiment; developing and executing a campaign for an $\mathcal{O}(10^4)$ counts per isotope ratio measurement of ^6He and ^{19}Ne allowing for a proof-of-principle measurement of b_{Fierz} with $\sim 10^{-1}$ statistical sensitivity⁸ There is significant simulation and experimental work to do before a fully defensible measurement of b_{Fierz} can be quoted, but nonetheless, the work presented here represents an important step towards a high precision measurement of b_{Fierz} with CRES.

6.3.2.2 Cutting events from below

It is also worth conducting an analysis in which one attempts to cut all events originating below the visible bandwidth (below 18.0 GHz). This was the original He6-CRES experimental concept and has the advantage that the count rate of events born within the active frequency bandwidth (18.0 – 19.1 GHz) is directly proportional to the spectral density at a

⁸See Subsection 3.1.1 for the full list of experimental benchmarks of *Phase 1*.




Figure 6.8: Spectral ratio plot of $^{19}\text{Ne}/^6\text{He}$ observed with CRES including residuals with respect to Monte Carlo (lower inset). As described in the text, the data presented here is the average of the four ratios obtained in the SNR study described above. Systematic uncertainties were estimated by calculating the range of ratios obtained from the different Katydid SNR thresholds (8, 9, 10, 11) and added in quadrature with the statistical uncertainties. Different scattering environments for the two isotopes prevent the cancellation of efficiencies at low fields (not shown).

Fit parameter (unit)	Fit result
b_{Fierz} (unitless)	0.028 ± 0.051
C (unitless)	0.3155 ± 0.0074
t_{trap} (ms)	38.85 ± 0.68

Table 6.1: Fit results from ratio measurement, including all events coming up from below. The reduced χ^2 of the fit was 2.41, indicating that the errors could still be somewhat underestimated. See text for details.

given field and frequency bandwidth; there is no need to specify the trap-on time (t_{trap}) in Monte Carlo as there is in the analysis of Subsection 6.3.2.1. We simply use the (completely analytic) predicted ratio given by Equation 3.5 to compare prediction to observation.

With this advantage and (apparent) simplicity comes the need to confidently cut all β events that originated below the visible bandwidth. As was discussed above and illustrated visually in Figure 6.2, this is difficult, and why the approach in Subsection 6.3.2.1 that doesn't rely on robust event cuts was favored. That given, there was significant effort put into how to make these cuts in a sensible way so this Subsection is included for completeness and to assist future efforts along these lines.

First we look at what ratio would result after a cut on event start frequency only. We utilize the same analysis approach as is described in Subsection 6.3.2.1: using the ratios and errors obtained from four different Katydid Fourier bin SNR thresholds (8, 9, 10, 11) to approximate the systematic error.

Figure 6.9 shows the ratio obtained when all events with a start frequency less than 200 MHz are cut. This leaves a 100 MHz buffer above the high pass filter at 100MHz for events to be detected and vetoed (making the relevant frequency bandwidth 18.1 – 19.1 GHz). We see that the shape of the ratio presented in Figure 6.9 is in clear disagreement with prediction (reduced $\chi^2 = 7.46$). This is caused by the imperfect vetoing of low SNR events originating

below 18.0 GHz. Since the SNR response of the detector is oscillating, and peaked in the center of the bandwidth (see Figure 6.5) it is possible for sufficiently low SNR events to not be observed at all until 200 to 600 MHz into the bandwidth, as is seen in the low SNR events of Figure 6.2. Such events are also often broken up as they come in and out of being detectable and misclassified as 2 or more separate events, further exacerbating the issue.⁹ This is not a problem that can be fully solved with better analysis strategies as low power events in low SNR frequency regions are often fully indistinguishable from noise.

The exact number of CRES events detected per isotope after this naive event start frequency cut is again dependent on the exact Katydid SNR threshold used, with lower SNR thresholds resulting in more detected events. For ^{19}Ne this cut leaves us with $\approx 60\%$ of events and for ^6He this cut leaves us with $\approx 55\%$ of events.

To understand why the ratio plot of Figure 6.9 looks the way it does we can reference Figure 5.5 which shows the spectral densities for ^{19}Ne and ^6He over our regions of experimental interest. Here we see that in the regions of high energy (and therefore high event slope) where the effect of events coming up from below is most significant, the ^6He spectral density far exceeds the ^{19}Ne spectral density (also demonstrated by Figure 5.11). This means that with imperfect vetoing we get an abundance of ^6He events at high fields, leading to a reduction of the ratio relative to expectation at high fields¹⁰.

A somewhat simplistic but instructive extension to the Monte Carlo presented in Section 5.2, includes (as an input parameter p_{veto}) a probability of being vetoed for each event that originates below the visible bandwidth and rises into being detectable¹¹. When we fit to the ratio data presented in Figure 6.9 we get a best fit veto probability of $p_{veto} = 0.76$, mean-

⁹The time-intercept clustering discussed in Section 6.1 was implemented to mitigate this issue but it only partially addresses the effect.

¹⁰Note that the normalization constant C allows for the ‘centering’ of the residuals in Figure 6.9, somewhat obscuring the experimental ratio’s reduction at high fields.

¹¹Technically p_{veto} in this simulation represents the probability of being vetoed for a simulated CRES event that traverses the entire frequency bandwidth. The veto probability is scaled linearly from 1 down to p_{veto} depending on the frequency extent of the event. This relies on the sensible but naive assumption that the larger the frequency extent of the event the more difficult it is to veto.




Figure 6.9: Ratio plot obtained after naively cutting all events originating below 18.1 GHz (200 MHz), in clear disagreement with prediction (reduced $\chi^2 = 7.46$).

ing the veto being 76% effective best explains the data observed. With this additional free parameter (p_{veto}) we obtain decent agreement between the model and observation (reduced $\chi^2 = 1.90$).

To address the issue of imperfect vetoing of events coming from below, a cut on SNR can be made. This SNR cut must be field dependent as the signal power of CRES events is highly field dependent. The “detectability” of an event is a complex function of time length, slope, and SNR. A full Monte Carlo study of how these event features map on to detectability is underway (see Section 5.3).

In order to make a field-wise SNR cut, a normalized mean event SNR feature was constructed which normalizes the mean event SNR to the range (0,1) based on the maximum and minimum event SNRs observed at the given field for the given isotope¹². The top right plot of Figure 6.10 shows a scatter plot of normalized event SNR vs event start frequency for all ¹⁹Ne and ⁶He ratio data (1.25 - 3.25 T in 0.25 T steps). The nonuniformities in the event density are caused by the oscillation in the SNR response of the detector (see Figure 6.5). As was discussed above, the oscillating SNR response leads to broken events, which means vetoing events originating from below becomes very difficult. To guarantee that all events from below are properly vetoed we need to ensure the SNR of events is high enough that they will not be broken. This means that the cuts chosen should leave us with a (relatively) uniform density of events¹³. With this in mind we choose to cut all events with a normalized SNR less than 0.5. The right column of Figure 6.10 shows the effect of these two cuts. The projection onto the event start frequency axis is now relatively uniform, though does fall off with frequency due to the frequency roll-off in event SNR.

After adding this normalized SNR cut we get the ratio presented in Figure 6.12. As is

¹²Technically the 1% and 99% percentiles are used instead of the minimum and maximum SNRs observed to make the normalization more robust to outliers. This explains why the normalized SNR can be somewhat outside of the range (0,1).

¹³One may (fairly) think that it should not be uniform but reflect the underlying probability distribution of the energy spectrum. The energy bands covered by a single field are narrow enough that this is not a large effect. See Figure 5.5.

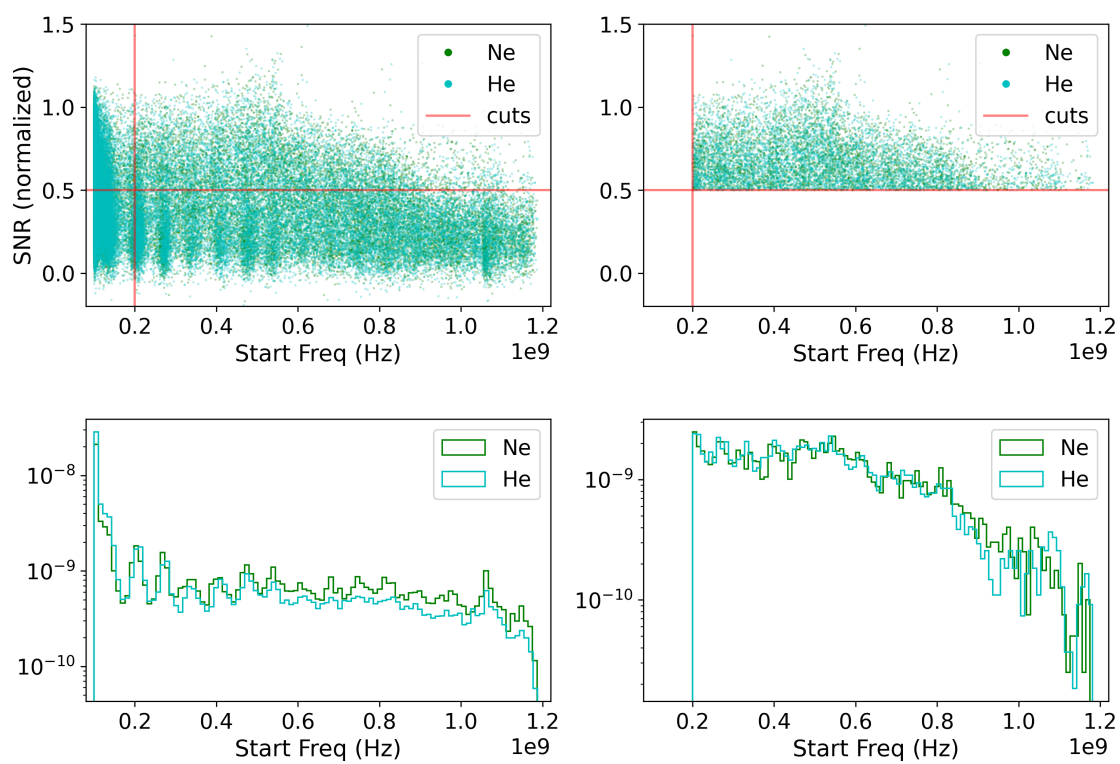


Figure 6.10: Event start frequency vs normalized SNR without (left) and with (right) cuts. Location of the cuts visualized in red.

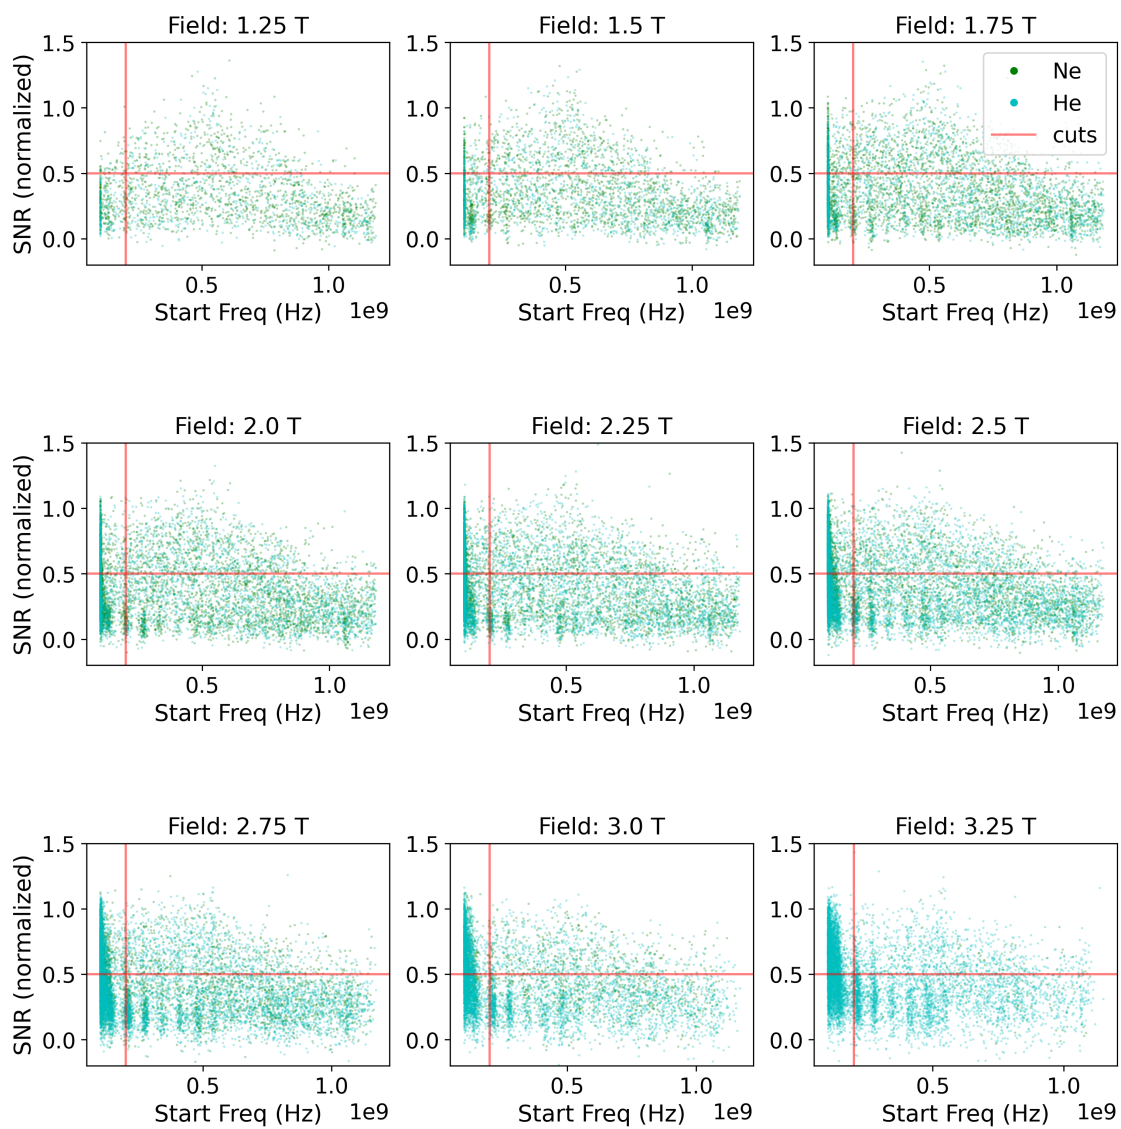


Figure 6.11: Field-wise event start frequency vs normalized SNR with cuts visualized in red.

evident in Figure 6.10, in cutting so aggressively in SNR we drastically reduce our statistics. For ^{19}Ne these two cuts leaves us with $\approx 17\%$ of events and for ^6He they leave us with $\approx 13\%$ of events. With b_{Fierz} is left fixed at 0 in the fit, we end up with a reduced $\chi^2 = 1.06$.

When we leave b_{Fierz} as a free parameter, we get $b_{Fierz} = -0.285 \pm 0.167$ (and a marginally better fit: reduced $\chi^2 = 0.89$). This 1.7 sigma deviation from zero is likely indicative of some imperfect vetoing of events from below. Figure 6.11 shows the field-wise scatter plots of normalized event SNR vs event start frequency for all ^{19}Ne and ^6He ratio data (1.25 - 3.25 T in 0.25 T steps). This figure could shed some light on this bias, as we see that a normalized SNR cut of 0.5 may not fully cut events from below at the highest fields (3.0 T and 3.25 T) based on where the uniform density of events begins in the scatter plots. We do see a dip in the ratio presented in Figure 6.12 compared to expectation at high fields, in accordance with a slight abundance of ^6He events. One could explore an even more aggressive SNR cut but we have already reduced our statistics by a factor of ≈ 7 and this section is meant to illustrate the current difficulty of cutting events from below not to present a measurement of b_{Fierz} .

Now that we have, to the best of our ability, cut events coming from below we can present the data via the alternate presentation described in Subsection 5.1.4 where one isotope (^6He) is used to calculate energy/field dependent detection efficiencies which are then applied to the other isotope (^{19}Ne). This allows for a direct comparison between the ^{19}Ne SM β -energy spectrum and efficiency-corrected ^{19}Ne data.

Figure 6.13 shows the efficiency-corrected energy domain plot for ^{19}Ne with $n_{chunks} = 1$. Note that this is the exact same analysis and fit presented above in Figure 6.12 (reduced $\chi^2 = 1.06$), it's just a different way to present the data. Also note that b_{Fierz} is fixed at 0 and is not left as a free parameter in the fit.

As was discussed in Subsection 5.1.4, systematics such as the wall-effect are greatly mitigated by using finer frequency binning when constructing the ratio (breaking the detection region up into n_{chunks} equal frequency-width bins and build n_{chunks} distinct ratios). Despite this benefit, there is also a reduction in counts per bin when using finer frequency binning.




Figure 6.12: Ratio plot obtained after cutting all events originating below 18.1 GHz (200 MHz) and with normalized SNR less than .5. The addition of the SNR cut leads to significantly better agreement with prediction (reduced $\chi^2 = 1.06$).




Figure 6.13: Efficiency-corrected energy domain plot for ^{19}Ne with $n_{\text{chunks}} = 1$. See text and Subsection 5.1.4 for more details.

Figure 6.14 shows the efficiency-corrected energy domain plot for ^{19}Ne with $n_{\text{chunks}} = 2$, meaning we break our bandwidth into two 500 MHz regions (200 – 700 MHz, and 700 – 1200 MHz). Figure 6.14 clearly suffers from increased statistical fluctuations due to the decreased frequency bin size and is presented here primarily to demonstrate that the approach discussed in the abstract in Subsection 5.1.4 can be applied sensibly to real CRES data.

6.4 Discussion of systematics

The preliminary ratio measurement presented above provides valuable insight into the systematic effects most important to understand and quantify before a competitive 10^{-3} determination of b_{Fierz} via CRES is possible. A detailed projected systematic error budget for a 10^{-3} measurement of b_{Fierz} is beyond the scope of this work but we provide a brief discussion of the important systematic effects encountered during this proof of principle ratio




Figure 6.14: Efficiency-corrected energy domain plot for ^{19}Ne with $n_{chunks} = 2$. See text and Subsection 5.1.4 for more details.

measurement.

1. **Magnetic field uncertainties:** As was discussed in Subsection 6.3.1, the current field matching between ^{19}Ne and ^6He field values would result in a $\approx 2 \times 10^{-3}$ systematic error on a determination of b_{Fierz} . Field shimming could easily result in an order of magnitude reduction of this effect.
2. **Residual gas effects:** Scattering is an important effect at low fields that needs to either be addressed via a more sophisticated analysis that accurately reconstructs scattered CRES events or with significantly decreased residual gas pressure in the decay cell. Note that if both of these approaches prove too difficult it is still possible to have sensitivity to b_{Fierz} without taking data below 1.75 T (where the scattering is most prominent) as was illustrated in Subsection 5.1.1.
3. **Differing β^+ and β^- cross sections:** Even with precise control of the residual gas pressures and species present in the decay cell, differing β^\pm cross sections (with the residual gases) would not cancel in the ratio. This further motivates the need for better analysis at low fields and lower residual gas pressures overall.
4. **Differing SNR responses between ^{19}Ne and ^6He runs:** As was discussed in Subsection 6.3.2, differing SNR responses between isotope runs will not cancel in the ratio. This issue should be fully mitigated by keeping the cryocooler on between runs and taking ^{19}Ne and ^6He data in as short an amount of time as possible.
5. **β monitor instability:** The current version of the β -monitor is stable to within $\sim 10^{-3}$ over 24 hours which corresponds to a $\approx 10^{-3}$ systematic uncertainty. To obtain the next level of statistical uncertainty on b_{Fierz} will take longer data runs. It will need to be ensured that the β -monitor is sufficiently stable over these longer time periods (see Subsection 5.1.3).

6. **False positive event detection:** As far as we can tell via a manual inspection of many hundreds of seconds of data there is an insignificant probability of noise falsely being detected as a CRES event given our analysis procedure. However, as these false positives wouldn't cancel in the ratio (they don't scale with the true event density) a study in which the full analysis is run on many seconds of pure noise data should be run to quantify the size of this effect.
7. **Large line-width and/or over-counting for low SNR events:** This is one of the most important systematics moving forward. The claim that one can obtain high-resolution with CRES ($\delta E/E \sim 10^{-3}$) is challenged for low SNR events. Since the SNR response of the detector is oscillating the start frequency may be misidentified for low SNR events (large line-width) or in the worst case the low SNR event may be misidentified as multiple events (over-counted). Ideally, the SNR response could be flattened by reducing the reflections within the RF system. But even with improved SNR response, this issue will still be present to some extent and so SNR cuts may need to be implemented.
8. **Wall-effect or other detection efficiencies imperfectly cancelling:** This effect is related to the previous effect. If the line shape isn't narrow the arguments about the cancellation of systematics in Section 5.1.5 (which covers the binned ratio technique) are challenged. Exactly how wide a line-shape the experiment can allow for while still having good cancellation of systematics is something that needs to be studied in Monte Carlo.
9. **Events from below:** As was discussed in Subsection 5.2, even with shorter trap-on periods, events rising from below will still be an issue. This issue intersects with the low SNR line-shape issue in that the density of events coming from below is highly isotope dependent and therefore non-vetoed events don't cancel in the ratio. Whatever strategy is used to mitigate this effect (flatter SNR response, better analysis, or an SNR

cut) could be cross-checked experimentally by using different trap-on time periods and ensuring that the resulting ratios are consistent.

10. **Trap emptying and time-on uncertainties:** Currently, due to eddy currents the trap doesn't turn on and off instantaneously. This is an effect that should cancel in the ratio but a Monte Carlo study should verify this if the current trap emptying method is kept in place.
11. **Event density dependence of DBSCAN clustering:** An analysis that relies on density-based clustering to obtain an event count will result in some bias. Concretely, the fraction of time intercepts of independent events within a given Δ_t will increase with the event density and therefor the fraction that are improperly clustered is event-density dependent. In the work presented here, the number of events in a given second is small (~ 5), so this effect is very small but as we push to higher statistics a study needs to be done to quantify the size of this effect.
12. **Radioactive contaminants:** As was discussed in Section 3.3, there could be trace amounts of radioactive contaminants produced in the creation of ${}^6\text{He}$ and/or ${}^{19}\text{Ne}$ which could distort a spectral measurement. Studies putting precise limits on the abundance of contaminants need to be conducted as we push to higher precisions.

The above represents the most important systematics encountered so far. These will need to be considered and accounted for carefully as He6-CRES works towards higher precision.

Chapter 7

DEEP LEARNING IMAGE SEGMENTATION FOR CRES EVENT CLASSIFICATION

In order to precisely measure β -energy spectra over a large range in energy ($\mathcal{O}(keV)$ - $\mathcal{O}(MeV)$) utilizing CRES an analysis must resolve start frequencies for a diverse set of event properties. At low energies ($\mathcal{O}(keV)$) there is significant scattering and a wide distribution of track slopes present due to RF resonances (see Chapter 4). At higher energies ($\mathcal{O}(MeV)$) there is little scattering (longer tracks) and much steeper slopes. In either case, tracks often have some level of curvature as a result of RF resonances. The event reconstruction analysis described in Chapter 6 requires roughly 20 hyperparameters be tuned in order for the reconstruction to work well which makes analysis development time consuming as currently a human must do this hyperparameter tuning manually. It also takes a long time to run (50 seconds of analysis for 1 second of data), struggles to resolve events with significant curvature, only works well for high SNR tracks due to the SNR oscillations with respect to frequency, and isn't capable of differentiating between sidebands and main bands.

A supervised¹ *machine learning* (ML) based analysis has the potential to alleviate many of the above failures of the baseline Katydid (plus python-post processing) model. The following work does not represent a polished presentation of a completely ML-based analysis pipeline that is ready to be applied to real data. The direction of our experiment is still uncertain in terms of what our data may look like in the next iteration of the experiment (time domain vs frequency domain, etc.). Instead this work presents three studies that illustrate the potential of an ML-based analysis to address many of our most complex analysis issues and can be used as a starting point and framework for future ML-based analysis

¹Supervised learning refers to modeling that relies on learning from labeled training data.

development. As long as representative labeled training data can be constructed, an ML-based approach is very flexible with respect to qualitatively different looking data. For example, consider a future iteration of the experiment where we apply match-filtering and our spectrogram tracks become completely flat. As long as training data reflects the new signal structure an ML based model can be quickly retrained without much manual tuning, whereas Katydid’s hyperparameters would need to be completely re-tuned manually and may even break entirely, necessitating the creation of new Katydid processors.

These studies rely heavily on the ground work laid by Luis Saldana and the Project 8 ML team [98]. It is important to note that despite both collaborations utilizing CRES for determination of β -energy spectra, the experimental goals of Project 8 and He6-CRES differ significantly. Our experiment relies on a high statistics determination of spectra over a wide range of energies and on the cancellation of systematics in the ratio of ^{19}Ne and ^6He spectra. P8 is interested in very precise energy determination in a very narrow energy band around the tritium endpoint. As such, our analysis needs differ significantly. For example He6-CRES could currently sacrifice frequency resolution in order to gain statistics and data throughput.

Following from the work done by Luis Saldana of Project 8, we investigate the utility of deep learning algorithms to help solve two of our most difficult analysis issues: SNR oscillations causing broken tracks and the classification of sidebands². CRES signals are voltage time series’ in their simplest form. However, our FPGA (Roach) continually takes FFTs and writes time/frequency spectrograms to disk. These spectrograms are effectively 8-bit gray-scale images and the extraction of meaningful event objects from these images can be framed as a computer vision task.

²Note that we do believe the presence of sidebands cancels in the ratio. However, how many sidebands we observe is currently an unknown. As such, identifying and counting them would enable another important comparison between Monte Carlo and observation.

7.1 Introduction to UNET for CRES

7.1.1 Convolutional neural networks

In recent years deep learning, which refers to the use of Neural Networks (NNs) featuring multiple “hidden” layers between the input and output, has emerged as the most powerful means of approaching image classification, image segmentation, object detection, and instance segmentation tasks [113]. Image classification refers to classifying an object in an image; for example whether an image contains a cat or a dog. Image segmentation refers to the task of assigning each pixel in an image a class label; for example providing a mask that colors all dog pixels as red, all cat pixels as blue, and all pixels that correspond to neither (background) as white. Object detection refers to the task of locating (typically with a rectangular bounding box) and classifying individual objects in an image; for example putting red boxes around each dog and blue boxes around each cat in an image. Instance segmentation is identical to object detection but a mask that overlays each objects bounds in space is also provided in addition to the rectangular bounding box; for example putting red boxes around each dog and a red mask over the contours of the dogs body for each dog in an image and likewise with blue for cats.

Each of the four classes of computer vision tasks described above have a variety of viable different deep learning architectures associated with them. However, all of them utilize the Convolutional Neural Network (CNN) as an essential building block. CNNs use filters to extract high level features from images and can be stacked on top of one another to build more complex and abstract feature sets such as shapes and textures. These extracted features are then used for a variety of object classification and detection tasks, as described above. It is the filters that are optimized during training to extract the most important features relating to the computer vision task at hand [54, 96].

Instance segmentation is most aligned to our physics goals because it would provide us with a count of the number of side-bands and main-bands present in our spectrogram along with masks for each detected object and a score (between 0 and 1) for how confident the

model is about a given detection. With this information we can make cuts on the score and extract signal-to-noise ratios (SNRs), start frequencies, etc via the provided masks. We can also then veto the sidebands or do something more sophisticated with them. The downside is that state of the art instance segmentation models such as Mask-RCNN and YOLO are enormous models ($\mathcal{O}(10M)$ parameters) that take large amounts of training data to obtain good performance [93, 54]. A future study could look at applying this class of ML model to CRES event reconstruction.

7.1.2 UNET for CRES event reconstruction

A simpler approach than applying an instance segmentation model is to utilize an image segmentation model such as UNET (with as few as $\mathcal{O}(1K)$ parameters, described below) to our spectrograms and then use scikit-image (an open-source python image processing package) to transform the image segmentation model into an instance segmentation model. This idea was initially proposed by Luis Saldana. The procedure goes as follows:

1. Apply UNET to a spectrogram. The output consists of one softmax dimension for each class (including background). The n th softmax dimension consists of each pixel being assigned a value between 0 and 1, indicating the probability that that pixel corresponds to class n . We take the maximum probability across the class dimension (`torch.nn.functional.argmax(output, dim = 1)`) to assign a unique class label to each pixel [84].
2. Apply `skimage.measure.label` to the predicted class mask. This will label connected regions with a unique event ID and assign 0's to background. Note that this process will need to be done individually for each (non-background) class if utilizing a multi-class UNET (for example 0 = background, 1 = main bands, 2 = sidebands). This provides us with a mask and unique label for each event in the spectrogram [98, 85].
3. Apply `skimage.measure.regionprops` to the output of the above process. This

method measures properties of the labeled image regions. The properties calculated include the mean intensity of the region which can be used to get the event SNR³, the pixel coordinates of the region which can be used to get the start and end frequencies and times of the region. From these properties we can apply cuts and work towards constructing β -energy spectra [98, 85].

There are two cases in which the above track reconstruction process would fail; disconnected mask regions corresponding to the same track, and overlapping tracks. Disconnected mask regions corresponding to the same event would be assigned two different event IDs by `skimage.measure.label`. This issue is largely mitigated by the fact that a UNET model can learn to ‘connect’ tracks as they traverse regions with low SNR as is illustrated below. The issue of overlapping tracks is not as easily mitigated. Overlapping tracks of differing classes (for example a side-band overlapping a main-band) may be addressed by using the softmax output from UNET intelligently, because in this case both classes will have large scores. This could be used to reconstruct over intersections. A more sophisticated instance segmentation model such as Mask RCNN could address the issue of intersecting tracks of the same class, but in the current iteration of the experiment this is a relatively rare occurrence. Also note that the event reconstruction algorithm described above (similar to the event reconstruction algorithm described and applied in Chapter 6) doesn’t properly reconstruct events that undergo scattering. This could be mitigated with a simple heads-to-tails event building processor as Project 8 has implemented in the past.

7.1.3 UNET architecture

The UNET architecture was first developed in 2015 for use in biomedical image segmentation [96]. The UNET model significantly outperformed the sliding window convolutional network which was the previous best model for such biomedical image segmentation tasks.

³An added benefit of this technique is that the SNR calculated is not dependent on the Katydid Fourier bin threshold chosen.

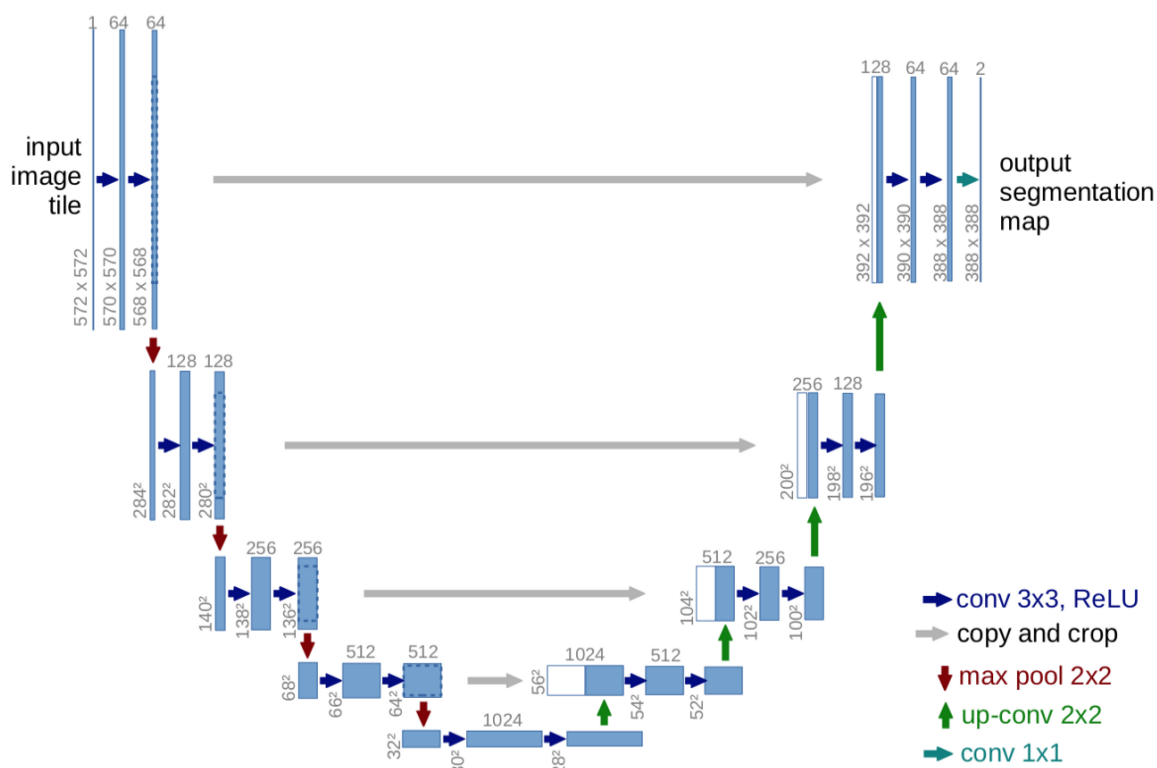


Figure 7.1: UNET model architecture used in the original publication [96]. Each box denotes a feature map with the number of channels on the top of the box. The image shape is given in the lower left of each feature map box. The white boxes represent the copied feature maps resulting from the skip connections and the colored arrows represent the different operations performed on the feature maps. This figure is taken from Reference [96].

Though many advances in utilizing deep learning for computer vision tasks have been made since 2015, UNET remains a competitive algorithm for image segmentation.

UNET is a fully convolutional network consisting of a contracting and an expanding path. The contracting path consists of a typical CNN feature extractor similar to what may be used in a simple object classification model (determining what class of object an image contains). The expanding path builds back up from the extracted features towards a segmentation mask of the same $x - y$ shape as the original image using transpose convolutions (often referred to as up-sampling).

Our specific UNET implementation is python based and utilizes the PyTorch library[84]. The implementation relies heavily on a Double Convolution Layer (DCL) which consists of two consecutive instances of a two dimensional convolution, a two dimensional batch normalization, and a rectified linear unit (ReLU). Our implementation (see Appendix E for the source code) takes the arguments `first_feature_num` and `num_layers` which determines the depth and the number of features to be extracted at each layer.

This flexibility makes it easy to add complexity to the architecture via both the depth and the number of features extracted at each layer. Each layer of the contracting path of the UNET consists of an application of the DCL where the number of features output is twice that of the input number of features input. After each application of the DCL, a max pooling layer of kernel size and stride equal to 2 is applied. A max pool with kernel size (N,N) and stride K implies taking an N x N subset ("pooling window") of the image or feature map and selecting the maximum value, then taking a stride of K to the right and doing the same; repeating until the the entire image has been processed. This results in a reduced spacial size, but an increased in the depth of the feature dimension (originally the color dimension).

At the bottom of the architecture after `num_layers` applications of the DCL and doublings of the number of features there is a bottle-neck which doubles the number of features without the application of a max pool. Then in the expanding path each layer contains a transpose convolution that serves to double the spacial dimension of the feature map (acting as the inversion of the max pool), a skip-connection where the features of the corresponding layer

of the down are appended to the current tensor, and a DCL where the feature number is cut in half. This eventually leads to a feature map with spacial dimensions equal to the original image but with `first_feature_num` features (matching the first layer of the down side).

Here a final 2D convolution with kernel size and stride 1 maps this tensor into a tensor with the same x-y shape as the original image but with `num_classes` features. The elements of this output tensor will be raw logits and we can apply a softmax layer (`torch.nn.functional.softmax(output, dim = 1)`) in order to normalize them across the class dimension; now there are `num_classes` layers, the i^{th} layer representing the probability of the pixel belonging to class i . We can then apply a `torch.nn.functional.argmax(ouptut, dim = 1)` in order to obtain a set of 2D labels where now each pixel's value represents the predicted class of that pixel.

The skip connection adds spacial precision to the model. We have added in the ability to turn off the skip connections in the model and it is striking how much more poorly it performs. It is totally unable to provide clear borders on object borders without the skip connections. Intuitively, it adds the spacial information to the semantic information (abstract information about texture and shape) so that the most accurate segmentation is possible.

The above is a description of the specific UNET implementation we built for this work. Figure 7.1 shows the original UNET architecture presented in [96]. Our architecture allows for a variable number of features and layers which makes it easy to compare models of different complexities without writing additional code. The only substantive difference between our implementation and the originally proposed UNET is that our Double Convolution Layer uses batch normalization after each convolutional layer. Batch normalization has been found to accelerate the speed of training by enabling larger learning rates via mitigating the fact that the distribution of each layer's inputs in a deep neural net are affected as the previous layers are altered⁴ [60].

⁴Batch normalization also helps to alleviate the 'vanishing gradient problem' [60].

7.1.4 Loss function: cross-entropy loss

We use the cross-entropy loss as our objective function for this study. Cross-entropy loss, also known as negative log-likelihood loss, is used often for classification tasks in machine learning. It is a measure of the dissimilarity or distance between the predicted probability vector X and the ground truth probability vector Y . The ground truth probability vector Y is represented by the one-hot encoded labels, while the predicted vector X is represented by the predicted probability of each class.

The weighted cross-entropy loss is given by:

$$CE(X, Y) = \sum_{i=1}^{i=N} \sum_{j=1}^{j=C} -w_j Y_{ij} \ln X_{ij}, \quad (7.1)$$

where N is the number of training examples, C is the number of classes, and Y and X are the ground truth labels and prediction respectively and are both two dimensional probability vectors of shape (N, C) . Both X, Y reduce to a vector of N 1's when summed along the second (class) dimension of length C . The predicted probability vector X is obtained by applying a softmax layer to the raw model outputs (logits). For us the prediction vector Y has a single 1 and the rest 0's because we assign only one class per pixel. Each class is given a weight w_j which can be used to address class imbalances. For our current purposes the number of training examples N isn't the number of training images n_{images} but is instead given by, $N = HxWxn_{images}$, because each pixel is treated as a unique example.

7.1.5 Gradient descent algorithm: Adam

A gradient descent algorithm is a crucial component of any machine learning model; it's used to find the model parameters that minimize the loss function (cross-entropy loss for this study). During training the parameters of the neural net (or machine learning model generally) are initialized with a set of values and then the gradient descent algorithm is used to iteratively alter the parameters according to the gradient and a learning rate. A simple gradient descent algorithm step can be described as

$$\theta_{i+1} = \theta_i - \ell \nabla_{\theta} J(X, \theta)|_{\theta_i}, \quad (7.2)$$

where θ_i is the vector of all parameters in the model at step i of the, ℓ is the learning rate, X is the entire training data set, and J is the loss function whose gradient is being evaluated at θ_i . This process is continued until a specified early stopping criteria is met or the specified number of training steps have been executed. The intuition for the algorithm is that you are simply “walking” down the loss function via the steepest path (given by the negative of the gradient) with a given speed (given by the learning rate ℓ).

Stochastic gradient descent is a variant of the gradient descent algorithm where the gradient is calculated not over the entire data set but for each element in the training set individually. This means that the steps taken are generically more noisy but since each step can be calculated more quickly with much less computation. One can imagine that now the path towards the minimum of the loss function is now more wandery. Despite this less direct path towards the optima, stochastic gradient descent typically works well and is generically much faster than vanilla gradient descent because of the increased speed of each step.

In this study we used the PyTorch implementation of the Adam (Adaptive Moment Estimation) optimizer `torch.optim.Adam()`. Adam is a type of stochastic gradient descent that relies on adaptive learning rates to improve convergence. Adam uses the moving average of the gradient and the squared gradient to scale the learning rate and make intelligent parameter updates as training progresses. It has been found to be faster and more efficient than many other implementations of stochastic gradient descent such as AdaGrad and AdaDelta. A full explanation of the algorithm can be found in Reference [67].

7.1.6 Machine learning metrics

Below we summarize the three machine learning metrics used in the performance assessment of the UNET model in the sections to follow. In machine learning metrics are used to quantify the similarity between the targets (ground truth labels) and the model predictions. In this

work we use Accuracy, F1-score, and Intersection over Union (IoU), which we define below. We define the following: TP = True Positives, TN = True Negatives, FP = False Positives, and FN = False Negatives. Here TP would be the count of all times the target value is 1 and the predicted value is 1. The following definitions only hold for binary classification tasks, meaning only class labels of 0 and 1 are used. However, these definitions can be generalized to multi-class classification tasks in a few different ways, the most common being to convert the multi-class problem into a series of binary tasks by using a one-versus-the-rest approach. We won't cover multi-class classification metrics but an in-depth overview can be found in Reference [63].

Accuracy is the number of correct predictions over the total number of predictions. Mathematically it is defined as:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}. \quad (7.3)$$

The F1-Score is the harmonic mean of the precision and recall (see Ref. [63] for definitions), and is defined mathematically as:

$$F_1 = \frac{2TP}{2TP + FP + FN}. \quad (7.4)$$

The IoU, also called the Jaccard index, is intuitively the ratio of the intersection and the union of the positive targets and the positive predictions. It is commonly used in object detection computer vision tasks. For a binary classification task the IoU is defined as:

$$IoU = \frac{TP}{TP + FP + FN}, \quad (7.5)$$

The F1-Score and the IoU are generally more suited to classification tasks with severe class imbalance than the accuracy. As an example, imagine a set of targets with 99% false or 0 values. A prediction of all false will have 99% accuracy. F1-score and IoU will both be 0, as there are no true positives ($TP = 0$), which is more desirable and more representative of the lack of sophistication of the prediction.

7.2 Overview of the three studies conducted

To explore the utility of a shallow UNET architecture for the specifics of our CRES signal structure and experimental goals three separate studies, each containing simulated training data with different characteristics, were conducted. Instead of starting with a data set that simultaneously contains all of the features we would like to ultimately detect/resolve, we start with the simplest demonstration of each feature individually, verifying that the model can perform well in these individual circumstances. We hope to move on to creating a training dataset that is true to our CRES data and contains all of these individual features represented simultaneously and realistically. This will likely require much more training data and correspondingly more computational resources than were needed for these studies. Below is a brief description of the simulated training dataset used for each study:

1. **Study 1:** Events with constant SNR, without the presence of sidebands.
2. **Study 2:** Events affected by severe SNR fluctuations, without the presence of sidebands.
3. **Study 3:** Events with constant SNR, with the presence of sidebands.

The first study is used to validate that a shallow UNET architecture can provide high quality segmentation masks for the simplest type of signal we encounter. This was, in principle, already verified by the Project 8 Machine Learning Team, but the specifics of their implementation were different so it is still a useful first step and sanity check. The baseline model, utilizing the sequential track finder processor of Katydid to reconstruct short track segments of a maximum length (typically set to 5 bins) and then a layer of DBSCAN clustering of adjacent track segments, often performs poorly when large SNR fluctuations are present and is completely unable to differentiate sidebands from main bands. The second and third studies aim to explore the deep learning model's performance in these cases when the baseline model described in Section 6.1 fails.

A higher precision iteration of our experiment or a more advanced analysis of our data will need to be able to precisely determine the start frequencies of tracks with severe SNR fluctuations and, ideally, the presence of sidebands. The deep learning UNET approach has the promise to accurately segment spectrograms with SNR fluctuations and sidebands.

7.2.1 *Simulated training data*

Synthetic signals are produced in the time domain and the action of the Raorch FPGA board and 17.9 GHz downmixer are simulated to produce realistic spec files identical to those written to disk during experimental data acquisition in terms of bin and slice structure, and noise distribution. The simulated training data for these three studies were built to resemble the data set taken in the Autumn of 2022. Each simulated spec file consisted of 4096 frequency bins corresponding to $N = 2^{12}$ points per FFT and a 2400 MHz sampling rate. This results in a 1200 MHz bandwidth, and roughly 300 kHz per frequency bin. Consecutive $3.4 \mu\text{s}$ time slices of the FFT output are summed, resulting in a spec file time slice of approx $6.8 \mu\text{s}$.

For each simulated spectrogram a corresponding mask was generated that contains the labels for the individual pixels. For example pixels corresponding to background are labeled 0 and pixels corresponding to signal are labeled 1. The details of the specific masks generated for the different experiments are presented below.

The simulated training datasets were created using Python vectorized Numpy operations for optimal computational efficiency. However, the simulations, which rely on taking the FFT of high frequency time domain signals, are currently slow; averaging $\mathcal{O}(30 \text{ s})$ to simulate 1 s of data. A future iteration of these simulations may be better suited to C++ which is generally faster.

Each of the three datasets shown below were simulated without scattering and with normally distributed slopes with mean 10^{11} Hz/s and standard deviation 10^{10} Hz/s . This was to approximately resemble CRES data taken at 2.75 T (given our 18.0 - 19.1 GHz bandwidth). If a larger dataset with many different slopes was used (requiring more computing power)

then the UNET should be able to simultaneously segment events of all different slopes without extensive hyperparameter tuning.

7.2.2 Preprocessing data with max pooling

One of the major challenges facing the He6-CRES experiment is the enormous amount of data being written to disk. We have 600 MB/s being written to disk. Over the 11 fields observed for ^{19}Ne (and similarly for ^6He) we observed $\approx 10^4$ counts with 500 s of data taken at each field. This corresponds to 3.3 TB for 10^4 counts or .33 GB/count (over the full spectrum of ^{19}Ne 0 - 2.2 MeV). In order to reach 10^{-3} precision on b_{Fierz} roughly 10^8 counts are needed per isotope. This corresponds to 33 petabytes of data per isotope. Even if a factor of 10 increase in event rate is achievable through less trap slewing downtime and increased density of radioactivity in the decay cell, this is an untenable amount of data to write to disk.

There are many potential ways to address this issue. One of them is to apply some sort of max or mean pooling to the raw spectrograms before they are written to disk. Another is to apply a power cut and then apply a zero-suppression algorithm to the data before writing to disk. In either case information is lost, however this loss in information may or may not have a significant impact on the CRES event reconstruction efficiency which is what effects the quality of the measurement.

Since the features and textures of our tracks are very simple (lines with potential curves) the information contained in the spectrograms could be mostly maintained while potentially reducing the size of the images significantly. In the three studies presented below a max pooling layer with kernel size (16,16) and stride 16 has been applied to the raw spectrograms before they are used to train the model or used for validation and testing. This is, among other reasons, to simulate the decision to apply this max pooling layer before data is written to disk. A max pool of this type applied before data is written to disk would save a factor of $16^2 = 256$ in memory. This would take the 33 petabytes necessary for a competitive measurement to a more reasonable 129 TB.

There is risk in applying a max pooling layer before data is written to disk in that the frequency resolution of the track start frequency is degraded by a factor of roughly 16. There is also a risk of losing low SNR events below the noise floor, as the max pooling will degrade the SNR of the tracks. In my experimentation I found an aggressive max pooling application as is described here which reduces data by a factor of 256 may reduce the SNR of signals by a factor of 2.5 - 3. However, the SNR isn't the only factor in track detection and the max pooling layer also reduces the variance of the noise which generally helps detection efficiency. Additionally, having smaller images may make detection easier in certain ways: less class imbalance between background and signal for example. A more quantitative study of the effects of max pooling the raw spectrogram on detection efficiency needs to be conducted before we decide to apply this to a full ratio experiment.

The motivation to apply this max pooling layer was two-fold. One, to explore if an analysis could perform well on these smaller lightweight images. And two, it makes experimentation with deep learning models much easier if images are small. Fewer computation resources and less time to train makes it easier to experiment and iterate which is crucial to this initial exploratory phase of analysis development. If it is decided that too much spacial knowledge is lost with this aggressive of a max pool then the models presented here will need to be retrained, and may need to be deepened, but in principle they should still perform just as well.

The spectrograms simulated for each of the studies are all 35 ms long. This length was chosen to match the amount of time the trap was on (t_{trap}) during the Autumn 2022 data run. In principle an analysis that works well on tracks of maximum length 35 ms should perform just as well on a full 1 s spec file. The 35 ms file corresponds to an array of shape (4096, 5126), and after the max pooling is applied the shape is (256, 320).

7.2.3 UNET implementation details

Our flexible, variable depth implementation of the UNET architecture, as described in Section 7.1.3, is written in Python with the PyTorch library that was initially released in 2016

and was developed by Meta AI [84]. We utilize the Pytorch Lightning Python library which provides a high-level interface for writing PyTorch code, tracking metrics, and saving model checkpoints. Specifically, we used tensorboard for tracking model metrics during training and saving model checkpoints.

There are three Pytorch Lightning classes that are crucial to training. The LightningModule is an object that contains the model, the training and validation steps, metrics, the optimizer, and the loss function. The Lightning DataModule contains methods for loading data from disk into PyTorch dataloaders and also does any preprocessing (such as the max pool step described above). The third important class is the Lightning Trainer which governs which devices (CPU, GPU, Cluster,...) are used for training, how many epochs⁵ to run, early stopping, and how often to log model checkpoints. Once all three of these objects have been appropriately instantiated training is as simple as `Trainer(LightningModule,DataModule)` (pseudo-code). In this way the many components of the entire deep learning research pipeline are containerized into three main objects and boiler-plate code is drastically reduced.

Both the simulation of the training data and the model training was done on a personal Linux laptop CPU. For each study presented below it took less than 1 hr to build the simulated dataset and train. Future development using more training data may need to utilize the University of Washington GPU research cluster HYAK to keep simulation and training wall-times reasonable.

Appendix E contains the source code for the PyTorch implementation of the UNET model as well as the LightningModule, which contains all of the specific training components used such as the Adam optimizer and cross-entropy loss function. All of the code used to conduct these studies as well as instructions for reproducing or expanding upon this work can be found here: <https://github.com/Helium6CRES/he6-cres-deep-learning>.

⁵The number of epochs is the number of passes taken through the entire training set during training.

7.3 Study 1: Constant SNR tracks, no sidebands

This first study was used to verify the ability for our implementation of a shallow UNET architecture to segment spectrograms containing simple, constant SNR tracks with no sidebands present. This was already demonstrated by Luis Saldana with a much deeper UNET architecture via a TensorFlow implementation [98]. It is presented as a sanity check and a verification that our PyTorch implementation is working well.

7.3.1 Specifics of simulated training data

Here we describe the specific simulation parameters of the data set used in this study. There were 50 spectrograms simulated, each 35 ms long. In total 200 events were simulated, uniformly distributed over the 50 simulated spectrograms for an average of 4 events per file. All of the events were simulated as a pure chirp (straight, not curvy), without scattering, with a flat noise floor, and with a flat SNR. The track lengths were uniformly distributed between 0 and 35 ms though many were shorter than their allocated length due to reaching the top of the 1200 MHz frequency bandwidth. The start frequency of the events were uniformly distributed from the top of the high pass filter to the top of our bandwidth; 100 - 1200 MHz. This is not technically physical as there is always an underlying start frequency distribution due to the energy spectrum but this is a small effect and not crucial to these studies. The slopes were normally distributed with mean $1e11$ Hz/s and standard deviation $1e10$ Hz/s (to approximate 2.75 T data). The SNR of the event was constant along the extent of the event and the SNR for each individual event was normally distributed with mean 8.2 and standard deviation ≈ 1.2 . Note that the above is a description of the spectrogram before the application of the max pooling. The SNR values were reduced by a factor of 2.5-3 as mentioned above by the action of the max pooling.

7.3.2 Model hyperparameters and training configuration

Here we describe the specific UNET hyperparameters and training settings used in the training of the model. A proper hyperparameter grid search wasn't conducted because such good performance was obtained with these parameters, which were found with relative ease. The PyTorch Lightning DataModule used a batch size of 8, a preprocessing max pool with kernel and stride 16 (as described above), and split the 50 spec file dataset into training, validation, and test dataloaders with following relative proportions: 0.6, 0.3, 0.1. The Adam optimizer used for training had a learning rate of $5 * 10^{-4}$. The cross-entropy loss used a weight of 10:1 on signal:noise pixels, which accelerated and rewarded the identification of the signal pixels. The number of output classes from the model, `num_classes`, was two: signal, noise. The input `first_feature_number` was 4, and the input `num_layers` was 2, meaning the number of features extracted was 4, then 8 in the two layers of the contracting side of the UNET. The model was trained for 250 epochs.

7.3.3 Study results

The model performed extremely well with the hyperparameters and training configuration described above and it took only took 16 minutes to complete the 250 epochs of training on a laptop CPU. A visualization of the model's performance on a single spectrogram from the test set (not seen during training or validation) is presented in Fig. 7.2, and a zoomed in version of the prediction is shown in Fig. 7.3. To qualitatively assess the model's performance, compare the top right subfigure (target or ground truth), to the bottom right subfigure (model prediction after full course of 250 epochs of training) in Fig. 7.2, and Fig. 7.3. In both the targets and predictions the transparent overlay uses the class mapping: white = class 0 = background, red = class 1 = track. The prediction masks were obtained by applying `torch.nn.functional.argmax(output, dim = 1)` to the raw logits output by the model.

In the bottom right subfigure of Fig. 7.3 we see that the model misidentifies the end of the event by a few pixels which would result in an inaccurate start frequency determination. Luis

Saldana and the Project 8 Machine Learning Team developed a method of trimming these extra pixels. We don't present that process here, but we could also in principle implement such a method [98].

Figure 7.4 shows the Loss, Accuracy, F1-score, and IoU of the model for both training and validation over the course of the 250 epochs of training. These metrics are defined in Section 7.1.6, and as was discussed, the F1-score (Eq. 7.4) and the IoU (Eq. 7.5) are most representative of model performance for classification tasks with severe class imbalance such as this. For most computer vision tasks an IoU of 0.7 or even 0.5 is considered strong performance, and here the validation IoU exceeded 0.95, indicating excellent performance.

7.4 Study 2: Tracks with oscillating SNR, no sidebands

The second study was used to verify the ability for a shallow UNET architecture to segment spectrograms containing tracks with severely oscillating SNR and with no sidebands present. Intuitively, this asks the model to provide a continuous mask for tracks that are only periodically 'visible', a task that the baseline model of connecting adjacent high power bins cannot do.

7.4.1 Specifics of simulated training data

Here we describe the specific simulation parameters of the data set used in this study. The simulation parameters were identical to the ones used in Study 1 (Sec. 7.3) except for the simulated gain and SNR. The mean SNR of the event was normally distributed with mean 7.0 and standard deviation ≈ 1.05 . A sinusoidal gain function with mean 1, amplitude 1 (resulting in a min value of 0 and a max value of 2), and frequency of $\frac{2\pi*20}{1200e6Hz}$ was applied to the post-FFT signal (and not the noise). This resulted in an SNR that had 20 periods contained in the full frequency bandwidth of the experiment. At the trough the signal had zero amplitude.

7.4.2 Model hyperparameters and training configuration

Here we describe the specific UNET hyperparameters and training settings used in the training of the model. A proper hyperparameter grid search wasn't conducted because such good performance was obtained with these parameters, which were found with relative ease. The PyTorch Lightning DataModule used a batch size of 8, a preprocessing max pool with kernel and stride 16 (as described above), and split the 50 spec file dataset into training, validation, and test dataloaders with relative proportions: 0.6, 0.3, 0.1. The Adam optimizer used for training had a learning rate of $10 * 10^{-4}$. The cross-entropy loss used a weight of 10:1 on signal:noise pixels. The number of output classes from the model, `num_classes`, was two (signal, noise). The input `first_feature_number` was 4, and the input `num_layers` was 3, meaning the number of features extracted was 4, then 8, then 16 in the three layers of the contracting side of the UNET. It was found that the model didn't perform as well with only 2 layers. The model was trained for 250 epochs.

7.4.3 Study results

The model performed extremely well with the hyperparameters and training configuration described above and it took only took 11 minutes to complete the 250 epochs of training on a laptop CPU. A visualization of the model's performance on a single spectrogram from the test set (not seen during training or validation) is presented in Fig. 7.5, and a zoomed in version of the prediction is shown in Fig. 7.6. One can see that the model 'understands' that the track still exists between two regions of high SNR.

Figure 7.7 shows the Loss, Accuracy, F1-score, and IoU of the model for both training and validation over the course of the 250 epochs of training. Here, similarly to the first study, the validation IoU exceeded 0.95, indicating excellent model performance.

7.5 Study 3: Constant SNR tracks, sidebands

The third study was used to verify the ability for a shallow UNET architecture to perform a three class classification; differentiating between mainband, sideband, and background pixels. Intuitively, this asks the model to use the surrounding pixels and features to decide if a given pixel belongs to a mainband or sideband, a task that the baseline model, which is inherently binary (differentiating between signal or noise), cannot do.

7.5.1 Specifics of simulated training data

Here we describe the specific simulation parameters of the data set used in this study. There were 100 spectrograms simulated, each 35 ms long. In total 200 events were simulated, uniformly distributed over the 100 simulated spectrograms for an average of 2 events per file. The track lengths, start frequencies, and slopes of the events were generated in an identical fashion to the first two studies. Each event was simulated as a frequency modulated chirp with modulation index h uniformly distributed between .9 - 1.1 and an axial frequency f_{axial} uniformly distributed between 60 MHz - 100 MHz. All of the events were simulated without scattering, with a flat noise floor, and with a flat SNR. The amplitude of the signal was fixed such that the main band had a mean SNR of 11.68, and the sidebands had a mean SNR of 6.62, with the spread in track SNRs originating from the spread in modulation indices. Note that, just as above, this is a description of the spectrogram before the application of the max pooling. The SNR values were reduced by a factor of 2.5-3 as mentioned above by the action of the max pooling.

7.5.2 Model hyperparameters and training configuration

Here we describe the specific UNET hyperparameters and training settings used in the training of the model. A proper hyperparameter grid search wasn't conducted because such good performance was obtained with these parameters, which were found with relative ease. The PyTorch Lightning DataModule used a batch size of 8, a preprocessing max pool with ker-

nel and stride 16 (as described above), and split the 100 spec file dataset into training, validation, and test dataloaders with relative proportions: 0.6, 0.3, 0.1. The Adam optimizer used for training had a learning rate of $8 * 10^{-4}$. The cross-entropy loss used a weight of 20:10:1 on mainband:sideband:background pixels. The number of output classes from the model, `num_classes`, was three (mainband, sideband, background). The input `first_feature_number` was 4, and the input `num_layers` was 3, meaning the number of features extracted was 4, then 8, then 16 in the three layers of the contracting side of the UNET. The model was trained for 500 epochs.

7.5.3 Study results

The model performed extremely well with the hyperparameters and training configuration described above and it took 56 minutes to complete the 500 epochs of training on a laptop CPU. A visualization of the model's performance on a single spectrogram from the test set (not seen during training or validation) is presented in Fig. 7.8, and a zoomed in version of the prediction is shown in Fig. 7.9. One can see that the model can precisely differentiate between sidebands and mainbands. In Fig. 7.10, we see an example where there are overlapping events. We see that even in this confusing example the model segments the spectrogram accurately. It is interesting to note that since the weight given to the main bands in the loss function was higher, the classification of the mainband pixels occurs earlier, as can be seen in the bottom left subfigures of Fig. 7.9, and Fig. 7.10. In both the targets and predictions shown in Fig. 7.8, 7.9 and Fig. 7.10 the class mapping is: white = class 0 = background, red = class 1 = main band, green = class 2 = sideband.

Figure 7.11 shows the Loss, Accuracy, F1-score, and IoU of the model for both training and validation over the course of the 500 epochs of training. Here, similarly to the first and second study, the validation IoU exceeded 0.95, indicating excellent model performance. For the metrics shown in Figure 7.11, the macro-average is used. This means that the given metric for each class is calculated using a one-versus-the-rest strategy and then the average across each class is taken.

This study represents the first sideband classification algorithm developed by our collaboration. Despite this being a very limited demonstration, given that such a limited range of modulation indices and axial frequencies were simulated, it illustrates that a lightweight UNET model such as this could perform well to classify sidebands.

7.6 Conclusion of UNET studies

The process of image segmentation using UNET followed by using the sklearn method (`skimage.measure.label`) to obtain instance segmentation, as discussed in Section 7.1.2, could be a powerful CRES event reconstruction algorithm. We’ve demonstrated that the approach could accurately reconstruct tracks over frequency regions with an oscillating SNR response and could be used to detect and therefor count or veto sidebands. Though this was not explicitly demonstrated, the UNET model should also perform well with curvy tracks, as long as the training data contains curvy tracks.

We have not put significant work into incorporating an ML-based model such as this into our Katydid based analysis pipeline but Chapter 5.5 of Luis Saldana’s thesis contains a good description of how this could be done [98]. It is important to note that currently a heads-to-tails approach as is implemented in Katydid is still necessary to deal with clustering scattered β s into one event object. This was largely ignored here for simplicity and because our data contains little scattering above 1.5 T.

The above ML-based event reconstruction using UNET has distinct advantages over the baseline Katydid model but it doesn’t intrinsically address scattering, it fails when main events overlap, and it could only be used to veto sidebands but not to group associated sidebands and mainbands into one event. An instance segmentation model such as Mask-RCNN or YOLO could solve these issues [54, 93]. An instance segmentation model such as this could accurately group scattered tracks, sidebands, and overlapping tracks into events. The main downside to these models is that they are $\mathcal{O}(10 \text{ million})$ parameters compared to our shallow UNET models which have $\mathcal{O}(10 \text{ thousand})$ parameters, so are slower to train and take longer to make predictions on data. What to pursue ultimately depends on the

physics goals of He6-CRES. Winston DeGraw of He6-CRES has made some progress towards the application of the Mask-RCNN instance segmentation on CRES data (<https://github.com/wdegrow93/he6-cres-deep-learning>). Figure 7.12 demonstrates the model's ability to properly resolve scattering.

These three studies and the flexible modules presented in Appendix E are meant to serve as a starting point for future analysis development. In order to have a data-ready model we will need more realistic and abundant training data.





Figure 7.2: Track identification performance on a 35 ms simulated spectrogram from the test set (not seen during training). Simulated tracks have fixed SNR and no side-bands. In both the targets and predictions the class mapping is: white = class 0 = background, red = class 1 = track. (a) Raw 35 ms spectrogram from test set. (b) Spectrogram with overlaid targets ($\alpha = .4$). (c) Spectrogram with overlaid predictions from the 25th epoch of training ($\alpha = .4$). (d) Spectrogram with overlaid predictions from the 250th epoch of training ($\alpha = .4$).



(a) Raw spectrogram.

(b) Overlaid targets.


(c) Overlaid predictions from from 25th epoch. (d) Overlaid predictions from 250th epoch.

Figure 7.3: Track identification performance on a 35 ms simulated spectrogram from the test set (not seen during training); zoomed in. In both the targets and predictions the class mapping is: white = class 0 = background, red = class 1 = track. (a) Raw 35 ms spectrogram from test set. (b) Spectrogram with overlaid targets ($\alpha = .4$). (c) Spectrogram with overlaid predictions from the 25th epoch of training ($\alpha = .4$). (d) Spectrogram with overlaid predictions from the 250th epoch of training ($\alpha = .4$).

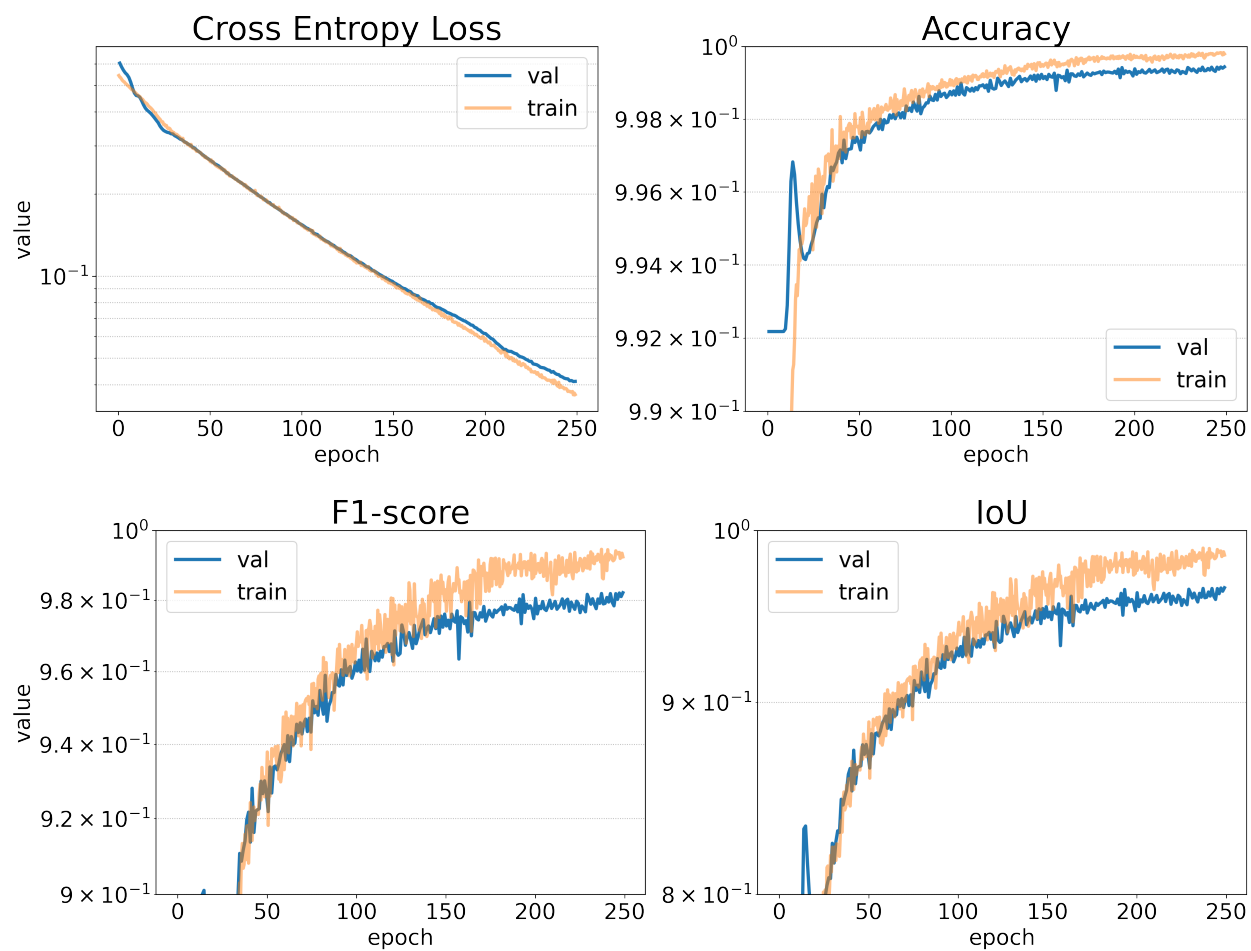


Figure 7.4: Summary of model performance for study of tracks with fixed SNR and no sidebands.





Figure 7.5: Track identification performance on a 35 ms simulated spectrogram from the test set (not seen during training). Simulated tracks have severe oscillations in SNR and no side-bands. In both the targets and predictions the class mapping is: white = class 0 = background, red = class 1 = track. (a) Raw 35 ms spectrogram from test set. (b) Spectrogram with overlaid targets ($\alpha = .4$). (c) Spectrogram with overlaid predictions from the 25th epoch of training ($\alpha = .4$). (d) Spectrogram with overlaid predictions from the 250th epoch of training ($\alpha = .4$).



(a) Raw spectrogram.

(b) Overlaid targets.


(c) Overlaid predictions from from 25th epoch. (d) Overlaid predictions from from 250th epoch.

Figure 7.6: Track identification performance on a 35 ms simulated spectrogram from the test set (not seen during training); zoomed in. Simulated tracks have severe oscillations in SNR and no side-bands. In both the targets and predictions the class mapping is: white = class 0 = background, red = class 1 = track. (a) Raw 35 ms spectrogram from test set. (b) Spectrogram with overlaid targets ($\alpha = .4$). (c) Spectrogram with overlaid predictions from the 25th epoch of training ($\alpha = .4$). (d) Spectrogram with overlaid predictions from the 250th epoch of training ($\alpha = .4$).

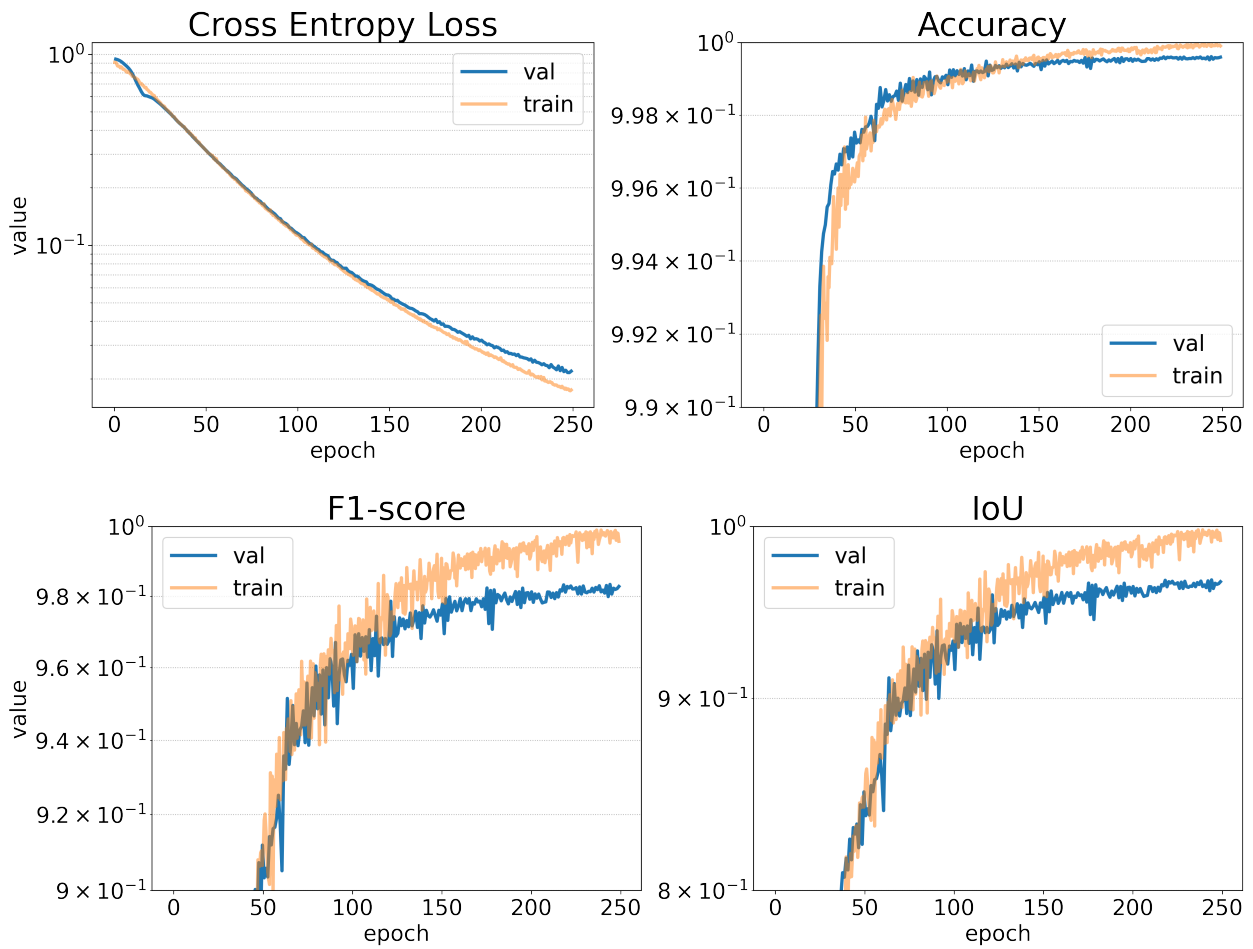


Figure 7.7: Summary of model performance for study of tracks with sever oscillations in SNR and no side-bands.

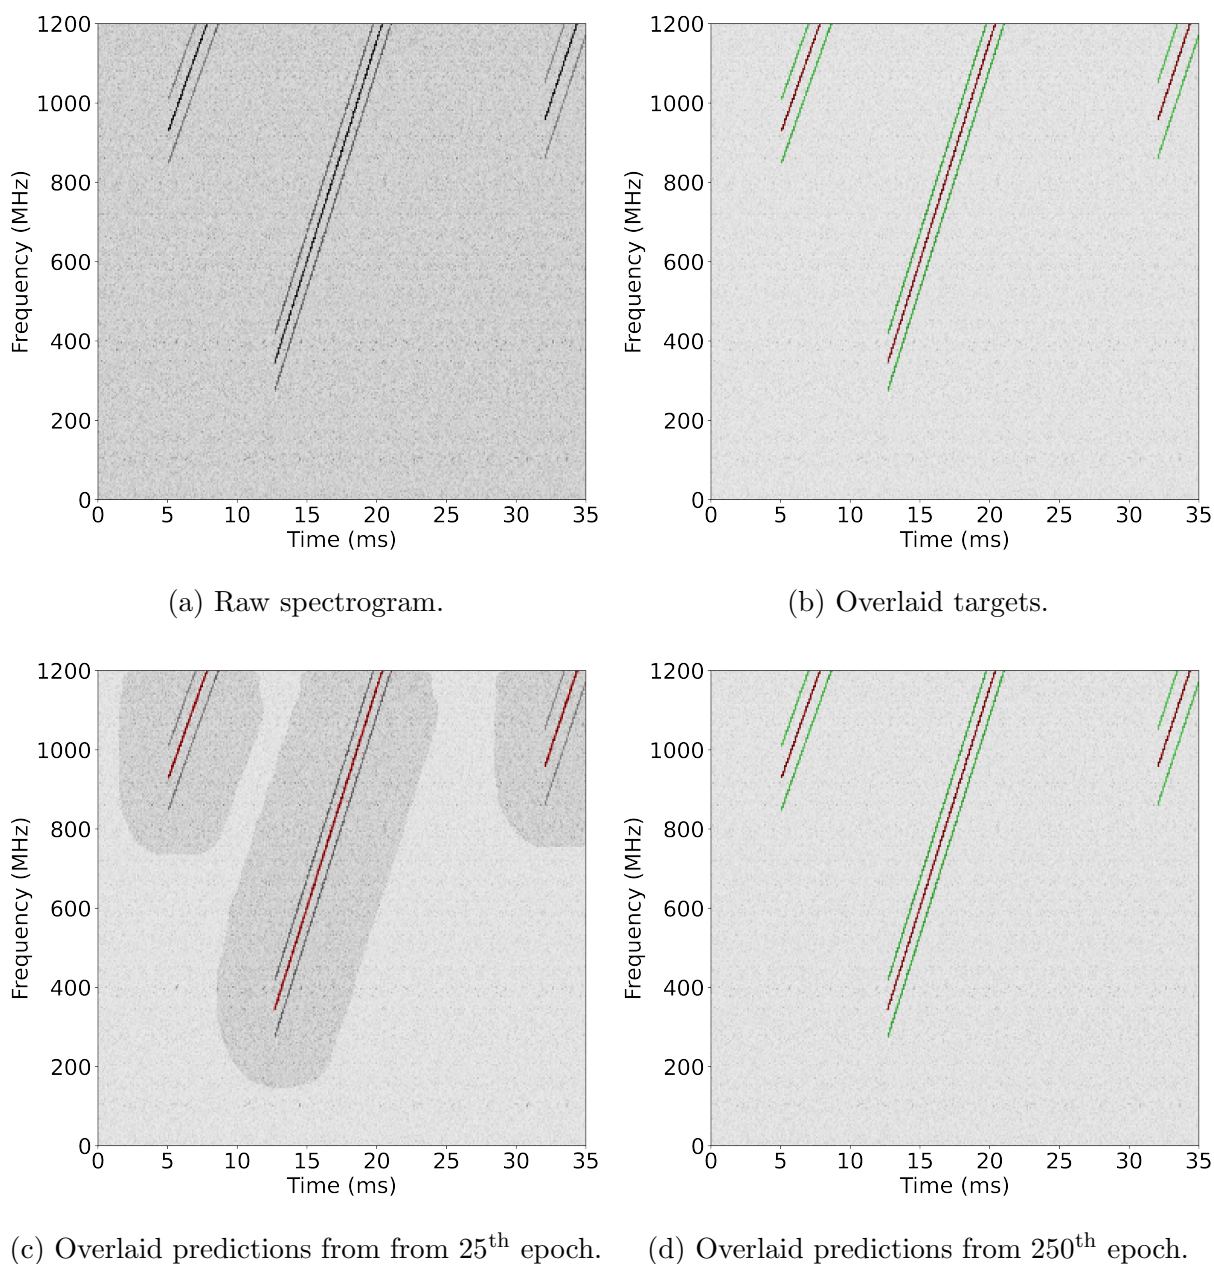


Figure 7.8: Track identification performance on a 35 ms simulated spectrogram from the test set (not seen in training). Simulated tracks have fixed SNR and a single set of side-bands. In both the targets and predictions the class mapping is: white = class 0 = background, red = class 1 = main band, green = class 2 = sideband. (a) Raw 35 ms spectrogram from test set. (b) Spectrogram with overlaid targets ($\alpha = .4$). (c) Spectrogram with overlaid predictions from the 25th epoch of training ($\alpha = .4$). (d) Spectrogram with overlaid predictions from the 250th epoch of training ($\alpha = .4$).

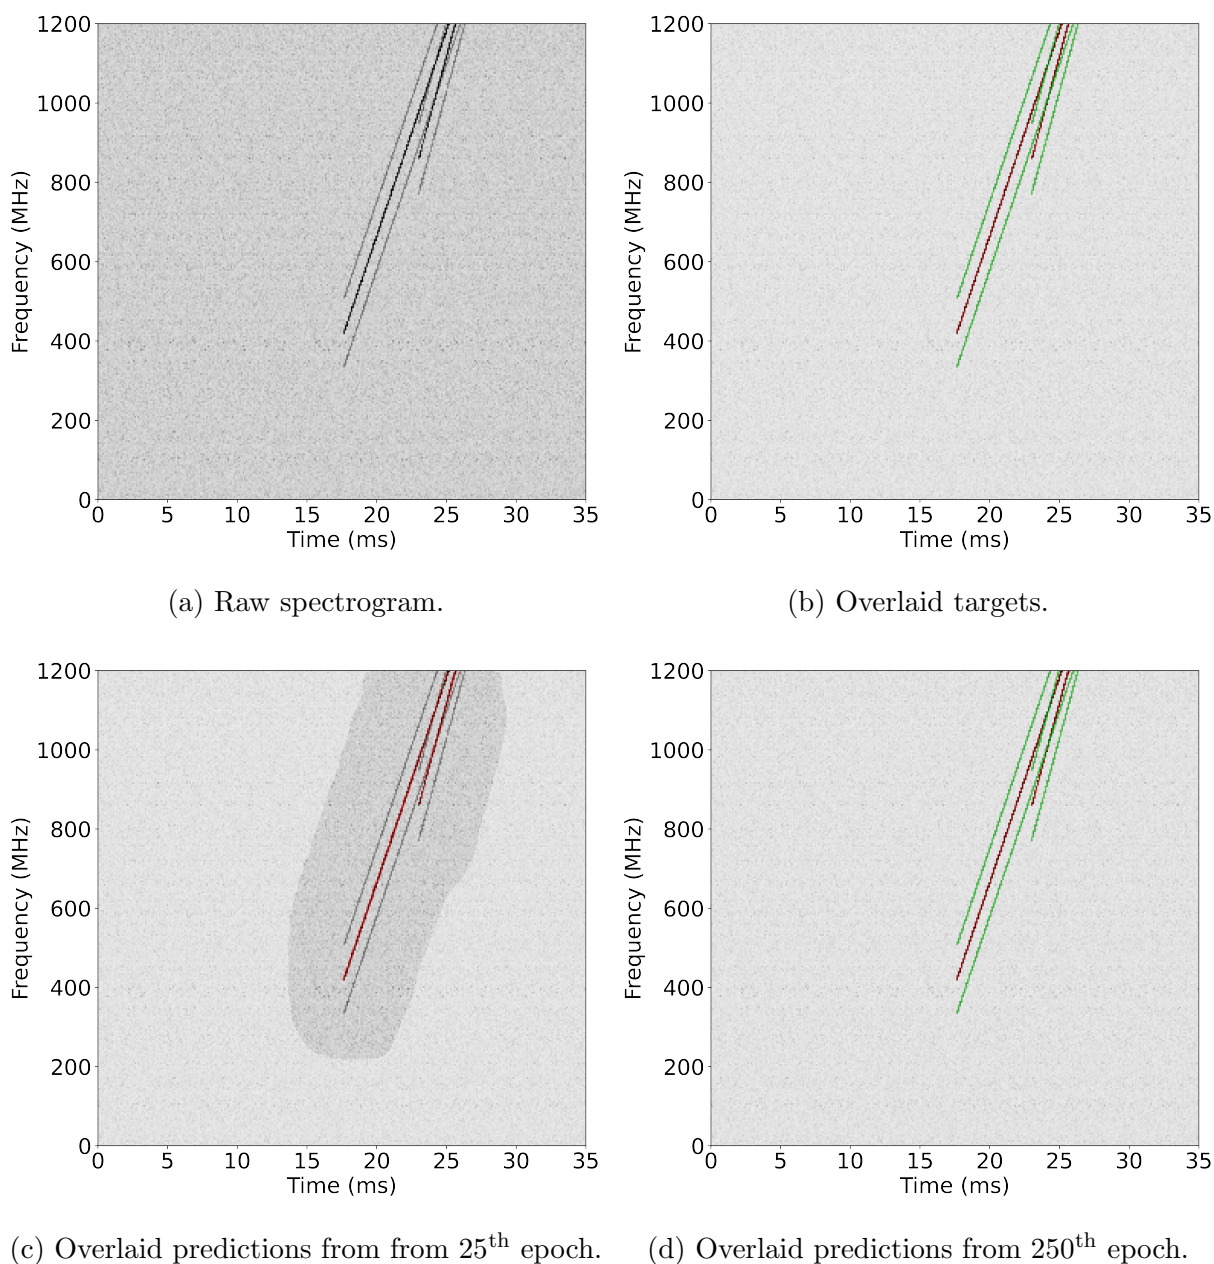



Figure 7.9: Track identification performance on a 35 ms simulated spectrogram from the test set (not seen during training) with overlapping tracks. Simulated tracks have fixed SNR and a single set of side-bands. In both the targets and predictions the class mapping is: white = class 0 = background, red = class 1 = main band, green = class 2 = sideband. (a) Raw 35 ms spectrogram from test set. (b) Spectrogram with overlaid targets ($\alpha = .4$). (c) Spectrogram with overlaid predictions from the 25th epoch of training ($\alpha = .4$). (d) Spectrogram with overlaid predictions from the 250th epoch of training ($\alpha = .4$).



(a) Raw spectrogram.

(b) Overlaid targets.


(c) Overlaid predictions from from 25th epoch.(d) Overlaid predictions from from 250th epoch.

Figure 7.10: Track identification performance on a 35 ms simulated spectrogram from the test set (not seen in training) with overlapping tracks; zoomed in. Simulated tracks have fixed SNR and no side-bands. In both the targets and predictions the class mapping is: white = class 0 = background, red = class 1 = track. (a) Raw 35 ms spectrogram from test set. (b) Spectrogram with overlaid targets ($\alpha = .4$). (c) Spectrogram with overlaid predictions from the 25th epoch of training ($\alpha = .4$). (d) Spectrogram with overlaid predictions from the 250th epoch of training ($\alpha = .4$).

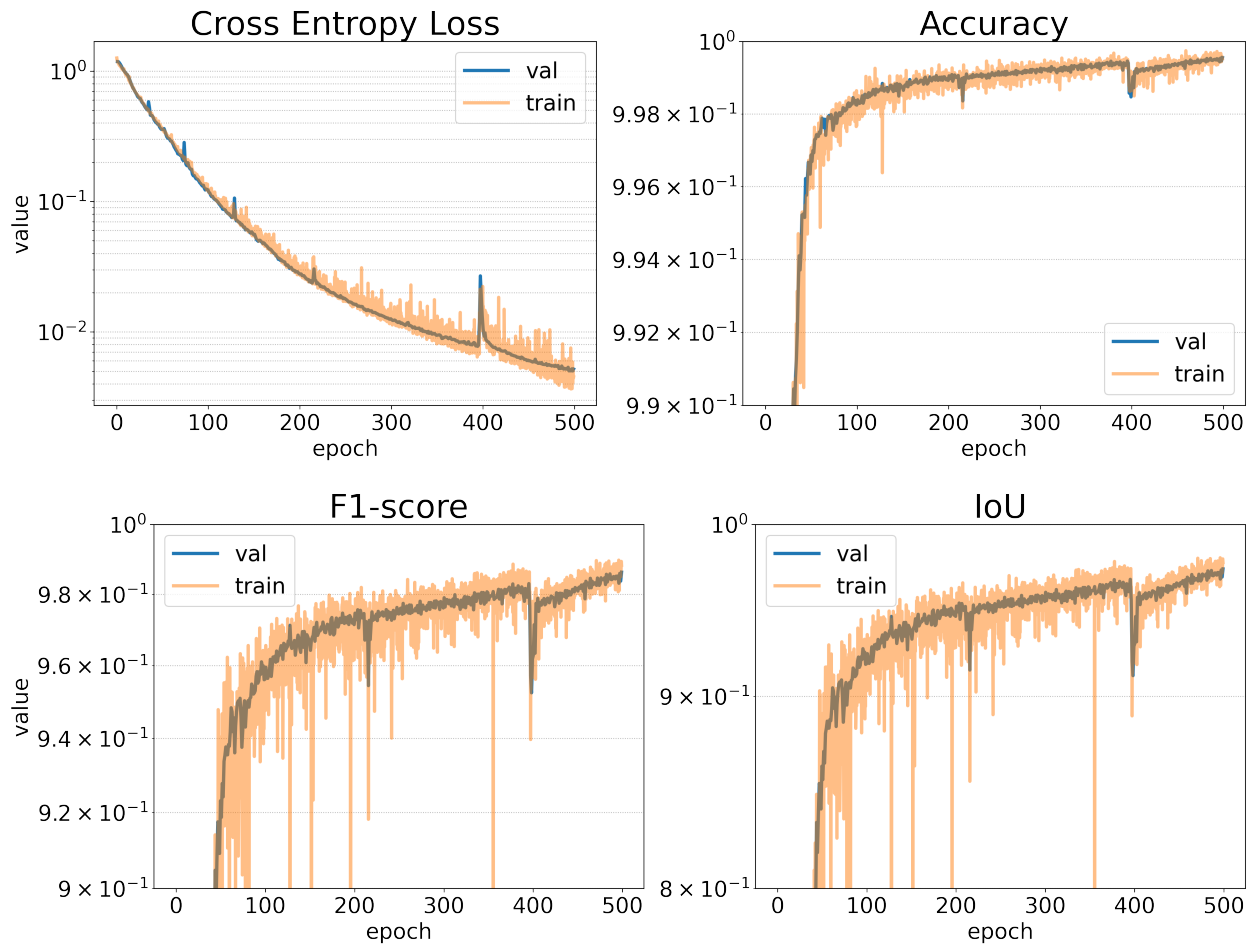


Figure 7.11: Summary of model performance for study of tracks with constant SNR and a single set of sidebands.




Figure 7.12: Example of Mask-RCNN applied to CRES data with scattering.

Chapter 8

CONCLUSIONS

8.1 *Experimental Summary*

The He6-CRES collaboration is working towards competitive sensitivity to tensor currents in the weak interaction via a precise determination of b_{Fierz} . This thesis presents much of the work done to complete *Phase 1* of the He6-CRES experiment whose overarching goal was to demonstrate that CRES can plausibly be used to achieve a 10^{-3} measurement of b_{Fierz} via the ratio of ${}^6\text{He}$ and ${}^{19}\text{Ne}$ spectra and to discover what experimental design, analysis, or simulation bottlenecks may hinder such a measurement. In particular, this dissertation presents the work done to complete the hardware, software, and experimental goals outlined in Section 3.1.

To summarize, we have built the second-ever functioning CRES apparatus and observed initial CRES signals from ${}^{83\text{m}}\text{Kr}$ (see Chapter 4). We present the first-ever observations of $\mathcal{O}(\text{MeV})$ β^\pm 's via CRES with ${}^6\text{He}$ and ${}^{19}\text{Ne}$ (see Chapter 4). And finally, we present a campaign for an $\mathcal{O}(10^4)$ counts per isotope ratio measurement of ${}^6\text{He}$ and ${}^{19}\text{Ne}$ allowing for a proof-of-principle measurement of b_{Fierz} with $\sim 10^{-1}$ statistical sensitivity (see Chapter 6).

8.2 *Towards a 10^{-3} determination of b_{Fierz} with CRES*

Taking the ratio of ${}^6\text{He}$ and ${}^{19}\text{Ne}$ provides an immense suppression of systematic uncertainties via the cancellation of energy-dependent efficiencies. Despite this, moving towards a 10^{-3} determination of b_{Fierz} with CRES via the ratio technique will require significant upgrades relative to the *Phase 1* iteration of the experiment presented in this thesis. The He6-CRES collaboration is working step-wise towards the goal of a 10^{-3} determination of b_{Fierz} , with the next goalpost being a $\sim 10^{-2}$ determination of b_{Fierz} . Here we briefly describe the hardware,

Monte Carlo, and analysis improvements necessary to attain this next level of sensitivity.

1. Hardware:

- (a) **Reduced residual gas pressures:** Low energy β^\pm experience significant scattering. Ideally we would help mitigate this effect by having significantly lower residual gas pressures in the decay cell.
- (b) **Increased frequency bandwidth:** A larger frequency bandwidth would serve the dual purpose of increasing event counts for a given source intensity and enabling an easier time constructing overlapping sections of spectra to verify the functioning of the source intensity normalization.
- (c) **$\mathbf{E} \times \mathbf{B}$ trap emptying:** A trap emptying mechanism relying on having a decay cell that is made of two electrically isolated half-cylinders and applying a voltage across them in order to empty the trap in $\mathcal{O}(100 \mu\text{s})$ is in development. Note that since the gap between the two pieces is much less than the wavelength of the radiation, the RF properties are left relatively unaffected. This will fully resolve the eddy current issue in that the trap coil will no longer be slewed off to empty the trap.
- (d) **Increased β -monitor stability:** The current β -monitor is stable to within $\sim 10^{-3}$ over 24 hours. However, to obtain the next level of sensitivity on b_{Fierz} a longer data run will be necessary (likely around 50 hours per isotope assuming other data taking processes are sped up) so we will need to be certain the monitor is sufficiently stable over that time period.
- (e) **Lighter weight spec files:** Currently each second of data is 0.6 GB, meaning that the ratio measurement presented here entailed $.6 \text{ GB} \times 500 \text{ seconds} \times 11 \text{ fields} \times 2 \text{ isotopes} = 6.6 \text{ TB}$ of data written to disk. Using a conservative thresholding of the spectrogram data and then applying zero-suppression to the data

before it's written to disk results in a ≈ 100 reduction in data written to disk and is currently being developed.

- (f) **Flatter SNR response:** The fact that the SNR oscillates with respect to frequency makes vetoing events from below difficult and causes lower SNR events to be falsely detected as two or more distinct events. Flattening the SNR response of the detector by reducing RF reflections would significantly simplify analysis.

2. Monte Carlo:

- (a) **Detailed particle-level simulation:** Further developing the particle-level Monte Carlo presented in Section 5.3 will be necessary to make precise comparisons between observation and expectation. This will include correctly simulating the (potentially still oscillating) SNR response of the detector and inputting realistic simulated `spec` files through the entire analysis pipeline so as to properly account for any effects lurking there. Scaling a Monte Carlo such as this to have much better statistics than the desired experimental dataset will be a significant challenge even with smaller data files.

3. Analysis:

- (a) **Better event reconstruction at low fields:** As was illustrated in Chapter 5, the best sensitivity to b_{Fierz} is obtained when one measures the ratio over a wide breadth of energies. To properly resolve events at low energies, where scattering will be present even with reductions in residual gas pressures, a different analysis approach will need to be applied there that properly stitches together scattered tracks into events.

With these relatively minimal improvements He6-CRES should be able to attain a 10^{-2} determination of b_{Fierz} in the near future.

8.3 *Final remarks*

Traditional spectroscopy techniques using solid state detectors that rely on β energy loss in matter contend with percent-level corrections due to β (back-) scattering and bremsstrahlung losses. Spectral measurements using conventional magnetic spectrometers also have $\mathcal{O}(1\%)$ corrections due to scattering in collimators and other effects related to the distance β s travel before detection [116]. Utilizing CRES for the determination of $\mathcal{O}(\text{MeV})$ -endpoint β -energy spectra is a novel approach to precise β spectroscopy with an entirely new set of considerations and systematics many of which were first encountered via the work presented in this thesis. A new experimental technique will always have challenges to overcome but the outlook of a competitive 10^{-3} determination of b_{Fierz} via CRES in the next 5-6 years is promising.

Appendix A

CHIRP SNR OPTIMIZATION

A.1 The concept

A.1.1 A balance




Figure A.1: Imagine the Track Slope expressed in units of Freq Bin/Slice.

In order to maximize SNR a balance between bin width and slice length, which are inversely related to one another, must be struck. If the slices are long relative to the CRES track slope then the track leaves the bin before the end of the FFT, which decreases the power in that bin. If the frequency bin width is too wide you are adding noise power unnecessarily because the noise power per bin scales with the bin width. See Fig. A.1 for a visualization of this. In order to optimize SNR we want the slope to be unity in units of

frequency bin width divided by time slice length.

A.1.2 A more precise model

In order to construct a more precise model of the effect of the bin structure on the SNR of a track we need to consider the average SNR over a track and we also need to consider tracks that don't start in the bottom left corner of a bin.

We define the SNR of a given bin to be the signal power in that bin over the noise power in that bin. The mean SNR of such a track will be:

$$SNR_{Mean} = \underbrace{\frac{Power * \# slices}{N}}_{Signal\ per\ Bin} \times \underbrace{\frac{1}{kT_e \Delta f}}_{1/Noise\ per\ Bin} \quad (A.1)$$

Where N is the total number of bins that the track traverses through, Δf is the size of a frequency bin, and T_e is the equivalent noise temperature of the entire RF system. Since the power and equivalent noise temperature T_e are constant for a given track and the time duration of a track is proportional to b , the number of slices it passes through, we can reduce our expression to:

$$SNR_{Mean} = \underbrace{\frac{Power * b}{N}}_{Signal} \times \underbrace{\frac{1}{kT_e \Delta f}}_{1/Noise} = \underbrace{\frac{Power}{kT_e}}_{physics} \times \underbrace{\frac{b}{N \Delta f}}_{daq} \quad (A.2)$$

To gain some intuition let's first consider a track that starts and ends in opposite corners of an $a \times b$ grid. See Fig. A.2. For such a track N , the total number of bins the track passes through will be:

$$N = a + b - GCD(a, b) \quad (A.3)$$

Where $GCD(a,b)$ is the greatest common divisor of a and b . This formula can be understood when you consider that there are a horizontal lines and b vertical lines that must

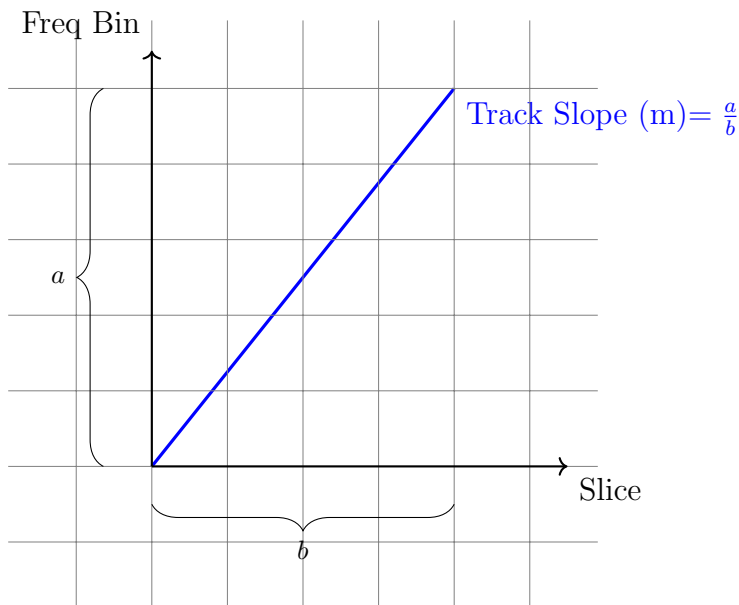


Figure A.2: A track traversing an $a \times b$ grid.

be crossed by the track as it crosses the grid and for each line it crosses it enters a new bin. However, whenever it crosses a grid vertex it crosses a vertical and horizontal line simultaneously but only enters one bin, so you are over counting by one every time you cross a vertex. And the number of vertices you cross (including the very last one in the top right corner) is equal to the greatest common divisor of a and b .

In the more general case where the track starts in an arbitrary position on the frequency axis (for ease let's consider only tracks that begin and end on a vertical grid line) we add in one more bin in the top right corner. And for physical tracks, the GCD of a and b will be 1. So for a track starting at an arbitrary frequency, that starts and stops on vertical grid lines:

$$N = a + b = b(m + 1) \tag{A.4}$$

For a track slope of $m = \frac{a}{b}$. Note that including tracks that start at arbitrary times will only increase N by 1 or 2. Let us consider the “daq” part of Equation A.7 because we would like to consider how to optimize the daq for SNR. Note that this isn't a unitless quantity

(it has units of seconds) but we simply multiply by the “physics” part of that equation to obtain a unitless SNR. We see that

$$SNR_{Mean} \propto \frac{b}{N\Delta f} = \frac{b}{\Delta f b(m+1)} = \frac{1}{\Delta f(m+1)} \quad (\text{A.5})$$

Now consider that m here is expressed in the units of frequency bins per slice (unitless) which is the physical slope divided by Δf^2 (because an FFT requires $\Delta t = \frac{1}{\Delta f}$).

$$SNR_{Mean} \propto \frac{1}{\Delta f \left(\frac{m_{Physical}}{\Delta f^2} + 1 \right)} = \left(\frac{m_{Physical}}{\Delta f} + \Delta f \right)^{-1} \quad (\text{A.6})$$

To find the Δf associated with the max mean SNR for a given physical slope $m_{Physical}$ we take the derivative of the above expression with respect to Δf and set it equal to zero. We get:

$$\frac{\partial(SNR_{Mean})}{\partial \Delta f} = 0 = - \left(\frac{m_{Physical}}{\Delta f} + \Delta f \right)^{-2} \left(\frac{-m_{Physical}}{\Delta f^2} + 1 \right) \rightarrow \boxed{\Delta f = \sqrt{m_{Physical}}} \quad (\text{A.7})$$

Which is the same result you get based on intuition but now we have a function to work with.

A.1.3 Is this the best we can do?

Is it true that $\Delta f = \sqrt{m_{Physical}}$ corresponds to the best possible SNR obtainable for a fixed track slope, track power, and noise floor? There are two knobs one can turn on the daq: the points per FFT (N) and the sampling rate (f_s). The choice of these fully determine Δf , Δt , and the total measurement bandwidth (BW).

$$\Delta f = \frac{f_s}{N} \quad (\text{A.8})$$

$$\Delta t = \frac{N}{f_s} = \frac{1}{\Delta f} \quad (\text{A.9})$$

$$BW = \frac{f_s}{2} \tag{A.10}$$

One can see here that the ratio $\frac{f_s}{N} = \Delta f$ is what ultimately determines the SNR of a track. You can't modify the daq in such a way to beat the optimal SNR achieved by $\Delta f = \sqrt{m_{Physical}}$ because you can't escape the fact that Δf and Δt are always inversely related. What you can do is increase the points per FFT (N) and the sampling rate (f_s) such that the ratio $\frac{f_s}{N}$ is maintained. This will result in an increased experimental BW without any loss of SNR.

Appendix B

KAPTON ADSORPTION OF HELIUM

Reference to this appendix is in ch 2, sec 2.3 (cres apparatus) FOR THIS APPENDIX (WORK ON LATER): Appendix A: Kapton absorption of helium Take main stuff from here: <https://www.overleaf.com/project/5e1e1310be536d0001442cbb>, And also include just the experimental stuff from here: <https://www.overleaf.com/project/5cc0a036f5be224069f3e321>
 MENTION THAT THIS SHOULD GENERALIZE TO NEON AS WELL.

B.1 Background

B.1.1 General governing equations

In order to address the issue of ${}^6\text{He}$'s absorption into Kapton we first need to present the theory that governs the diffusion of noble gases through a permeable thin foil of thickness l . The permeability K establishes the net flow rate Q through a solid and is given by

$$Q = K \frac{A}{l} \Delta P \quad (\text{B.1})$$

for a film of thickness l and cross-sectional area A where ΔP is the change in pressure across the film. It is important to note that Equation B.1 is only valid in equilibrium, when the concentration profile in the film is constant in time. The permeability K can be considered as a product of the diffusivity D and solubility b .

$$K = D b \quad (\text{B.2})$$

The diffusivity determines the rate at which a gas flows in a material. The solubility defines the concentration of gas in the solid at a given partial pressure (or, equivalently, the chemical potential of the gas in the solid relative to it in vacuum)[9].

In order to measure K experimentally one can measure an equilibrium state flow Q in conjunction with Equation B.1. In order to measure D experimentally one must observe the time evolution of gas flow after establishing a pressure gradient across the film. Solubility b can be obtained via measurements of K and D in conjunction with Equation B.2.

In presenting the equations that govern the time evolution of the gas flow we follow the notation of Reference [9]. We begin by assuming no decay of the diffusing atoms. The diffusion current of a substance with concentration (number of atoms per unit volume in 3d or number of atoms per unit distance in 1d) C is given by:

$$J = -D \nabla C \quad (\text{B.3})$$

where we implicitly assume that C is a function of position and time and D stands for the diffusion coefficient. The conservation of atoms requires that the concentration at any given point satisfies:

$$\partial C / \partial t = -\nabla \cdot J. \quad (\text{B.4})$$

These two equations yield ‘‘Fick’s Equation’’, which in one dimension is:

$$\frac{\partial C(x, t)}{\partial t} = D \frac{\partial^2 C(x, t)}{\partial x^2} \quad (\text{B.5})$$

where we now show the explicit dependence on x and t . To fix ideas, the dimensions of the variables and quantities in 1d are:

$$\begin{aligned} [C(x, t)] &\equiv \frac{1}{\text{cm}^1} \\ [D] &\equiv \frac{\text{cm}^2}{\text{s}} \\ [J(x, t)] &\equiv \frac{1}{\text{s}}. \end{aligned}$$

The solution to Fick’s equation, which describes the concentration as a function of position and time expressed in terms of the concentration at $t = 0$, $f(x')$, is given by[9]:

$$C(x, t) = \frac{1}{2} \sqrt{\frac{1}{\pi D t}} \int_{-\infty}^{\infty} dx' f(x') e^{-\frac{(x'-x)^2}{4Dt}}. \quad (\text{B.6})$$

A detailed derivation of this solution is given in Appendix B.4.

B.1.2 Finite lifetime effect on diffusion

In order to take into account the decay of ${}^6\text{He}$ we add a decay term to Equation B.4:

$$\frac{\partial C(x, t)}{\partial t} = D \frac{\partial^2 C(x, t)}{\partial x^2} - \lambda C(x, t) \quad (\text{B.7})$$

and use as a trial function the same structure as Equation B.5, but this time with:

$$\phi(x) = A \cos kx + B \sin kx \quad (\text{B.8})$$

$$F(t) = \alpha \exp\{-(k^2 D + \lambda) t\}. \quad (\text{B.9})$$

As above, one can verify:

$$\frac{\partial C(x, t)}{\partial t} = -(k^2 D + \lambda) C(x, t) \quad (\text{B.10})$$

$$\frac{\partial^2 C(x, t)}{\partial x^2} = -k^2 C(x, t). \quad (\text{B.11})$$

The proposed trial function is thus a solution of the *modified* Fick's equation, Equation B.7.

The solutions are just like for the case without decay with the exception of the modification:

$$k^2 D \rightarrow k^2 D + \lambda. \quad (\text{B.12})$$

The consequent modification of Equation B.6 is:

$$C(x, t) = \frac{1}{\pi} \int_{-\infty}^{\infty} dx' f(x') \int_0^{\infty} dk e^{-(k^2 D + \lambda)t} \cos k(x - x') = e^{-\lambda t} \times C_{\lambda=0}(x, t) \quad (\text{B.13})$$

where $C_{\lambda=0}(x, t)$ is the solution with $\lambda = 0$.

In summary, the effect of a finite lifetime is simply to multiply the solution for $\lambda = 0$ by the factor $e^{-\lambda t}$. It effectively increases the diffusivity [Equation B.12]. Thus, the distribution profiles are not affected.

B.1.3 Modeling our experimental configuration

Our own situation can be modeled most easily as a volume V , at temperature T , with a constant flow R of ${}^6\text{He}$ particles ($\tau \approx 1s$) into the volume. An area A of the volume wall is a film of width ℓ and permeability K and the rest of the volume is impenetrable. Our initial conditions are that the volume and film are at zero concentration/pressure and that at $t = 0$ the in-flow R is turned on. The space on the other side of the film is kept at zero pressure. So long as equilibrium is reached in seconds or less all we would like to know is the steady state pressure P_{steady} inside the volume.

Only toy model of the time evolution of the situation is presented but the steady state pressure found should be correct.

At first we take the flow into the film as negligible because the flow into the film is a function of the pressure on the film which is initially zero [Equation B.43]. In this case we have:

$$\begin{aligned}\frac{dN}{dt} &= -N\lambda + R \\ \frac{V}{kT} \frac{dP}{dt} &= -P \frac{V}{kT} \lambda + R\end{aligned}\tag{B.14}$$

where the ideal gas law is used, $\lambda = \frac{1}{\tau}$, and k is the Boltzmann constant. Solving this separable differential equation we find

$$P = \frac{RkT}{V\lambda}(1 - e^{-\lambda t}).\tag{B.15}$$

So on the order of $\tau \approx 1s$ an equilibrium pressure of $P = \frac{RkT}{V\lambda}$ is reached.

Now in a second or so ($\tau_b = .56s$) an equilibrium flow Q is established into the film. A new equilibrium pressure in the volume is reached, again on the order of the lifetime. This new equilibrium condition is:

$$\frac{dN}{dt} = 0 = -N\lambda + R - Q\tag{B.16}$$

Where Q in steady state is given by [Equation B.1]:

$$Q = \Delta P K \frac{A}{\ell} \quad (\text{B.17})$$

Where the ΔP represents the pressure change over the film. Our final equilibrium pressure is:

$$\begin{aligned} P &= \frac{(R - Q) kT}{V \lambda} \\ P &= \left(\frac{R kT}{V \lambda} \right) \frac{1}{1 + \frac{kT K A}{V \lambda \ell}} \\ P &= P_0 \frac{1}{1 + \frac{kT K A}{V \lambda \ell}} \end{aligned} \quad (\text{B.18})$$

With P_0 being the pressure in the volume given no permeation into the film. Let us assume that Schowalter et al [101] are correct (at least in order of magnitude) in their measurement of the permeability K of helium into Kapton. Then the effect of the Kapton on the steady state pressure of our system P_{steady} is a 10^{-7} effect.

B.2 Literature review

With knowledge of the solubility b and diffusivity D of ^4He , and an understanding of how a finite half-life affects diffusion dynamics, we would be able to calculate the pressure of ^6He in our decay volume as a function of the flux entering the volume.

B.2.1 Harvard-group determination of the Diffusivity of ^4He

In search of data on the permeability b of ^4He through Kapton at pressures and temperatures relevant to our experiment, we found only Schowalter et al [101]. The group measures K and D and uses the relation $K = D b$ to come up with an indirect measurement of the solubility b . The group measures the diffusivity D by placing a constant high pressure (about 1 atm.)

of gas on one side of the Kapton foil and using a Residual Gas Analyzer (RGA) to measure the partial pressure of helium on the low pressure (initially evacuated) side.

Figure B.1 shows the setup and Figure B.2 shows a sample of their data. In Equation (3)

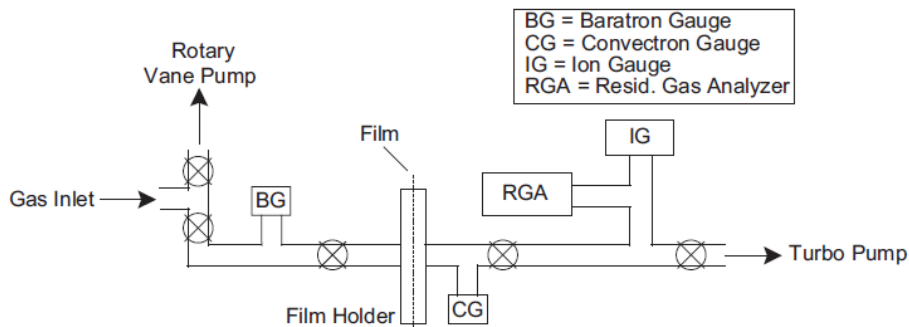


Figure B.1: Figure 1 of Schowalter et al (Reference [101]).

of Reference [101] the authors show what they used for calculating the gas flow versus time, $Q(t)$,

$$Q(t) = Q_0 \left(1 + 2 \sum_1^{\infty} e^{-(n\pi/\ell)^2 D t} (-1)^n \right). \quad (\text{B.19})$$

There is an apparent typo in the Harvard-group paper: where the paper reads $\ell^2 t/D$ we have used $D t/\ell^2$ [9]. We use ℓ for the membrane thickness, which is labeled d in the Harvard-group paper. In Appendix B.5 we present the solution to Barrer’s Equation particular to the boundary condition that the left side of the membrane is held at constant pressure and the right side is held at zero pressure. Our solution (Equation B.43) corresponds exactly to Equation (3) of Reference [101] (Equation B.19 in this work).

In Reference [101] they claim that the “diffusivity D is determined by fitting the solution to the one-dimensional diffusion equation (Equation (3)) to the partial pressure data shown in Figure 2.” This data is shown in Figure B.2 here (for Argon). However, (Equation (3)) is not the correct equation to use for partial pressures. In the limit of large t in (Equation (3)) a constant flow rate Q_0 is reached. This corresponds to a linear increase in pressure in the




Fig. 2. (Color online) Sample data used to determine permeability and diffusivity. At $t=0$, gas is introduced to the high-pressure chamber and allowed to come in contact with the film. Due to the pressure difference across the film, the gas begins to permeate the film. The argon gas partial pressure rises asymptotically to a steady-state pressure after a characteristic breakthrough time, t_b , defined in Eq. (4).

Figure B.2: Figure 2 of Schowalter et al (Reference [101]).

low pressure chamber not to an equilibrium pressure as they find. The issue does not seem like a typo in their label, since, in addition to the label on the figure, the end of the figure caption reads “The argon gas pressure rises asymptotically to a steady-state pressure after a characteristic breakthrough time, t_b ”. The characteristic time t_b that they are identifying with the time it takes the system to reach an equilibrium pressure is defined in the literature [9] as the t -intercept of the pressure in the linear limit [Equation B.46]:

$$t_b = \frac{\ell^2}{6D} \quad (\text{B.20})$$

The Harvard-group is pumping the low pressure side of the membrane which results in this steady-state pressure. The time to reach a steady state pressure (which they identify with Equation B.20) is a function of this pumping rate, not the diffusion D , as they claim.

The time evolution is governed by (derivation in Appendix B.6):

$$\frac{P_1}{P_0} = \frac{kTQ}{P_0 S} - \left(\frac{kTQ}{P_0 S} - 1 \right) e^{-St/V} \quad (\text{B.21})$$

Where P_1 is the pressure on the right side of the membrane, P_0 is the fixed pressure on the left, S is the pumping speed on the right in m^3/s , Q is the diffusion rate into the right volume in $atoms/s$, T is the temperature, and k is the Boltzmann constant. We see that the time to reach the steady-state pressure is only a function of the pumping speed S and the volume V .

The fact that Harvard-Group uses a function describing flow to fit to pressure data, and that they seem to misidentify t_b leads to the conclusion that the diffusion coefficient extracted via their method is incorrect. The solubility b of He that they quote is a function of this diffusivity D so is also of questionable reliability.

B.2.2 Motivating our measurement

There is one other group that has investigated the absorption of helium into Kapton at pressures relevant to our purposes that was conducted by A. Eggenberger et al [29]. However they only found the permeability K , not the diffusivity D or the solubility b of helium through Kapton. Given the lack of reliable measurements of the parameters we need, we have conducted our own experimental investigation of the absorption of ${}^6\text{He}$ into Kapton.

B.3 Experimental determination of ${}^6\text{He}$ absorption into kapton

B.3.1 Methods

We have designed an experiment to quantify the absorption of ${}^6\text{He}$ into Kapton. We will measure the apparent mean lifetime of ${}^6\text{He}$ in a decay volume with its inner surface area lined with Kapton film. Because the Kapton is at a distance from the two parallel scintillating panels that detect decay electrons, ${}^6\text{He}$ atoms that decay after being absorbed onto the

Kapton film will not be detected. This will lead to a decrease in measured (apparent) lifetime.

The time derivative of the number of particles in the volume without Kapton in the volume N is given by:

$$\frac{dN}{dt} = -\frac{N}{\tau} \quad (\text{B.22})$$

Where τ is the mean lifetime of ${}^6\text{He}$. The time derivative of the number of particles in the volume with Kapton in the volume N' is given by:

$$\frac{dN'}{dt} = -\frac{N'}{\tau} - P_{abs} \frac{A}{V} v_z N' \quad (\text{B.23})$$

where v_z is the velocity of the gas perpendicular to the film which is given by the temperature of the gas, A/V is the ratio of surface area (of Kapton) to volume of the decay cell, and P_{abs} is the probability of absorption upon collision with the film. Thus the relative difference between this apparent lifetime and the lifetime measured without Kapton in the volume is given by:

$$\frac{\tau' - \tau}{\tau} \approx P_{abs} \tau v_z \frac{A}{V} \quad (\text{B.24})$$

In May 2018 lifetime data was taken with Kapton in the decay volume and without. Each run was made up of 100-200 14-s cycles. A single 14-s cycle was comprised of 1 s of pumping to clear the volume, 3 s of ${}^6\text{He}$ filling, and 10 s in which the volume was closed and decay data was taken. The histogram of the decay spectra was recorded cumulatively over the course of one run, meaning that a given time bin contains the total counts that occurred at that time over all cycles in the run. Data from a run is shown in Figure B.3.

The decision to histogram cumulatively made the analysis challenging. There was significant deadtime and variability in ${}^6\text{He}$ production between cycles within a given run. The

inconsistency in lifetime found between runs with the same run parameters (either with Kapton or without) was of the same order as the difference across runs with different run parameters. The experiment was inconclusive. However, it did indicate that the effect of the absorption cannot be large.

B.4 Solution to Barrer's equation

We start with "Fick's Law" in one dimension:

$$\frac{\partial C(x, t)}{\partial t} = D \frac{\partial^2 C(x, t)}{\partial x^2} \quad (\text{B.25})$$

The solution can be expressed as a product of a position-dependent part and a time-dependent part:

$$C(x, t) = \phi(x) F(t) \quad (\text{B.26})$$

with

$$\phi(x) = A \cos kx + B \sin kx \quad (\text{B.27})$$

$$F(t) = \alpha \exp\{-k^2 D t\} \quad (\text{B.28})$$

with k a real positive number. One can verify:

$$\frac{\partial C(x, t)}{\partial t} = -k^2 D C(x, t) \quad (\text{B.29})$$

$$\frac{\partial^2 C(x, t)}{\partial x^2} = -k^2 C(x, t) \quad (\text{B.30})$$

confirming that the proposed trial function is indeed a solution.

Because any value of k yields a possible solution, the sum of terms of the form of Equation B.27 is also a solution. Since k can take any positive real value, the solutions can be expressed as:

$$C(x, t) = \int_0^\infty dk [g(k) \cos kx + h(k) \sin kx] e^{-k^2 D t} \quad (\text{B.31})$$

where $g(k)$ and $h(k)$ can be determined by boundary conditions at $t = 0$ via:

$$g(k) = \frac{1}{\pi} \int_{-\infty}^{\infty} dx' f(x') \cos k x' \quad (\text{B.32})$$

$$h(k) = \frac{1}{\pi} \int_{-\infty}^{\infty} dx' f(x') \sin k x'. \quad (\text{B.33})$$

Here $f(x')$ represents the concentration profile at $t = 0$. Replacing into Equation B.31 the solution can be expressed as:

$$C(x, t) = \frac{1}{\pi} \int_0^{\infty} dk e^{-k^2 D t} \int_{-\infty}^{\infty} dx' f(x') \cos k(x - x') \quad (\text{B.34})$$

Inverting the order of the integrals:

$$C(x, t) = \frac{1}{\pi} \int_{-\infty}^{\infty} dx' f(x') \int_0^{\infty} dk e^{-k^2 D t} \cos k(x - x'). \quad (\text{B.35})$$

From a table of integrals we obtain:

$$\int_0^{\infty} dk e^{-k^2 D t} \cos k(x - x') = \frac{1}{2} \sqrt{\frac{\pi}{D t}} e^{-\frac{(x'-x)^2}{4 D t}}. \quad (\text{B.36})$$

So, the final equation for the concentration as a function of position and time, $C(x, t)$, expressed in terms of the concentration at $t = 0$, $f(x')$, is given by:

$$C(x, t) = \frac{1}{2} \sqrt{\frac{1}{\pi D t}} \int_{-\infty}^{\infty} dx' f(x') e^{-\frac{(x'-x)^2}{4 D t}}. \quad (\text{B.37})$$

This is Barrer's Eq [9].

B.5 Particular situations: The finite solid with fixed pressure

A more relevant solution for the concentration of solute in a membrane that is in contact with a gas at constant pressure is that corresponding to the boundary conditions:

$$f(x') = \begin{cases} C_0 & \text{for } x' = 0 \text{ at all } t \\ 0 & \text{for } 0 < x' < \ell \text{ at } t = 0 \\ 0 & \text{for } x' = \ell \text{ at all } t \end{cases} \quad (\text{B.38})$$

The third condition is reasonable assuming that the right side of the membrane is always kept at low pressure. For this case, the solution is

$$C(x, t) = C_0 \left[(1 - x/\ell) - \frac{2}{\pi} \sum_1^{\infty} \frac{1}{n} e^{-(n\pi/\ell)^2 Dt} \sin \frac{n\pi x}{\ell} \right]. \quad (\text{B.39})$$

This equation fulfills the boundary conditions. Figure B.4 shows the concentration versus position at $t = 0$ and at some chosen subsequent times, expressed as a ratio of the squared of the membrane thickness divided by the diffusion coefficient, ℓ^2/D .

In a typical experiment a gas is kept at constant pressure on one side of a membrane and the pressure is monitored on the outflow volume on the opposite side. In the high-pressure side, the relation between the pressure, p_1 , and the gas concentration in the membrane is given by Henry's law [9]:

$$C_0 = p_1 S \quad (\text{B.40})$$

where S is the solubility. The equation for the concentration can be used to extract the flow of particles on the low pressure side:

$$J = -D \partial C(x = \ell, t) / \partial x. \quad (\text{B.41})$$

Then, assuming this flow is all going to the outflow volume, its pressure should change with time at a rate:

$$\frac{dp}{dt} = -\frac{A k T}{V} D \partial C(x = \ell, t) / \partial x. \quad (\text{B.42})$$

Taking into account Equations B.39, B.40 and B.42 we get:

$$\frac{dp}{dt} = p_1 \frac{A k T D S}{V \ell} \left[1 + 2 \sum_1^{\infty} e^{-(n\pi/\ell)^2 Dt} \cos n\pi \right]. \quad (\text{B.43})$$

The time integral is useful:

$$\begin{aligned}
p &= \kappa \left[t + \frac{2\ell^2}{\pi^2 D} \sum_1^{\infty} \left(1 - e^{-(n\pi/\ell)^2 D t} \right) \frac{\cos n\pi}{n^2} \right] \\
&= \kappa \left[t - \frac{\ell^2}{6 D} - \frac{2\ell^2}{\pi^2 D} \sum_1^{\infty} e^{-(n\pi/\ell)^2 D t} \frac{(-1)^n}{n^2} \right], \tag{B.44}
\end{aligned}$$

where we made use of

$$\sum_1^{\infty} \frac{(-1)^n}{n^2} = -\frac{\pi^2}{12}. \tag{B.45}$$

Equation B.44 is used to point out that, at long enough times, the pressure grows linearly with time, and a tangent to the pressure curve intersects the $p = 0$ axis at a time t_b that depends on the diffusivity:

$$t_b = \frac{\ell^2}{6 D} \tag{B.46}$$

which can be used to determine D experimentally.

B.6 Time evolution towards steady state pressure

Consider a volume V of pressure P_1 being pumped at a rate of $S \text{ m}^3/\text{s}$ with a constant input of $Q \text{ atoms}/\text{s}$. Note that in our case the rate of diffusion through the membrane Q is given by $Q = K \frac{A}{l} \Delta P$ where K is the permeability, A is the foil area, and l is the foil thickness. Thus, in the case when the pressure on the right side of the membrane P_1 is much less than the fixed high pressure on the left, ΔP (and therefor Q) is approximately constant.

Treating the helium as an ideal gas, we find the time evolution of the number of particles n in the volume is given by:

$$\frac{dn}{dt} = \frac{V}{kT} \frac{dP_1}{dt} = Q - \frac{P_1}{kT} \tag{B.47}$$

where T is the temperature, and k is the Boltzmann constant. Using separation of variables we find:

$$\frac{V dP_1}{kTQ - P_1 S} = dt \rightarrow$$
$$P_1(t) = \frac{kTQ}{S} - \left(\frac{kTQ}{S} - P_0 \right) e^{-St/V}. \quad (\text{B.48})$$

where P_0 is the initial pressure in the volume.

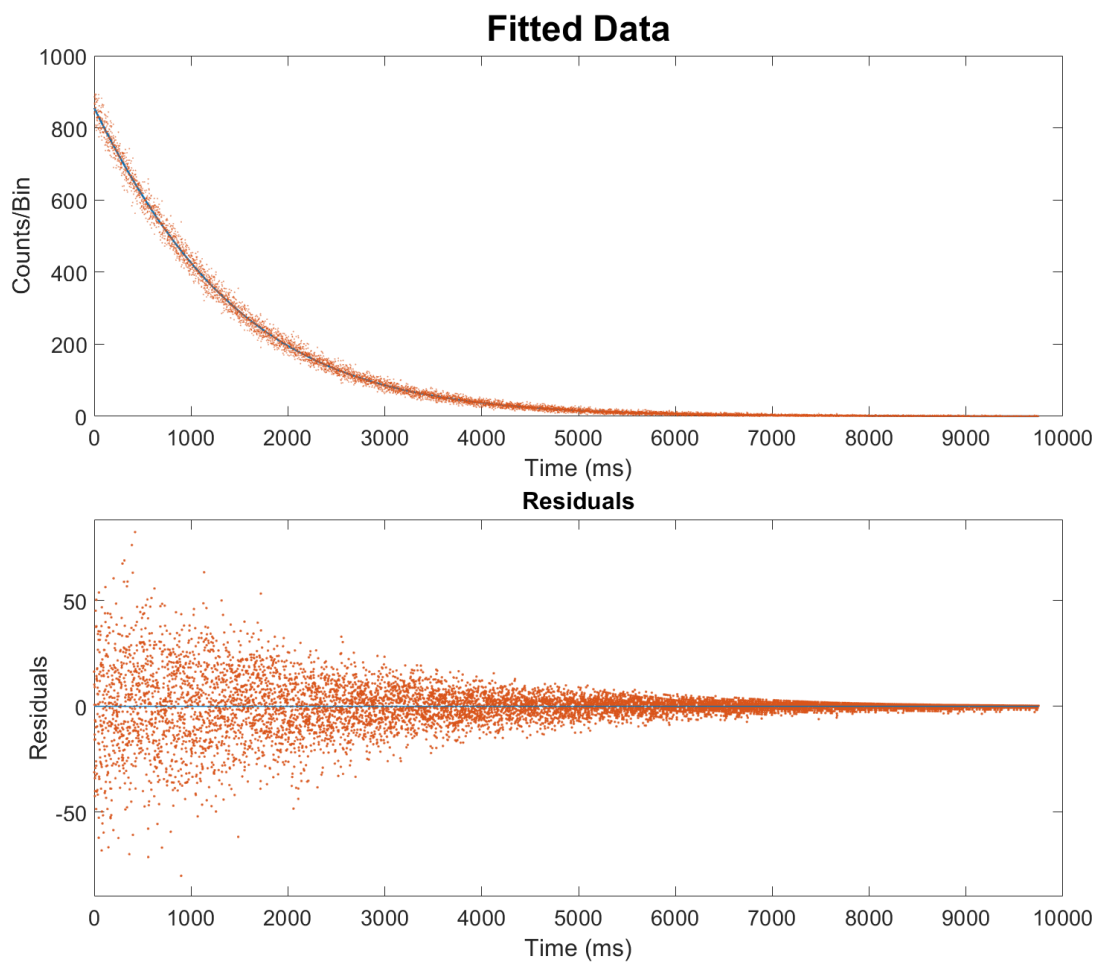


Figure B.3: A histogram of decay events (ms bin size) recorded during a run of ~ 200 14s cycles fit to a dead time effected exponential. Here the fitted half-life was $\tau_{1/2} = 812 ms$.




Figure B.4: Plot of Equation B.39 for different times: magenta: $t = 0$; black: $t = 0.01 (\ell^2/D)$; red: $t = 0.1 (\ell^2/D)$; blue: $t = (\ell^2/D)$. For these plots $\ell^2/D = 4$. The wiggles in the $t = 0$ distribution are due to taking the sum only up to $n = 200$.

Appendix C

POSTGRESQL DATABASE (`he6cres_db`) DETAILS

Details on each column of the four PostgreSQL tables (`run_log`, `spec_files`, `nmr`, and `monitor`) that are written to during CRES data acquisition by the He6-CRES DAQ-UI. See Section 3.8 for more context.

run_log	Table column	Type	Description
	run_id	INT	Unique identifier of data run.
	created_at	TIME	Date and time of row creation.
	num_spec_acq	INT	Number of acquisitions (spec files) in data run.
	analog_inputs	INT	Number of analog inputs (always 1 as of 2023).
	freq_ch	INT	Number of frequency bins (4096 typically).
	run_notes	TEXT	Notes on the run.
	roach_avg	INT	Number of slices summed by ROACH2 (2 typically).
	roach_nyquist	INT	Nyquist frequency of ROACH2 (1.2 GHz typically).
	set_field	REAL	Main field set value.
	trap_config	TEXT	Indicates set current and trap cycle structure.
	true_field	REAL	Current nmr field value.
	requant_gain	INT	Factor used by ROACH2 to put powers in 0-255 range.
	acq_length_ms	REAL	Length of each file (in ms).
	rga_*	REAL	RGA partial pressures of 11 residual gas species (Torr).
	rga_tot	REAL	Sum of all RGA partial pressures (Torr).
	monitor_on	BOOL	True if β monitor table is being written to.
	nmr_on	BOOL	True if nmr table is being written to.
	isotope	TEXT	Radioactive isotope being studied (Ne19, He6, or Kr83m).
	rf_side	INT	0 = U-side, 1 = I-side.

Table C.1: Description of each column in the run_log table of the he6cres_db.

<code>spec_files</code>	Table column	Type	Description
	<code>spec_id</code>	INT	Unique identifier of spec file.
	<code>created_at</code>	TIME	Date and time of row creation.
	<code>run_id</code>	INT	<code>run_id</code> associated with this file.
	<code>file_path</code>	TEXT	File path of file on DAQ computer.
	<code>file_size_mb</code>	REAL	File size (MB).
	<code>packets</code>	INT	Number of packets in file.
	<code>num_gaps</code>	INT	Number of chunks of missing packets.
	<code>mean_gap</code>	REAL	Mean gap size.
	<code>max_gap</code>	INT	Max gap size.
	<code>num_dropped_packets</code>	INT	Total number of dropped packets.
	<code>frac_of_packets_dropped</code>	REAL	Dropped packets over total packets in file.

Table C.2: Description of each column in the `spec_files` table of the `he6cres_db`.

<code>nmr</code>	Table column	Type	Description
	<code>nmr_id</code>	INT	Unique identifier of nmr value.
	<code>created_at</code>	TIME	Date and time of row creation.
	<code>field</code>	REAL	Field value (always positive).
	<code>locked</code>	BOOL	True if nmr probe is locked.

Table C.3: Description of each column in the `nmr` table of the `he6cres_db`.

<code>monitor</code>	Table column	Type	Description
	<code>monitor_id</code>	INT	Unique identifier of monitor value.
	<code>created_at</code>	TIME	Date and time of row creation.
	<code>rate</code>	REAL	Number of acquisitions (spec files) in data run.

Table C.4: Description of each column in the `monitor` table of the `he6cres.db`.

Appendix D

PARTICLE-LEVEL MONTE CARLO SIMULATION DOCUMENTATION

D.1 Purpose

This Appendix serves to document the analytic formulas used to obtain CRES event properties in the python-based He6-CRES Monte Carlo simulation package (*he6-cres-spec-sims*) as well as to compare the properties of trapped CRES electrons generated by Kassiopeia and *he6-cres-spec-sims*. The python simulations have many advantages over Kassiopeia (speed, easier to interpret, easier to alter) but the python simulations analytically calculate CRES event properties, relying on the validity of certain assumptions like the adiabatic approximation ($\omega_{axial} \ll \omega_{cyclotron}$). We would like to verify that there is agreement between the python simulations and Kassiopeia to the level necessary for our first determinations of β -energy spectra.

D.2 Overview

Section D.3 presents a comparison between the three coil trap field obtained in Kassiopeia and the python simulations. Section D.4 presents analytical formulas for all important CRES electron properties (such as cyclotron frequency, axial frequency, etc.) that are used in the python simulations. Section D.5 compares the results from the analytical formula's derived in Section D.4 (using the trap field profile presented in Section D.3) to CRES properties extracted from Kassiopeia particle tracks. Section D.6 summarizes the situation with respect to the viability of using these analytical expressions to build a sufficiently detailed particle-level Monte Carlo for the He6-CRES experiment.

D.3 Trap Field

D.3.1 Comparison between Kassiopeia and Python

Here we present a comparison between Kassiopeia and python simulations of the trap field magnitudes within the CRES decay cell. The python simulations use Biot-Savart and assume each of the three trap coils is made of 20 concentric sheets of current filling the volume taken up by the physical coil. A fine grid of points are calculated (a rectangular grid with a 0.1mm side-length, taking advantage of rotational symmetry) and then an interpolation is conducted using the *Interp2d* method from the python library *SciPy*. Kassiopeia uses a harmonic expansion method which takes advantage of axial symmetry to calculate the fields. Kassiopeia doesn't use interpolation but instead uses the field expansion to calculate the field at each specific point.

Below are absolute and relative comparisons between the field values calculated by Kassiopeia and Python. The comparison was done for a trap current of 0.87 A with a 1 T main field.

We find that the relative difference is within 10^{-6} for all locations within the decay cell. Given that we are using relative trap depths on the order of 10^{-3} at the smallest, this is sufficient agreement. Keep in mind that the trap field scales linearly with current and the background field is taken to be constant in both simulations, so it is the relative difference between the simulations that is meaningful. Figures D.1, D.2, and D.3 show the comparison between the Python-based field simulations and Kassiopeia simulations for a 1.0 T background field and a 0.87 A current in the trapping coil ($4 * 10^{-4}$ relative trap depth).

D.3.2 Trap field ρ dependence

In Figure D.2, one notices that, for a given z position, the field's dependence on ρ can be roughly modeled as a constant plus a quadratic term. This makes intuitive sense given the rotational symmetry of our geometry.

One interesting observation is that at $z = 7$ mm there is an inflection point in the ρ




Figure D.1: Field values vs z position (at $\rho = 0$ mm) for Python-based field simulations and Kassiopeia simulations for a 1.0 T background field and a 0.87 A current in the trapping coil ($4 * 10^{-4}$ relative trap depth).

dependence of the field (see Figure D.2, and D.4). It is this inflection that allows for grad B motion of both clockwise and counter-clockwise rotation about the center axis. It is the time averaged ρ gradient over one axial oscillation that determines the direction of the grad B drift.

However, for all of the electrons we observe, the maximum z position rarely exceeds 5 mm due to the Doppler effect. The modulation index is directly proportional to the maximum z position; thus, at a certain z position the CRES power is distributed across so many sidebands that we are no longer able to observe it above our noise floor. This means that in practice we only detect electrons with one handedness of grad B motion.

In Section D.4.2 we will utilize the fact that the ρ dependence of the field can be modeled by Equation D.1.

$$B(r, z) = B(z)(1 + A(z)r^2) \quad (\text{D.1})$$

When we use Equation D.1 and fit the ρ dependence of the field (at a give z position)




Figure D.2: Field values vs ρ position (at $z = 0, 2, 4$ mm) for Python-based field simulations and Kassiopeia simulations for a 1.0 T background field and a 0.87 A current in the trapping coil ($4 * 10^{-4}$ relative trap depth).

we end up with an $R^2 \approx .9994$ and residuals all within $10^{-5} T$. We find $A \approx -10 T m^2$ for relevant z positions (less than 5 mm).

D.4 Analytic calculations of CRES dynamics

D.4.1 Dynamics assuming no ρ dependence in B

Let us first derive the dynamics of an electron in a symmetric, single well trap where one assumes the magnetic field is purely in the z direction and that the magnitude of the field is only dependent on z ($\vec{B} = B(z) \hat{z}$). We begin with the basics of cyclotron motion [30, 5].

$$f_c = \frac{1}{2\pi} \frac{e B_{avg}}{\gamma m}, \quad \gamma = 1 + \frac{K}{mc^2} = \frac{E}{mc^2} \quad (D.2)$$

Where B_{avg} is the average instantaneous field experienced by the CRES electron. We can now take advantage of the fact that p_{\perp}^2/B is an adiabatic constant of the motion for $f_{axial} \ll f_{cyclotron}$ (Eq 12.69 [62]).




Figure D.3: Relative difference in field values vs ρ position (at $z = 0, 2, 4$ mm) for Python-based field simulations and Kassiopeia simulations for a 1.0 T background field and a 0.87 A current in the trapping coil ($4 * 10^{-4}$ relative trap depth).

$$\frac{p_{\perp}^2}{B} = \frac{\gamma m v_{\perp}^2}{B} = \frac{\gamma m v \sin^2 \theta}{B} = \text{constant} \quad \longrightarrow \quad \boxed{\frac{\sin^2 \theta}{B} = \text{constant}} \quad (\text{D.3})$$

Now we would like to use the above condition to derive the axial frequency of the electron (f_{axial}) and average magnetic field experienced by the electron (B_{avg}). We begin by deriving the period of the axial motion. Here we take advantage of the symmetric nature of the trap, implying that the period is 4 times the time it takes the electron to travel from $z = 0$ to $z = z_{max}$.

$$T = 4 \int_0^{z_{max}} \frac{dt}{dz} dz = \int_0^{z_{max}} \frac{1}{v_z} dz = \frac{4}{v} \int_0^{z_{max}} \frac{1}{\cos(\theta)} dz \quad (\text{D.4})$$

Now we notice that the above constant of the motion (Equation D.16), can be used to make progress on the above integral.

$$\frac{\sin^2 \theta}{B} = \text{Const} \quad \longrightarrow \quad \frac{\sin^2 \theta(z)}{B(z)} = \frac{1}{B(z_{max})} \quad (\text{D.5})$$

$$\cos(\theta) = \left[1 - \frac{B(z)}{B(z_{max})} \right]^{.5}$$

$$T = \frac{4}{v} \int_0^{z_{max}} \left[1 - \frac{B(z)}{B(z_{max})} \right]^{-.5} dz \quad (\text{D.6})$$

One can then use the relation $f_{axial} = \frac{1}{T}$. We can evaluate B_{avg} in a similar manner:

$$B_{avg} = \frac{4}{T} \int_0^{T/4} B(t) dt = \frac{4}{T} \int_0^{z_{max}} B(z) \frac{1}{v_z} dz \quad (\text{D.7})$$

$$B_{avg} = \frac{4}{Tv} \int_0^{z_{max}} B(z) \left[1 - \frac{B(z)}{B(z_{max})} \right]^{-.5} dz \quad (\text{D.8})$$

D.4.2 Dynamics assuming quadratic ρ dependence on B

As was illustrated above (Section D.3.2) our field's dependence on r is well approximated by

$$B(r, z) = B(z)(1 + A(z)r^2). \quad (\text{D.9})$$

For sufficiently small cyclotron radii the effect of this radial dependence is negligible. However, for our field and energy ranges of interest we observe cyclotron radii around 2 – 3 mm. This means that for many CRES electrons with large pitch angles (close to 90°) z_{max} is comparable to r_c . For our trap geometry, this means that the field variation experienced by the electron in one cyclotron orbit (within the xy – plane) can be on the same order of magnitude as the field variation experienced in z over one axial period. Thus for our trap configuration it is not fair to assume, as was done above in Section D.4.1, that the field experienced by the electron can be well approximated by the field at the guiding center.

The fundamental relationship between frequency and the average field (B_{avg}) seen in Equation D.2 still holds but now we need to be more careful about how to interpret the B in the denominator of our relevant adiabatic invariant (p_{\perp}^2/B) when we calculate the average field experienced by the electron (B_{avg}). This B needs to be the average field over the cross-section of the cyclotron orbit, not the B at the center of the cyclotron orbit. This is because the adiabatic invariant we invoke is derived from the conserved flux through a cyclotron orbit (see Equation 12.67 in [62]; note that the relevant steps from one adiabatic invariant to the other can be found in Equation 12.71). So the relevant B can be derived from:

$$\Phi_B = \int_A B(r, z) dA. \quad (D.10)$$

Here r denotes the distance from the center of the waveguide to a point of interest, let ρ denote the distance from the center of the waveguide to the guiding center of the cyclotron orbit, and let r' denote the distance from the guiding center of the cyclotron orbit to a point of interest. Let ϕ be the angle between the x-axis and r' . See Fig. D.5 for a visualization of the coordinate system. We then have the following relationship:

$$r^2 = (\rho + r' \cos \phi)^2 + (r' \sin \phi)^2 \quad (D.11)$$

Using the above we can make progress on our integral (Equation D.10).

$$\Phi_B = \int_0^{2\pi} \int_0^{r_c} B(z) (1 + A((\rho + r' \cos \phi)^2 + (r' \sin \phi)^2)) r' d\theta dr' \quad (D.12)$$

$$\begin{aligned} \Phi_B &= \pi r_c^2 B(z) + B(z) A \int_0^{2\pi} \int_0^{r_c} ((\rho + r' \cos \phi)^2 + (r' \sin \phi)^2) r' d\theta dr' \\ &= \pi r_c^2 B(z) + B(z) A \int_0^{2\pi} \int_0^{r_c} (\rho^2 r' + 2\rho r'^2 \cos \phi + r'^3 \cos^2(\phi) + r'^3 \sin^2(\phi)) d\theta dr' \quad (D.13) \\ &= \pi r_c^2 B(z) + B(z) A \int_0^{2\pi} \int_0^{r_c} (\rho^2 r' + r'^3) d\theta dr' \end{aligned}$$

We obtain the third equality of Equation D.13 by noting that $\cos^2(\phi) + \sin^2(\phi) = 1$ and that $\cos(\phi)$ evaluates to 0 when integrated over a period.

$$\Phi_B = \pi r_c^2 B(z) \left(1 + A\left(\rho^2 + \frac{r_c^2}{2}\right)\right) \quad (\text{D.14})$$

$$\Phi_B = \pi r_c^2 B(\rho', z) \quad w/ \quad \rho' = \left(\rho^2 + \frac{r_c^2}{2}\right)^{.5} \quad (\text{D.15})$$

What we find is that the B relevant to the constant of the motion (p_{\perp}^2/B) corresponds to a field value slightly displaced from the guiding center. Intuitively, this means that for positive A in Equation D.9 the value of $\Phi_B/\pi r_c^2$ is slightly larger than the field at the guiding center. Therefore, we find that the constant of the motion is given by:

$$\frac{\sin^2 \theta}{B(\rho', z)} = \text{constant} \quad (\text{D.16})$$

We can use the above to find the axial frequency $f_{axial} = 1/T$ as we did in Equation D.6.

$$T = \frac{4}{v} \int_0^{z_{max}} \left[1 - \frac{B(\rho', z)}{B(\rho', z_{max})}\right]^{-.5} dz \quad (\text{D.17})$$

Now we would like to find the average magnetic field experienced by the electron. The instantaneous magnetic field is also effected by the radial dependence of the field. We find that the average field experienced by a CRES electron in a single cyclotron orbit whose guiding center is at ρ , which we denote $B_{one\ orbit}(\rho, z)$, is given by:

$$\begin{aligned} B_{one\ orbit}(\rho, z) &= \frac{1}{2\pi} \int_0^{2\pi} B(r, z) d\theta \\ &= \frac{B(z)}{2\pi} \int_0^{2\pi} (1 + A((\rho + r_c \cos\phi)^2 + (r_c \sin\phi)^2)) d\theta \end{aligned} \quad (\text{D.18})$$

$$B_{one\ orbit}(\rho, z) = B(\rho'', z) \quad w/ \quad \rho'' = (\rho^2 + r_c^2)^{.5} \quad (\text{D.19})$$

Putting it all together we see that B_{avg} (from which we can find f_{cycl}) is given by:

$$B_{avg} = \frac{4}{Tv} \int_0^{z_{max}} B(\rho'', z) \left[1 - \frac{B(\rho', z)}{B(\rho', z_{max})} \right]^{-.5} dz \quad (\text{D.20})$$

As a final note, for our desired levels of uncertainty, it is fine to assume:

$$r_c = \frac{p \sin \theta}{eB(\rho, z)} = \frac{\gamma m v \sin \theta}{eB(\rho, z)} \quad (\text{D.21})$$

It is also worth noting that the definition of ρ' and ρ'' include a r_c , and r_c changes as a function of z . However, the difference between $r_{c,max}$ and $r_{c,min}$ is less than a part in one thousand and so we just use $r_{c,max}$ in the definitions of ρ' and ρ'' with no observable loss of precision. We have now constructed a set of expressions that can be used to precisely define the dynamics of a CRES electron without explicitly simulating particle trajectories. For completion, we list all of these expressions explicitly:

$$r_c = \frac{p \sin \theta}{eB(\rho, z)} \quad (\text{Cyclotron Radius})$$

$$\frac{\sin^2 \theta}{B(\rho', z)} = \frac{1}{B(\rho', z_{max})} \quad (\text{Maximum } z \text{ position})$$

$$f_{axial} = \left(\frac{4}{v} \int_0^{z_{max}} \left[1 - \frac{B(\rho', z)}{B(\rho', z_{max})} \right]^{-.5} dz \right)^{-1} \quad (\text{Axial Frequency})$$

$$B_{avg} = \frac{4f_{axial}}{v} \int_0^{z_{max}} B(\rho'', z) \left[1 - \frac{B(\rho', z)}{B(\rho', z_{max})} \right]^{-.5} dz \quad (\text{Average Magnetic Field})$$

$$f_c = \frac{1}{2\pi} \frac{e B_{avg}}{\gamma m}, \quad \gamma = 1 + \frac{K}{mc^2} = \frac{E}{mc^2} \quad (\text{Average Cyclotron Frequency})$$

D.5 Comparison to *Kassiopeia*

Kassiopeia relies on tracking particle trajectories and is therefore very slow within the context of needing to ultimately simulate more than 10^{10} CRES electrons. We would prefer to use the analytical expressions derived in Sections D.4.1 and D.4.2, which are considerably faster and easier to interact with, to build our Monte Carlo simulation. We conclude that the analytical expressions are sufficiently accurate for our current desired level of precision. Below we present a comparison between the high level CRES electron properties (axial frequency, maximum z position, time averaged magnetic field experienced, cyclotron frequency, and cyclotron radius) found with *Kassiopeia* and with the analytical expressions presented in Sections D.4.1 and D.4.2. We conducted the comparison at 3 main fields (1 T, 3 T, 5 T), two cyclotron frequencies (18 GHz, 19 GHz), and two pitch angles (88° , 89.5°), all at a fixed relative trap depth of $4 * 10^{-3}$. Accounting for the ρ dependence in the magnetic field (using the expression from Section D.4.2) resulted in improved agreement for all CRES electron properties. The level of improvement observed depends somewhat on the CRES property, but were on the order of a factor of 50 reduction in relative difference in comparison to *Kassiopeia* (see Figures D.6, D.7, and D.8 below).

We find that the cyclotron frequencies, the most important property for the extraction of the electron energy, agree to within 2.5 ppm between the improved analytical expressions (Section D.4.2) and *Kassiopeia* for all simulation parameters tested. A 2.5 ppm error in frequency at 18 GHz corresponds to a 10 eV error in energy at 5 T and even less at smaller fields. All other CRES electron properties agree to within 4 parts per thousand. This is more than sufficient agreement for our current desired level of precision.

Figures D.6, D.7, and D.8 show the relative difference between the CRES properties found with *Kassiopeia* and the analytical calculations (Equation D.22), for three specific sets of simulation parameters, presented in parts per million. Note that in the comparison tables below, “original python simulations” refers to the use of expressions from Section D.4.1, and “improved python simulations” refers to the use of expressions from Section D.4.2.

$$\textit{Relative Difference} = \frac{\textit{property}_{Kassiopeia} - \textit{property}_{python}}{\textit{property}_{python}} \quad (\text{D.22})$$

D.6 Conclusion

We conclude that we can confidently use these analytical expressions for CRES properties derived in Section D.4.2 to build a fast and easy to alter Monte Carlo for simulating the He6-CRES experiment. Correctly accounting for the power lost to radiation, frequency jumps due to scattering with residual gas, and other such effects unique to CRES would be very difficult to model in Kassiopeia, further motivating the need for a more adaptable simulation.




Figure D.4: Field values vs ρ position (at $z = 1, 4, 7, 10, 13$ mm) for Python-based field simulations for a 1.0 T background field and a 0.87 A current in the trapping coil ($4 * 10^{-4}$ relative trap depth). One can see that the ρ dependence is negligible at $z = 7$ mm.




Figure D.5: Coordinate system used for integrals (Equation D.10, Equation D.18). The red circle indicates the cyclotron orbit of a CRES electron. The magnetic field, $B(r, z)$, is defined relative to the central axis of the trapping field.

Simulation Params:

	Main Field (T)	Trap Current (A)	Frequency (GHz)	Pitch Angle (deg)
0	1	0.5	18.0	88.0

Difference between Kasseoppia and original python simulations (ppm):

	axial_freq	z_max	cycl_freq	B_average	cycl_radius_min
rho					
0 mm	16143.65	15682.97	16.17	16.21	19.80
1 mm	15553.10	15504.65	15.92	15.96	19.72
2 mm	15529.29	15994.75	17.83	17.87	20.43
3 mm	15975.59	16909.21	17.98	18.04	21.75

Difference between Kasseoppia and improved python simulations (ppm):

	axial_freq	z_max	cycl_freq	B_average	cycl_radius_min
rho					
0 mm	347.20	326.11	0.46	0.42	0.39
1 mm	780.63	285.94	0.36	0.31	0.22
2 mm	1183.72	37.11	0.71	0.75	0.13
3 mm	1367.03	485.58	0.10	0.04	0.83

Figure D.6: Comparison between Kasseoppia and analytical expressions (original = Section D.4.1 expressions, improved = Section D.4.2 expressions) for all important CRES electron properties in ppm. **Simulation Parameters:** Main Field: 1 T, Trap Current: .5 A, Cyclotron Frequency: 18 GHz, Pitch Angle: 88°.

Simulation Params:

	Main Field (T)	Trap Current (A)	Frequency (GHz)	Pitch Angle (deg)
0	3	1.5	18.0	89.5

Difference between Kasseoppia and original python simulations (ppm):

	axial_freq	z_max	cycl_freq	B_average	cycl_radius_min
rho					
0 mm	21903.91	22132.18	47.87	47.89	3.09
1 mm	21185.04	21919.23	48.22	48.23	2.91
2 mm	24020.11	22002.68	50.24	50.26	2.91
3 mm	24533.93	21923.27	53.22	53.24	2.88

Difference between Kasseoppia and improved python simulations (ppm):

	axial_freq	z_max	cycl_freq	B_average	cycl_radius_min
rho					
0 mm	146.84	215.03	0.32	0.30	0.13
1 mm	614.00	150.63	0.08	0.07	0.27
2 mm	2543.72	213.46	1.03	1.04	0.23
3 mm	3559.46	666.35	2.37	2.39	0.19

Figure D.7: Comparison between Kasseoppia and analytical expressions (original = Section D.4.1 expressions, improved = Section D.4.2 expressions) for all important CRES electron properties in ppm. **Simulation Parameters:** Main Field: 3 T, Trap Current: 1.5 A, Cyclotron Frequency: 18 GHz, Pitch Angle: 89.5°.

Simulation Params:

	Main Field (T)	Trap Current (A)	Frequency (GHz)	Pitch Angle (deg)
0	5	2.5	19.0	88.0

Difference between Kasseoppia and original python simulations (ppm):

	axial_freq	z_max	cycl_freq	B_average	cycl_radius_min
rho					
0 mm	24324.58	23807.81	24.81	24.83	29.94
1 mm	25048.47	23678.76	24.88	24.90	29.96
2 mm	26673.19	24482.52	25.82	25.84	31.09
3 mm	26655.22	25920.55	27.63	27.67	33.14

Difference between Kasseoppia and improved python simulations (ppm):

	axial_freq	z_max	cycl_freq	B_average	cycl_radius_min
rho					
0 mm	333.64	287.15	0.07	0.05	0.33
1 mm	534.56	208.38	0.19	0.17	0.08
2 mm	1616.02	216.76	0.07	0.05	0.49
3 mm	628.51	1051.10	0.23	0.27	1.60

Figure D.8: Kassiopiea vs analytical expressions (original = Section D.4.1 expressions, improved = Section D.4.2 expressions) for all important CRES electron properties in ppm.

Simulation Parameters: Main Field: 5 T, Trap Current: 2.5 A, Cyclotron Frequency: 19 GHz, Pitch Angle: 88°.

Appendix E

DEEP LEARNING SOURCE CODE

Below are select portions of the source code used to conduct the studies presented in Chapter 7. Section E.1 contains the UNET Model implementation, written using the PyTorch library [84]. Section E.2 contains the UNET LightningModule written using the PyTorch Lightning library. Pytorch Lightning allows for the easy organization of a deep learning experiment and decreases boiler plate code. To see a specific demonstration of how the LightningModule is trained look in the `demo` directory in the associated github repo: <https://github.com/Helium6CRES/he6-cres-deep-learning>.

E.1 UNET model implementation

```
# Author: Drew Byron
# Date: 1/17/23

# Deep learning imports.
import torch
from torch import nn
import torch.nn.functional as F
from torch.utils.data import DataLoader, random_split, TensorDataset

import torchvision
from torchvision.utils import draw_bounding_boxes, draw_segmentation_masks, make_grid
from torchvision.ops import masks_to_boxes
import torchvision.transforms.functional as TF
from torchvision.models.detection.faster_rcnn import FastRCNNPredictor
from torchvision.models.detection.mask_rcnn import MaskRCNNPredictor

import pytorch_lightning as pl
from pytorch_lightning.callbacks import ModelCheckpoint
from pytorch_lightning.loggers import TensorBoardLogger

import torchmetrics

# Standard imports.
from typing import List, Union
import gc
```

```

import matplotlib.pyplot as plt
import numpy as np
import zipfile
from pathlib import Path

class UNET(nn.Module):
    """A PyTorch implementation of a UNET image segmentation model based
    on this work: https://arxiv.org/abs/1505.04597. Specifics based loosely
    on Aladdin Persson's Implimentation:
    https://github.com/aladdinpersson/
    Machine-Learning-Collection/blob/master/ML/Pytorch/image\_segmentation/
    semantic\_segmentation\_unet/model.py

    Args:
        in_channels (int): number of input channels. For example for an
            put of shape (batch_size, 3, img_size, img_size) in_channels
            is 3.
        num_classes (int): number of output classes desired. The output
            shape of the model will be (batch_size, num_classes, img_size,
            img_size). For example output[0][i] is a binary segmentation
            mask for class i. Note that class 0 is reserved for background.
        first_feature_num (int): An int specifying the number of features to
            be used in the first DoubleConv layer.
        num_layers (int): Number of layers to use in the UNET architecture.
            The ith layer contains first_feature_num * 2**i features. Note
            that if img_size // 2**num_layers < 1 then the model will break.
        kernel_size (int): A kernel of shape (kernel_size, kernel_size)
            will be applied to the imgs during both Conv2d layers of
            DoubleConv.
        bias (bool): whether or not to add a bias to the DoubleConv Conv2d
            layers.
        track_x_shape (bool): whether or not to track the shape of x.
    """

    def __init__(
        self,
        in_channels=3,
        num_classes=4,
        first_feature_num=8,
        num_layers=3,
        skip_connect=True,
        kernel_size=3,
        bias=True,
        track_x_shape=False,
    ):
        super().__init__()

        self.in_channels = in_channels
        self.num_classes = num_classes
        self.features = [first_feature_num * 2**i for i in range(num_layers)]
        self.skip_connect = skip_connect
        self.kernel_size = kernel_size
        self.bias = bias
        self.track_x_shape = track_x_shape

        self.ups = nn.ModuleList()
        self.downs = nn.ModuleList()

```

```

self.pool = nn.MaxPool2d(kernel_size=2, stride=2)

self.bottleneck = DoubleConv(self.features[-1], self.features[-1] * 2)
self.final_conv = nn.Conv2d(self.features[0], self.num_classes, kernel_size=1)

if self.track_x_shape:
    self.x_shape_tracker = []

# Down part of UNET.
for feature in self.features:
    self.downs.append(
        DoubleConv(
            in_channels, feature, kernel_size=self.kernel_size, bias=self.bias
        )
    )
    in_channels = feature

# Up part of UNET.
for feature in reversed(self.features):
    self.ups.append(
        nn.ConvTranspose2d(
            feature * 2,
            feature,
            kernel_size=2,
            stride=2,
        )
    )
    if self.skip_connect:
        self.ups.append(
            DoubleConv(
                feature * 2,
                feature,
                kernel_size=self.kernel_size,
                bias=self.bias,
            )
        )
    else:
        self.ups.append(
            DoubleConv(
                feature, feature, kernel_size=self.kernel_size, bias=self.bias
            )
        )

def track_shape(self, x, description):
    if self.track_x_shape:
        self.x_shape_tracker.append((f"{description}:\n\t{x.shape}"))

    return None

def forward(self, x):
    skip_connections = []

    self.track_shape(x, "input shape")

    for idx, down in enumerate(self.downs):
        x = down(x)

```

```

        self.track_shape(x, f"double_conv (down) {idx}")

        skip_connections.append(x)
        x = self.pool(x)

        self.track_shape(x, f"max_pool {idx}")

    x = self.bottleneck(x)
    self.track_shape(x, "bottleneck")

    # Reverse the list of skip connections.
    skip_connections = skip_connections[::-1]

    for idx in range(0, len(self.ups), 2):

        x = self.ups[idx](x)

        self.track_shape(x, f"conv_trans {idx//2}")

        skip_connection = skip_connections[idx // 2]

        if x.shape != skip_connection.shape:
            x = TF.resize(x, size=skip_connection.shape[2:])

        if self.skip_connect:
            x = torch.cat((skip_connection, x), dim=1)
            self.track_shape(x, f"skip connection {idx//2}")

        x = self.ups[idx + 1](x)

        self.track_shape(x, f"double_conv (up) {idx//2}")

    x = self.final_conv(x)

    self.track_shape(x, "output shape")

    return x

```

E.2 UNET lightning module

```

# Author: Drew Byron
# Date: 1/17/23

# Deep learning imports.
import torch
from torch import nn
import torch.nn.functional as F
from torch.utils.data import DataLoader, random_split, TensorDataset

import torchvision
from torchvision.utils import draw_bounding_boxes, draw_segmentation_masks, make_grid
from torchvision.ops import masks_to_boxes

```

```

import torchvision.transforms.functional as TF
from torchvision.models.detection.faster_rcnn import FastRCNNPredictor
from torchvision.models.detection.mask_rcnn import MaskRCNNPredictor

import pytorch_lightning as pl
from pytorch_lightning.callbacks import ModelCheckpoint
from pytorch_lightning.loggers import TensorBoardLogger

import torchmetrics

# Standard imports.
from typing import List, Union
import gc
import matplotlib.pyplot as plt
import numpy as np
import zipfile
from pathlib import Path

class LightningImageSegmentation(pl.LightningModule):
    def __init__(
        self,
        in_channels=3,
        num_classes=4,
        first_feature_num=4,
        num_layers=2,
        skip_connect=True,
        kernel_size=3,
        bias=False,
        weight_loss=None,
        learning_rate=1e-3,
    ):
        super().__init__()

        # LM attributes.
        self.in_channels = in_channels
        self.num_classes = num_classes
        self.first_feature_num = first_feature_num
        self.num_layers = num_layers
        self.skip_connect = skip_connect
        self.kernel_size = kernel_size
        self.bias = bias
        self.learning_rate = learning_rate

        if weight_loss == None:
            self.weight_loss = torch.ones(self.num_classes)
        else:
            self.weight_loss = weight_loss

        # Log hyperparameters.
        self.save_hyperparameters()

        # Metrics.
        self.train_acc = torchmetrics.Accuracy(
            full_state_update="false", mdmc_average="samplewise"
        )
        self.train_f1 = torchmetrics.F1Score(

```

```

        full_state_update="false",
        mdmc_average="global",
        average="macro",
        num_classes=self.num_classes,
    )
    self.train_iou = torchmetrics.JaccardIndex(
        full_state_update="false", num_classes=self.num_classes, average="macro"
    )
    self.val_acc = torchmetrics.Accuracy(
        full_state_update="false", mdmc_average="samplewise"
    )
    self.val_f1 = torchmetrics.F1Score(
        full_state_update="false",
        mdmc_average="global",
        average="macro",
        num_classes=self.num_classes,
    )
    self.val_iou = torchmetrics.JaccardIndex(
        full_state_update="false", num_classes=self.num_classes, average="macro"
    )

    # Loss function.
    self.loss = nn.CrossEntropyLoss(weight=self.weight_loss)

    # Define Model.
    self.model = UNET(
        in_channels=self.in_channels,
        num_classes=self.num_classes,
        first_feature_num=self.first_feature_num,
        skip_connect=self.skip_connect,
        kernel_size=self.kernel_size,
        bias=self.bias,
    )

    # Sample input. Used for logging the model graph.
    self.example_input_array = torch.rand(4, 1, 256, 256)

def forward(self, imgs):

    imgs_normed = imgs.float() / 255.0
    return self.model(imgs_normed)

def cross_entropy_loss(self, logits, labels):
    return self.loss(logits, labels)

def training_step(self, train_batch, batch_idx):

    # Grab images and labels from batch.
    images, labels = train_batch
    logits = self.forward(images)

    # Calculate loss.
    loss = self.cross_entropy_loss(logits, labels)

    # Log step metrics.
    self.train_acc(logits, labels)
    self.train_f1(logits, labels)

```

```
self.train_iou(logits, labels)

self.log("train/loss", loss)
self.log("train/acc", self.train_acc, on_step=True)
self.log("train/f1", self.train_f1, on_step=True)
self.log("train/iou", self.train_iou, on_step=True)

return loss

def validation_step(self, val_batch, batch_idx):

    # Grab images and labels from batch.
    images, labels = val_batch
    logits = self.forward(images)

    # Calculate loss.
    loss = self.cross_entropy_loss(logits, labels)

    # log step metrics.
    self.val_acc(logits, labels)
    self.val_f1(logits, labels)
    self.val_iou(logits, labels)

    self.log("val/loss", loss)
    self.log("val/acc", self.val_acc, on_step=True)
    self.log("val/f1", self.val_f1, on_step=True)
    self.log("val/iou", self.val_iou, on_step=True)

def configure_optimizers(self):
    optimizer = torch.optim.Adam(self.parameters(), lr=self.learning_rate)
    return optimizer
```

BIBLIOGRAPHY

- [1] B. Abi, T. Albahri, S. Al-Kilani, D. Allspach, L. P. Alonzi, A. Anastasi, A. Anisenkov, F. Azfar, K. Badgley, S. Baeßler, I. Bailey, V. A. Baranov, E. Barlas-Yucel, T. Barrett, E. Barzi, A. Basti, F. Bedeschi, A. Behnke, M. Berz, M. Bhattacharya, H. P. Binney, R. Bjorkquist, P. Bloom, J. Bono, E. Bottalico, T. Bowcock, D. Boyden, G. Cantatore, R. M. Carey, J. Carroll, B. C. K. Casey, D. Cauz, S. Ceravolo, R. Chakraborty, S. P. Chang, A. Chapelain, S. Chappa, S. Charity, R. Chislett, J. Choi, Z. Chu, T. E. Chupp, M. E. Convery, A. Conway, G. Corradi, S. Corrodi, L. Cotrozzi, J. D. Crnkovic, S. Dabagov, P. M. De Lurgio, P. T. Debevec, S. Di Falco, P. Di Meo, G. Di Sciascio, R. Di Stefano, B. Drendel, A. Driutti, V. N. Duginov, M. Eads, N. Eggert, A. Epps, J. Esquivel, M. Farooq, R. Fatemi, C. Ferrari, M. Fertl, A. Fiedler, A. T. Fienberg, A. Fioretti, D. Flay, S. B. Foster, H. Friedsam, E. Frlež, N. S. Froemming, J. Fry, C. Fu, C. Gabbanini, M. D. Galati, S. Ganguly, A. Garcia, D. E. Gastler, J. George, L. K. Gibbons, A. Gioiosa, K. L. Giovanetti, P. Girotti, W. Gohn, T. Gorringer, J. Grange, S. Grant, F. Gray, S. Haciomeroglu, D. Hahn, T. Halewood-Leagas, D. Hampai, F. Han, E. Hazen, J. Hempstead, S. Henry, A. T. Herrod, D. W. Hertzog, G. Hesketh, A. Hibbert, Z. Hodge, J. L. Holzbauer, K. W. Hong, R. Hong, M. Iacovacci, M. Incagli, C. Johnstone, J. A. Johnstone, P. Kammel, M. Kargiantoulakis, M. Karuza, J. Kaspar, D. Kawall, L. Kelton, A. Keshavarzi, D. Kessler, K. S. Khaw, Z. Khechadorian, N. V. Khomutov, B. Kiburg, M. Kiburg, O. Kim, S. C. Kim, Y. I. Kim, B. King, N. Kinnaird, M. Korostelev, I. Kourbanis, E. Kraegeloh, V. A. Krylov, A. Kuchibhotla, N. A. Kuchinskiy, K. R. Labe, J. LaBounty, M. Lancaster, M. J. Lee, S. Lee, S. Leo, B. Li, D. Li, L. Li, I. Logashenko, A. Lorente Campos, A. Lucà, G. Lukicov, G. Luo, A. Lusiani, A. L. Lyon, B. Mac-

- Coy, R. Madrak, K. Makino, F. Marignetti, S. Mastroianni, S. Maxfield, M. McEvoy, W. Merritt, A. A. Mikhailichenko, J. P. Miller, S. Miozzi, J. P. Morgan, W. M. Morse, J. Mott, E. Motuk, A. Nath, D. Newton, H. Nguyen, M. Oberling, R. Os-
ofsky, J.-F. Ostiguy, S. Park, G. Pauletta, G. M. Piacentino, R. N. Pilato, K. T. Pitts,
B. Plaster, D. Počanić, N. Pohlman, C. C. Polly, M. Popovic, J. Price, B. Quinn,
N. Raha, S. Ramachandran, E. Ramberg, N. T. Rider, J. L. Ritchie, B. L. Roberts,
D. L. Rubin, L. Santi, D. Sathyan, H. Schellman, C. Schlesier, A. Schreckenberger,
Y. K. Semertzidis, Y. M. Shatunov, D. Shemyakin, M. Shenk, D. Sim, M. W. Smith,
A. Smith, A. K. Soha, M. Sorbara, D. Stöckinger, J. Stapleton, D. Still, C. Stoughton,
D. Stratakis, C. Strohmaier, T. Stuttard, H. E. Swanson, G. Sweetmore, D. A. Sweigart,
M. J. Syphers, D. A. Tarazona, T. Teubner, A. E. Tewsley-Booth, K. Thomson,
V. Tishchenko, N. H. Tran, W. Turner, E. Valetov, D. Vasilkova, G. Venanzoni, V. P.
Volnykh, T. Walton, M. Warren, A. Weisskopf, L. Welty-Rieger, M. Whitley, P. Win-
ter, A. Wolski, M. Wormald, W. Wu, and C. Yoshikawa. Measurement of the posi-
tive muon anomalous magnetic moment to 0.46 ppm. *Phys. Rev. Lett.*, 126:141801,
Apr 2021. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.126.141801>,
doi:10.1103/PhysRevLett.126.141801.
- [2] P. Alfaut et al. WISArD : Weak Interaction Studies with ^{32}Ar Decay. *PoS*,
PANIC2021:449, 2022. doi:10.22323/1.380.0449.
- [3] Michael Artin. *Algebra*. Birkhäuser, 1998.
- [4] M. Askins, Z. Bagdasarian, N. Barros, E. W. Beier, A. Bernstein, E. Blucher, R. Bon-
ventre, E. Bourret, E. J. Callaghan, J. Caravaca, M. Diwan, S. T. Dye, J. Eisch,
A. Elagin, T. Enqvist, U. Fahrenholz, V. Fischer, K. Frankiewicz, C. Grant, D. Guf-
fanti, C. Hagner, A. Hallin, C. M. Jackson, R. Jiang, T. Kaptanoglu, J. R. Klein, Yu. G.
Kolomensky, C. Kraus, F. Krennrich, T. Kutter, T. Lachenmaier, B. Land, K. Lande,
L. Lebanowski, J. G. Learned, V. A. Li, V. Lozza, L. Ludhova, M. Malek, S. Manecki,
J. Maneira, J. Maricic, J. Martyn, A. Mastbaum, C. Mauger, M. Mayer, J. Migenda,

- F. Moretti, J. Napolitano, B. Naranjo, M. Nieslony, L. Oberauer, G. D. Orebi Gann, J. Ouellet, T. Pershing, S. T. Petcov, L. Pickard, R. Rosero, M. C. Sanchez, J. Sawatzki, S. H. Seo, M. Smiley, M. Smy, A. Stahl, H. Steiger, M. R. Stock, H. Sunej, R. Svoboda, E. Tiras, W. H. Trzaska, M. Tzanov, M. Vagins, C. Vilela, Z. Wang, J. Wang, M. Wetstein, M. J. Wilking, L. Winslow, P. Wittich, B. Wonsak, E. Worcester, M. Wurm, G. Yang, M. Yeh, E. D. Zimmerman, S. Zsoldos, and K. Zuber. Theia: Summary of physics program. snowmass white paper submission, 1 2022. [arXiv:2202.12839](https://arxiv.org/abs/2202.12839).
- [5] D. M. Asner, R. F. Bradley, L. de Viveiros, P. J. Doe, J. L. Fernandes, M. Fertl, E. C. Finn, J. A. Formaggio, D. Furse, A. M. Jones, J. N. Kofron, B. H. LaRoque, M. Leber, E. L. McBride, M. L. Miller, P. Mohanmurthy, B. Monreal, N. S. Oblath, R. G. H. Robertson, L. J. Rosenberg, G. Rybka, D. Rysewyk, M. G. Sternberg, J. R. Tedeschi, T. Thümmler, B. A. VanDevender, and N. L. Woods. Single-electron detection and spectroscopy via relativistic cyclotron radiation. *Phys. Rev. Lett.*, 114:162501, Apr 2015. doi:10.1103/PhysRevLett.114.162501.
- [6] John N. Bahcall, E. Lisi, D. E. Alburger, L. De Braekeleer, S. J. Freedman, and J. Napolitano. Standard neutrino spectrum from ${}^8\text{B}$ decay. *Phys. Rev. C*, 54:411–422, Jul 1996. URL: <https://link.aps.org/doi/10.1103/PhysRevC.54.411>, doi:10.1103/PhysRevC.54.411.
- [7] A. B. Balantekin, J. Carlson, D. J. Dean, G. M. Fuller, R. J. Furnstahl, M. Hjorth-Jensen, R. V. F. Janssens, Bao-An Li, W. Nazarewicz, F. M. Nunes, W. E. Ormand, S. Reddy, and B. M. Sherrill. Nuclear theory and science of the facility for rare isotope beams. *Modern Physics Letters A*, 29(11):1430010, apr 2014. URL: <https://doi.org/10.1142/Fs0217732314300109>, doi:10.1142/s0217732314300109.
- [8] R. Barbieri, C. Bouchiat, A. Georges, and P. Le Doussal. Quark-lepton non-universality in supersymmetric models. *Physics Letters B*, 156(5):348 – 352, 1985. doi:10.1016/0370-2693(85)91623-5.

- [9] Richard M. Barrer. *Diffusion in and through Solids*. Cambridge University Press, 1941. URL: <https://archive.org/details/in.ernet.dli.2015.469917>.
- [10] Evan Berkowitz, David Brantley, Chris Bouchard, Chia Cheng Chang, M. A. Clark, Nicholas Garron, Balint Joo, Thorsten Kurth, Chris Monahan, Henry Monge-Camacho, Amy Nicholson, Kostas Orginos, Enrico Rinaldi, Pavlos Vranas, and Andre Walker-Loud. An accurate calculation of the nucleon axial charge with lattice qcd, 2017. URL: <https://arxiv.org/abs/1704.01114>, doi:10.48550/ARXIV.1704.01114.
- [11] Tanmoy Bhattacharya, Vincenzo Cirigliano, Saul D. Cohen, Alberto Filipuzzi, Martín González-Alonso, Michael L. Graesser, Rajan Gupta, and Huey-Wen Lin. Probing novel scalar and tensor interactions from (ultra)cold neutrons to the LHC. *Physical Review D*, 85(5), mar 2012. URL: <https://doi.org/10.1103/PhysRevD.85.054512>, doi:10.1103/PhysRevD.85.054512.
- [12] Lowell S. Brown and Gerald Gabrielse. Geonium theory: Physics of a single electron or ion in a penning trap. *Rev. Mod. Phys.*, 58:233–311, Jan 1986. doi:10.1103/RevModPhys.58.233.
- [13] Lowell S. Brown and Gerald Gabrielse. Geonium theory: Physics of a single electron or ion in a penning trap. *Rev. Mod. Phys.*, 58:233–311, Jan 1986. URL: <https://link.aps.org/doi/10.1103/RevModPhys.58.233>, doi:10.1103/RevModPhys.58.233.
- [14] M. T. Burkey, G. Savard, A. T. Gallant, N. D. Scielzo, J. A. Clark, T. Y. Hirsh, L. Varriano, G. H. Sargsyan, K. D. Launey, M. Brodeur, D. P. Burdette, E. Heckmaier, K. Joerres, J. W. Klimes, K. Kolos, A. Laminack, K. G. Leach, A. F. Levand, B. Longfellow, B. Maaß, S. T. Marley, G. E. Morgan, P. Mueller, R. Orford, S. W. Padgett, A. Pérez Galván, J. R. Pierce, D. Ray, R. Segel, K. Siegl, K. S. Sharma, and B. S. Wang. Improved Limit on Tensor Currents in the Weak Interaction from ^8Li β Decay. *Phys. Rev. Lett.*, 128:202502, May 2022. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.128.202502>, doi:10.1103/PhysRevLett.128.202502.

- [15] W Byron, H Harrington, R J Taylor, W DeGraw, N Buzinsky, B Dodson, M Fertl, A García, G Garvey, B Graner, M Guigue, L Hayen, X Huyan, K S Khaw, K Knutsen, D McClain, D Melconian, P Müller, E Novitski, N S Oblath, R G H Robertson, G Rybka, G Savard, E Smith, D D Stancil, M Sternberg, D W Storm, H E Swanson, J R Tedeschi, B A VanDevender, F E Wietfeldt, A R Young, and X Zhu. First observation of cyclotron radiation from mev-scale e^\pm following nuclear decay. *Physical Review Letters*, 131(8):082502, 2023. URL: <https://app.dimensions.ai/details/publication/pub.1163491872>, doi:10.1103/physrevlett.131.082502.
- [16] F. P. Calaprice, E. D. Commins, and D. C. Girvin. New test of time-reversal invariance in ^{19}Ne beta decay. *Phys. Rev. D*, 9:519–529, Feb 1974. URL: <https://link.aps.org/doi/10.1103/PhysRevD.9.519>, doi:10.1103/PhysRevD.9.519.
- [17] Casper. CASPER, ROACH2 system. <https://casper.astro.berkeley.edu/wiki/ROACH2>, 2013.
- [18] M.J. Christensen and T. Richter. Achieving reliable udp transmission at 10 gb/s using bsd socket for data acquisition systems. *Journal of Instrumentation*, 15(09):T09005, sep 2020. URL: <https://dx.doi.org/10.1088/1748-0221/15/09/T09005>, doi:10.1088/1748-0221/15/09/T09005.
- [19] Vincenzo Cirigliano, David Díaz-Calderón, Adam Falkowski, Martín González-Alonso, and Antonio Rodríguez-Sánchez. Semileptonic tau decays beyond the Standard Model. *Journal of High Energy Physics*, 2022(4):152, apr 2022. URL: <http://arxiv.org/abs/2112.02087>[https://link.springer.com/10.1007/JHEP04\(2022\)152](https://link.springer.com/10.1007/JHEP04(2022)152), arXiv:2112.02087, doi:10.1007/JHEP04(2022)152.
- [20] Vincenzo Cirigliano, Susan Gardner, and Barry R. Holstein. Beta decays and non-standard interactions in the LHC era. *Progress in Particle and Nuclear Physics*, 71:93–118, jul 2013. URL: <https://doi.org/10.1016/Fj.pnpnp.2013.03.005>, doi:10.1016/j.pnpnp.2013.03.005.

- [21] Vincenzo Cirigliano, Susan Gardner, and Barry R. Holstein. Beta decays and non-standard interactions in the LHC era. *Progress in Particle and Nuclear Physics*, 71:93 – 118, 2013. doi:10.1016/j.pnpnp.2013.03.005.
- [22] Vincenzo Cirigliano, Martín González-Alonso, and Michael L. Graesser. Non-standard charged current interactions: beta decays versus the LHC. *Journal of High Energy Physics*, 2013(2), feb 2013. URL: <https://doi.org/10.10072Fjhep0228201329046>, doi:10.1007/jhep02(2013)046.
- [23] Christine Claessens. *Event detection in Project 8*. PhD thesis, 2020.
- [24] Dustin Combs, Gordon Jones, William Anderson, Frank Calaprice, Leendert Hayen, and Albert Young. A look into mirrors: A measurement of the β -asymmetry in ^{19}Ne decay and searches for new physics, 2020. URL: <https://arxiv.org/abs/2009.13700>, doi:10.48550/ARXIV.2009.13700.
- [25] John J. Cowan, Christopher Sneden, James E. Lawler, Ani Aprahamian, Michael Wiescher, Karlheinz Langanke, Gabriel Martínez-Pinedo, and Friedrich-Karl Thielemann. Origin of the heaviest elements: The rapid neutron-capture process. *Rev. Mod. Phys.*, 93:015002, Feb 2021. URL: <https://link.aps.org/doi/10.1103/RevModPhys.93.015002>, doi:10.1103/RevModPhys.93.015002.
- [26] D. Dobson. *The beta-decay asymmetry and nuclear magnetic moment of Neon-19*. PhD thesis, University of California, Berkeley, 1963. URL: <https://babel.hathitrust.org/cgi/pt?id=mdp.39015077591348;view=1up;seq=3>.
- [27] V. P. Dokuchaev. Synchrotron radiation in a circular metallic waveguide. *Radiophysics and Quantum Electronics*, 44(7):542–546, Jul 2001. doi:10.1023/A:1017909800636.
- [28] John F. Donoghue and D. Wyler. Isospin breaking and the precise determination of ν_{ud} . *Physics Letters B*, 241(2):243–248, 1990. URL: <https://www>.

- sciencedirect.com/science/article/pii/037026939091287L, doi:[https://doi.org/10.1016/0370-2693\(90\)91287-L](https://doi.org/10.1016/0370-2693(90)91287-L).
- [29] A. Eggenberger, I. Belosevic, A. Antognini, K. Kirch, and F.M. Piegsa. Temperature dependence of the leak rate and permeability of helium gas through kapton foils. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 798:140 – 143, 2015. URL: <http://www.sciencedirect.com/science/article/pii/S016890021500844X>, doi: <https://doi.org/10.1016/j.nima.2015.07.015>.
- [30] A. Ashtari Esfahani, V. Bansal, S. Böser, N. Buzinsky, R. Cervantes, C. Claessens, L. de Viveiros, P. J. Doe, M. Fertl, J. A. Formaggio, L. Gladstone, M. Guigue, K. M. Heeger, J. Johnston, A. M. Jones, K. Kazkaz, B. H. LaRoque, M. Leber, A. Lindman, E. Machado, B. Monreal, E. C. Morrison, J. A. Nikkel, E. Novitski, N. S. Oblath, W. Pettus, R. G. H. Robertson, G. Rybka, L. Saldaña, V. Sibille, M. Schram, P. L. Slocum, Y.-H. Sun, J. R. Tedeschi, T. Thümmeler, B. A. VanDevender, M. Wachten-donk, M. Walter, T. E. Weiss, T. Wendler, and E. Zayas. Electron radiated power in cyclotron radiation emission spectroscopy experiments. *Phys. Rev. C*, 99:055501, May 2019. URL: <https://link.aps.org/doi/10.1103/PhysRevC.99.055501>, doi: [10.1103/PhysRevC.99.055501](https://doi.org/10.1103/PhysRevC.99.055501).
- [31] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, KDD'96*, page 226–231. AAAI Press, 1996.
- [32] S. Chatrchyan et. al. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters B*, 716(1):30–61, sep 2012. doi:[10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021).
- [33] Adam Falkowski, Martín González-Alonso, and Oscar Naviliat-Cuncic. Comprehensive

- analysis of beta decays within and beyond the Standard Model. *Journal of High Energy Physics*, 2021(4):126, apr 2021. URL: <http://arxiv.org/abs/2010.13797>[https://link.springer.com/10.1007/JHEP04\(2021\)126](https://link.springer.com/10.1007/JHEP04(2021)126), arXiv:2010.13797, doi:10.1007/JHEP04(2021)126.
- [34] B. Fenker, A. Gorelov, D. Melconian, J. A. Behr, M. Anholm, D. Ashery, R. S. Behling, I. Cohen, I. Craiciu, G. Gwinner, J. McNeil, M. Mehlman, K. Olchanski, P. D. Shidling, S. Smale, and C. L. Warner. Precision Measurement of the β Asymmetry in Spin-Polarized ^{37}K Decay. *Phys. Rev. Lett.*, 120:062502, Feb 2018. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.062502>, doi:10.1103/PhysRevLett.120.062502.
- [35] Markus Fierz. Zur Fermischen Theorie des β -Zerfalls. *Zeitschrift fur Physik*, 104(7-8):553–565, July 1937. doi:10.1007/BF01330070.
- [36] Xavier Flechard. personal communication, 2020.
- [37] Fry, J., Alarcon, R., Baeßler, S., Balascuta, S., Palos, L. Barrón, Bailey, T., Bass, K., Birge, N., Blose, A., Borissenko, D., Bowman, J.D., Broussard, L.J., Bryant, A.T., Byrne, J., Calarco, J.R., Caylor, J., Chang, K., Chupp, T., Cianciolo, T.V., Crawford, C., Ding, X., Doyle, M., Fan, W., Farrar, W., Fomin, N., Frlez, E., Gericke, M.T., Gervais, M., Glück, F., Greene, G.L., Grzywacz, R.K., Gudkov, V., Hamblen, J., Hayes, C., Hendrus, C., Ito, T., Jezghani, A., Li, H., Makela, M., Macsai, N., Mammei, J., Mammei, R., Martinez, M., Matthews, D.G., McCrea, M., McGaughey, P., McLaughlin, C.D., Mueller, P., Petten, D. van, Penttilä, S.I., Perryman, D.E., Picker, R., Pierce, J., Pocanić, D., Qian, Y., Ramsey, J., Randall, G., Riley, G., Rykaczewski, K.P., Salas-Bacci, A., Samiei, S., Scott, E.M., Shelton, T., Sjue, S.K., Smith, A., Smith, E., Stevens, E., Wexler, J., Whitehead, R., Wilburn, W.S., Young, A., and Zeck, B. The Nab experiment: A precision measurement of unpolarized neutron beta decay. *EPJ Web Conf.*, 219:04002, 2019. doi:10.1051/epjconf/201921904002.

- [38] F. H. Garcia, C. Andreoiu, G. C. Ball, A. Bell, A. B. Garnsworthy, F. Nowacki, C. M. Petrache, A. Poves, K. Whitmore, F. A. Ali, N. Bernier, S. S. Bhattacharjee, M. Bowry, R. J. Coleman, I. Dillmann, I. Djianto, A. M. Forney, M. Gascoine, G. Hackman, K. G. Leach, A. N. Murphy, C. R. Natzke, B. Olaizola, K. Ortner, E. E. Peters, M. M. Rajabali, K. Raymond, C. E. Svensson, R. Umashankar, J. Williams, and D. Yates. Absence of low-energy shape coexistence in ^{80}Ge : The nonobservation of a proposed excited 0_2^+ level at 639 keV. *Phys. Rev. Lett.*, 125:172501, Oct 2020. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.172501>, doi:10.1103/PhysRevLett.125.172501.
- [39] E. A. George, P. A. Voytas, G. W. Severin, and L. D. Knutson. Measurement of the shape factor for the β decay of ^{14}O . *Phys. Rev. C*, 90:065501, Dec 2014. doi:10.1103/PhysRevC.90.065501.
- [40] D.C. Girvin. *Test of time-reversal invariance in ^{19}Ne beta decay*. PhD thesis, University of California, Berkeley, 1972.
- [41] Ayala Glick-Magid, Christian Forssén, Daniel Gazda, Doron Gazit, Peter Gysbers, and Petr Navrátil. Nuclear ab initio calculations of 6He -decay for beyond the standard model studies. *Physics Letters B*, 832:137259, 2022. URL: <https://www.sciencedirect.com/science/article/pii/S0370269322003938>, doi:<https://doi.org/10.1016/j.physletb.2022.137259>.
- [42] F. Glück, I. Joó, and J. Last. Measurable parameters of neutron decay. *Nuclear Physics A*, 593(2):125–150, 1995. URL: <https://www.sciencedirect.com/science/article/pii/0375947495003544>, doi:[https://doi.org/10.1016/0375-9474\(95\)00354-4](https://doi.org/10.1016/0375-9474(95)00354-4).
- [43] M. González-Alonso and J. Martin Camalich. Isospin breaking in the nucleon mass and the sensitivity of β decays to new physics. *Phys. Rev. Lett.*, 112:042501, Jan

2014. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.112.042501>, doi: 10.1103/PhysRevLett.112.042501.
- [44] M. González-Alonso and O. Naviliat-Cuncic. Kinematic sensitivity to the fierz term of β -decay differential spectra. *Phys. Rev. C*, 94:035503, Sep 2016. URL: <https://link.aps.org/doi/10.1103/PhysRevC.94.035503>, doi:10.1103/PhysRevC.94.035503.
- [45] Martin Gonzalez-Alonso, Oscar Naviliat-Cuncic, and Nathal Severijns. New physics searches in nuclear and neutron β decay. *Progress in Particle and Nuclear Physics*, xxx(x):xxxx, 2018. 1. URL: <https://arxiv.org/abs/1803.08732>.
- [46] A. Gorelov, D. Melconian, W. P. Alford, D. Ashery, G. Ball, J. A. Behr, P. G. Bricault, J. M. D’Auria, J. Deutsch, J. Dilling, M. Dombisky, P. Dubé, J. Fingler, U. Giesen, F. Glück, S. Gu, O. Häusser, K. P. Jackson, B. K. Jennings, M. R. Pearson, T. J. Stocki, T. B. Swanson, and M. Trinczek. Scalar Interaction Limits from the β - ν Correlation of Trapped Radioactive Atoms. *Phys. Rev. Lett.*, 94:142501, Apr 2005. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.94.142501>, doi:10.1103/PhysRevLett.94.142501.
- [47] D. Griffiths. *Introduction to Elementary Particles*. John Wiley Sons, New York, USA, 1987.
- [48] Francis Halzen and Alan Martin. *Quarks & Leptons: An introductory course in modern particle physics*. John Wiley & Sons, New York, USA, 1984.
- [49] J. C. Hardy and I. S. Towner. Superaligned $0^+ \rightarrow 0^+$ nuclear β decays: 2014 critical survey, with precise results for V_{ud} and ckm unitarity. *Phys. Rev. C*, 91:025501, Feb 2015. URL: <https://link.aps.org/doi/10.1103/PhysRevC.91.025501>, doi: 10.1103/PhysRevC.91.025501.
- [50] L. Hayen, J. Kostensalo, N. Severijns, and J. Suhonen. First-forbidden transitions in

- the reactor anomaly. *Phys. Rev. C*, 100:054323, Nov 2019. URL: <https://link.aps.org/doi/10.1103/PhysRevC.100.054323>, doi:10.1103/PhysRevC.100.054323.
- [51] L. Hayen and N. Severijns. Beta spectrum generator: High precision allowed spectrum shapes. *Computer Physics Communications*, 240:152–164, 2019. URL: <https://www.sciencedirect.com/science/article/pii/S0010465519300645>, doi:<https://doi.org/10.1016/j.cpc.2019.02.012>.
- [52] Leendert Hayen, Nathal Severijns, Kazimierz Bodek, Dagmara Rozpedzik, and Xavier Mougeot. High precision analytical description of the allowed β spectrum shape. *Rev. Mod. Phys.*, 90:015008, Mar 2018. URL: <https://link.aps.org/doi/10.1103/RevModPhys.90.015008>, doi:10.1103/RevModPhys.90.015008.
- [53] A. C. Hayes, J. L. Friar, G. T. Garvey, Duligur Ibeling, Gerard Jungman, T. Kawano, and Robert W. Mills. Possible origins and implications of the shoulder in reactor neutrino spectra. *Phys. Rev. D*, 92:033015, Aug 2015. URL: <https://link.aps.org/doi/10.1103/PhysRevD.92.033015>, doi:10.1103/PhysRevD.92.033015.
- [54] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn, 2017. URL: <https://arxiv.org/abs/1703.06870>, doi:10.48550/ARXIV.1703.06870.
- [55] Peter W. Higgs. Spontaneous symmetry breakdown without massless bosons. *Phys. Rev.*, 145:1156–1163, May 1966. URL: <https://link.aps.org/doi/10.1103/PhysRev.145.1156>, doi:10.1103/PhysRev.145.1156.
- [56] Barry R. Holstein and S. B. Treiman. Tests of spontaneous left-right-symmetry breaking. *Phys. Rev. D*, 16:2369–2372, Oct 1977. doi:10.1103/PhysRevD.16.2369.
- [57] Ran Hong, Matthew G. Sternberg, and Alejandro Garcia. Helicity and nuclear decay correlations. *American Journal of Physics*, 85(1):45–53, jan 2017. URL: <https://doi.org/10.11192F1.4966197>, doi:10.1119/1.4966197.

- [58] X. Huyan, O. Naviliat-Cuncic, P. Voytas, S. Chandavar, M. Hughes, K. Minamisono, and S.V. Paulauskas. Geant4 simulations of the absorption of photons in CsI and NaI produced by electrons with energies up to 4 MeV and their application to precision measurements of the β -energy spectrum with a calorimetric technique. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 879:134 – 140, 2018. URL: <http://www.sciencedirect.com/science/article/pii/S0168900217311440>, doi: <https://doi.org/10.1016/j.nima.2017.10.061>.
- [59] Xueying Huyan. *PRECISION MEASUREMENT OF THE β -ENERGY SPECTRUM IN ${}^6\text{He}$ DECAY*. PhD thesis, 2019.
- [60] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift, 2015. URL: <https://arxiv.org/abs/1502.03167>, doi:10.48550/ARXIV.1502.03167.
- [61] J. D. Jackson, S. B. Treiman, and H. W. Wyld. Possible tests of time reversal invariance in beta decay. *Phys. Rev.*, 106:517–521, May 1957. URL: <https://link.aps.org/doi/10.1103/PhysRev.106.517>, doi:10.1103/PhysRev.106.517.
- [62] John David Jackson. *Classical electrodynamics*. Wiley, New York, NY, 3rd ed. edition, 1999. URL: <http://cdsweb.cern.ch/record/490457>.
- [63] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An Introduction to Statistical Learning: with Applications in R*. Springer, 2013. URL: <https://faculty.marshall.usc.edu/gareth-james/ISL/>.
- [64] J.G. Jenkin, L.G. Earwaker, and E.W. Titterton. The $\text{F}^{19}(\text{p}, \text{n})\text{Ne}^{19}$ reaction between 5 and 11 mev. *Nuclear Physics*, 44:453 – 459, 1963. URL: <http://www.sciencedirect.com/science/article/pii/0029558263900397>, doi:[https://doi.org/10.1016/0029-5582\(63\)90039-7](https://doi.org/10.1016/0029-5582(63)90039-7).

- [65] M. Kanafani, X. Flé chard, O. Naviliat-Cuncic, G. D. Chung, S. Leblond, E. Liénard, X. Mougeot, G. Quéméner, A. Simancas Di Filippo, and J.-C. Thomas. High-precision measurement of the $t_{1/2}$ of ^6He . *Physical Review C*, 106(4), oct 2022. URL: <https://doi.org/10.1103/PhysRevC.106.045502>, doi:10.1103/PhysRevC.106.045502.
- [66] Garrett B. King, Alessandro Baroni, Vincenzo Cirigliano, Stefano Gandolfi, Leendert Hayen, Emanuele Mereghetti, Saori Pastore, and Maria Piarulli. Ab initio calculation of the decay spectrum of ^6He , 2022. URL: <https://arxiv.org/abs/2207.11179>, doi:10.48550/ARXIV.2207.11179.
- [67] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. URL: <https://arxiv.org/abs/1412.6980>, doi:10.48550/ARXIV.1412.6980.
- [68] A. Knecht, R. Hong, D. W. Zumwalt, B. G. Delbridge, A. García, P. Müller, H. E. Swanson, I. S. Towner, S. Utsuno, W. Williams, and C. Wrede. Precision measurement of the ^6He half-life and the weak axial current in nuclei. *Phys. Rev. C*, 86:035506, Sep 2012. URL: <https://link.aps.org/doi/10.1103/PhysRevC.86.035506>, doi:10.1103/PhysRevC.86.035506.
- [69] A. Knecht, R. Hong, D. W. Zumwalt, B. G. Delbridge, A. García, P. Müller, H. E. Swanson, I. S. Towner, S. Utsuno, W. Williams, and C. Wrede. Precision Measurement of the ^6He Half-Life and the Weak Axial Current in Nuclei. *Phys. Rev. Lett.*, 108:122502, Mar 2012. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.108.122502>, doi:10.1103/PhysRevLett.108.122502.
- [70] A. Knecht, D.W. Zumwalt, B.G. Delbridge, A. García, G.C. Harper, R. Hong, P. Müller, A.S.C. Palmer, R.G.H. Robertson, H.E. Swanson, S. Utsuno, D.I. Will, W. Williams, and C. Wrede. A high-intensity source of ^6He atoms

- for fundamental research. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 660(1):43–47, 2011. URL: <https://www.sciencedirect.com/science/article/pii/S0168900211017888>, doi:<https://doi.org/10.1016/j.nima.2011.09.033>.
- [71] L. D. Knutson, G. W. Severin, S. L. Cotter, Li Zhan, P. A. Voytas, and E. A. George. A superconducting beta spectrometer. *Review of Scientific Instruments*, 82(7):073302, 2011. arXiv:<https://doi.org/10.1063/1.3608454>, doi:10.1063/1.3608454.
- [72] Kuniharu Kubodera and Tae-Sun Park. The solar hep process. *Annual Review of Nuclear and Particle Science*, 54(1):19–37, 2004. arXiv:<https://doi.org/10.1146/annurev.nucl.54.070103.181239>, doi:10.1146/annurev.nucl.54.070103.181239.
- [73] A. Kurylov and M. J. Ramsey-Musolf. Charged current universality in the minimal supersymmetric standard model. *Phys. Rev. Lett.*, 88:071804, Feb 2002. doi:10.1103/PhysRevLett.88.071804.
- [74] D. Lubos, J. Park, T. Faestermann, R. Gernhäuser, R. Krücken, M. Lewitowicz, S. Nishimura, H. Sakurai, D. S. Ahn, H. Baba, B. Blank, A. Blazhev, P. Boutachkov, F. Browne, I. Čeliković, G. de France, P. Doornenbal, Y. Fang, N. Fukuda, J. Giovinazzo, N. Goel, M. Górska, S. Ilieva, N. Inabe, T. Isobe, A. Jungclaus, D. Kameda, Y. K. Kim, I. Kojouharov, T. Kubo, N. Kurz, Y. K. Kwon, G. Lorusso, K. Moschner, D. Murai, I. Nishizuka, Z. Patel, M. M. Rajabali, S. Rice, H. Schaffner, Y. Shimizu, L. Sinclair, P.-A. Söderström, K. Steiger, T. Sumikama, H. Suzuki, H. Takeda, Z. Wang, N. Warr, H. Watanabe, J. Wu, and Z. Xu. Improved value for the gamow-teller strength of the ^{100}Sn beta decay. *Phys. Rev. Lett.*, 122:222502, Jun 2019. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.122.222502>, doi:10.1103/PhysRevLett.122.222502.
- [75] Benjamin Monreal and Joseph A. Formaggio. Relativistic cyclotron radiation detection

- of tritium decay electrons as a new technique for measuring the neutrino mass. *Phys. Rev. D*, 80:051301, Sep 2009. URL: <https://link.aps.org/doi/10.1103/PhysRevD.80.051301>, doi:10.1103/PhysRevD.80.051301.
- [76] Daniel Moser, Hartmut Abele, Joachim Bosina, Harald Fillunger, Torsten Soldner, Xiangzun Wang, Johann Zmeskal, and Gertrud Konrad. NoMoS: An $R \times B$ drift momentum spectrometer for beta decay studies. *EPJ Web Conf.*, 219:04003, 2019. doi:10.1051/epjconf/201921904003.
- [77] P. Müller, Y. Bagdasarova, R. Hong, A. Leredde, K. G. Bailey, X. Fléchar, A. García, B. Graner, A. Knecht, O. Naviliat-Cuncic, T.P. O'Connor, M.G. Sternberg, D.W. Storm, H. E. Swanson, F. Wauters, and D.W. Zumwalt. β -nuclear-recoil correlation from ${}^6\text{He}$ decay in a laser trap. *Phys. Rev. Lett.*, Jul 2022. in press. URL: <https://journals.aps.org/prl/accepted/84075Ye4A171bc7a82ed55f023b7bd042d26bab91>.
- [78] O. Naviliat-Cuncic and N. Severijns. Test of the conserved vector current hypothesis in $t = 1/2$ mirror transitions and new determination of $|V_{ud}|$. *Phys. Rev. Lett.*, 102:142302, Apr 2009. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.102.142302>, doi:10.1103/PhysRevLett.102.142302.
- [79] O. Naviliat-Cuncic and N. Severijns. Test of the Conserved Vector Current Hypothesis in $T = 1/2$ Mirror Transitions and New Determination of $|V_{ud}|$. *Phys. Rev. Lett.*, 102:142302, Apr 2009. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.102.142302>, doi:10.1103/PhysRevLett.102.142302.
- [80] Matt Newville, Renee Otten, Andrew Nelson, Till Stensitzki, Antonino Ingargiola, Dan Allan, Austin Fox, Faustin Carter, Michał, Ray Osborn, Dima Pustakhod, Ineuhaus, Sebastian Weigand, Andrey Aristov, Glenn, Christoph Deil, mgunyho, Mark, Allan L. R. Hansen, Gustavo Pasquevich, Leon Foks, Nicholas Zobrist, Oliver Frost, Stuermer, azelcer, Anthony Polloreno, Arun Persaud, Jens Hedegaard Nielsen, Matteo Pom-

- pili, and Pieter Eendebak. `lmfit/lmfit-py`: 1.2.2, July 2023. doi:10.5281/zenodo.8145703.
- [81] Noah S. Oblath, Nick Buzinsky, Christine Claessens, Laura Gladstone, Mathieu Guigue, Jiarui Huang, Brittney Johnson, Jared Kofron, Benjamin LaRoque, Elise Novitski, Devyn Rysewyk, Luis Saldaña, Penny Slocum, Yu-Hao Sun, Brent Vandevender, Luiz de Viveiros, and Evan Zayas. `Katydid`: Project 8 analysis software package. <https://doi.org/10.5281/zenodo.2621293>[doi.org], April 2019. URL: <https://doi.org/10.5281/zenodo.2621293>[doi.org], doi:10.5281/zenodo.2621293.
- [82] K.A. Olive. Review of particle physics. *Chinese Physics C*, 38(9):090001, aug 2014. URL: <https://dx.doi.org/10.1088/1674-1137/38/9/090001>, doi:10.1088/1674-1137/38/9/090001.
- [83] P.D. O'Malley, M. Brodeur, D.P. Burdette, J.W. Klimes, A.A. Valverde, J.A. Clark, G. Savard, R. Ringle, and V. Varentsov. Testing the weak interaction using St. Benedict at the University of Notre Dame. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 463:488–490, 2020. URL: <https://www.sciencedirect.com/science/article/pii/S0168583X19302010>, doi:<https://doi.org/10.1016/j.nimb.2019.04.017>.
- [84] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- [85] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. `Scikit-learn`: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [86] Michael E. Peskin. On the trail of the higgs boson. *Annalen der Physik*, 528(1-

- 2):20–34, 2016. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.201500225>, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.201500225>, doi:<https://doi.org/10.1002/andp.201500225>.
- [87] Michael Edward Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Westview Press, 1995. Reading, USA: Addison-Wesley (1995) 842 p.
- [88] David M Pozar. *Microwave engineering; 3rd ed.* Wiley, Hoboken, NJ, 2005. URL: <https://cds.cern.ch/record/882338>.
- [89] Dinko Počanić, R. Alarcon, L.P. Alonzi, S. Baeßler, S. Balascuta, J.D. Bowman, M.A. Bychkov, J. Byrne, J.R. Calarco, V. Cianciolo, C. Crawford, E. Frlež, M.T. Gericke, G.L. Greene, R.K. Grzywacz, V. Gudkov, F.W. Hersman, A. Klein, J. Martin, S.A. Page, A. Palladino, S.I. Penttilä, K.P. Rykaczewski, W.S. Wilburn, A.R. Young, and G.R. Young. Nab: Measurement principles, apparatus and uncertainties. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 611(2):211 – 215, 2009. Particle Physics with Slow Neutrons. URL: <http://www.sciencedirect.com/science/article/pii/S0168900209015289>, doi:<https://doi.org/10.1016/j.nima.2009.07.065>.
- [90] S. Profumo, M. J. Ramsey-Musolf, and S. Tulin. Supersymmetric contributions to weak decay correlation coefficients. *Phys. Rev. D*, 75:075017, Apr 2007. doi:10.1103/PhysRevD.75.075017.
- [91] Project 8 Collaboration, A. Ashtari Esfahani, S. Böser, N. Buzinsky, M. C. Carmona-Benitez, C. Claessens, L. de Viveiros, P. J. Doe, M. Fertl, J. A. Formaggio, J. K. Gaison, L. Gladstone, M. Grando, M. Guigue, J. Hartse, K. M. Heeger, X. Huyan, J. Johnston, A. M. Jones, K. Kazkaz, B. H. LaRoque, M. Li, A. Lindman, E. Machado, A. Marsteller, C. Matthé, R. Mohiuddin, B. Monreal, R. Mueller, J. A. Nikkel, E. Novitski, N. S. Oblath, J. I. Peña, W. Pettus, R. Reimann, R. G. H. Robertson, D. Rosa De Jesús, G. Rybka, L. Saldaña, M. Schram, P. L. Slocum, J. Stachurska,

- Y. H. Sun, P. T. Surukuchi, J. R. Tedeschi, A. B. Telles, F. Thomas, M. Thomas, L. A. Thorne, T. Thümmler, L. Tvrznikova, W. Van De Pontseele, B. A. VanDevender, J. Weintraub, T. E. Weiss, T. Wendler, A. Young, E. Zayas, and A. Ziegler. Tritium beta spectrum and neutrino mass limit from cyclotron radiation emission spectroscopy, 2022. URL: <https://arxiv.org/abs/2212.05048>, doi:10.48550/ARXIV.2212.05048.
- [92] Michael J. Ramsey-Musolf, Shufang Su, and Sean Tulin. Pion leptonic decays and supersymmetry. *Phys. Rev. D*, 76:095017, Nov 2007. doi:10.1103/PhysRevD.76.095017.
- [93] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. 2018. cite arxiv:1804.02767Comment: Tech Report. URL: <http://arxiv.org/abs/1804.02767>.
- [94] F.L. Riffle, J.D. Goss, D.R. Parsignault, and J.C. Harris. Measurement of the absolute total cross section for the reaction $^{19}\text{F}(p, n)^{19}\text{Ne}$. *Nuclear Physics A*, 115(1):120–128, 1968. URL: <https://www.sciencedirect.com/science/article/pii/0375947468906465>, doi:[https://doi.org/10.1016/0375-9474\(68\)90646-5](https://doi.org/10.1016/0375-9474(68)90646-5).
- [95] William J. Riley and Physics Laboratory (U.S.). *Handbook of frequency stability analysis [electronic resource] / W.J. Riley*. U.S. Dept. of Commerce, National Institute of Standards and Technology Boulder, Colo, 2008.
- [96] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015. cite arxiv:1505.04597Comment: conditionally accepted at MICCAI 2015. URL: <http://arxiv.org/abs/1505.04597>.
- [97] E. Rutherford. The scattering of α and β particles by matter and the structure of the atom. *Philosophical Magazine*, 92(4):379–398, 2012. doi:10.1080/14786435.2011.617037.

- [98] Luis Saldana. *Novel Signal Reconstruction Techniques in Cyclotron Radiation Emission Spectroscopy for Neutrino Mass Measurement*. PhD thesis, Yale University, 2021.
- [99] H. Saul, C. Roick, H. Abele, H. Mest, M. Klopf, A. K. Petukhov, T. Soldner, X. Wang, D. Werder, and B. Märkisch. Limit on the Fierz Interference Term b from a Measurement of the Beta Asymmetry in Neutron Decay. *Phys. Rev. Lett.*, 125:112501, Sep 2020. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.112501>, doi:10.1103/PhysRevLett.125.112501.
- [100] M. B. Schneider, F. P. Calaprice, A. L. Hallin, D. W. MacArthur, and D. F. Schreiber. Limit on $\text{Im}(C_S C_A^*)$ from a Test of T Invariance in ^{19}Ne Beta Decay. *Phys. Rev. Lett.*, 51:1239–1242, Oct 1983. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.51.1239>, doi:10.1103/PhysRevLett.51.1239.
- [101] Steven J. Schowalter, Colin B. Connolly, and John M. Doyle. Permeability of noble gases through kapton, butyl, nylon, and “silver shield”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 615(3):267 – 271, 2010. URL: <http://www.sciencedirect.com/science/article/pii/S0168900210001099>, doi:<https://doi.org/10.1016/j.nima.2010.01.041>.
- [102] S. H. Seo, W. Q. Choi, H. Seo, J. H. Choi, Y. Choi, H. I. Jang, J. S. Jang, K. K. Joo, B. R. Kim, H. S. Kim, J. Y. Kim, S. B. Kim, S. Y. Kim, W. Kim, E. Kwon, D. H. Lee, Y. C. Lee, I. T. Lim, M. Y. Pac, I. G. Park, J. S. Park, R. G. Park, Y. G. Seon, C. D. Shin, J. H. Yang, J. Y. Yang, I. S. Yeo, and I. Yu. Spectral measurement of the electron antineutrino oscillation amplitude and frequency using 500 live days of reno data. *Phys. Rev. D*, 98:012002, Jul 2018. URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.012002>, doi:10.1103/PhysRevD.98.012002.
- [103] Nathal Severijns, Marcus Beck, and Oscar Naviliat-Cuncic. Tests of the standard electroweak model in nuclear beta decay. *Reviews of Modern Physics*, 78(3):991–

- 1040, sep 2006. URL: <https://doi.org/10.11032Frevmodphys.78.991>, doi:10.1103/revmodphys.78.991.
- [104] P.D. Shidling, M. Mehlman, V.S. Kolhinen, G. Chubarian, L. Cooper, G. Duran, E. Gilg, V.E. Jacob, K.S. Marble, R. McAfee, D. McClain, M. McDonough, M. Nasser, C. Gonzalez-Ortiz, A. Ozmetin, B. Schroeder, M. Soulard, G. Tabacaru, and D. Melconian. The TAMUTRAP facility: A Penning trap facility at Texas A&M University for weak interaction studies. *International Journal of Mass Spectrometry*, 468:116636, 2021. URL: <https://www.sciencedirect.com/science/article/pii/S1387380621001160>, doi:<https://doi.org/10.1016/j.ijms.2021.116636>.
- [105] X. Sun, E. Adamek, B. Allgeier, Y. Bagdasarova, D. B. Berguno, M. Blatnik, T. J. Bowles, L. J. Broussard, M. A.-P. Brown, R. Carr, S. Clayton, C. Cude-Woods, S. Currie, E. B. Dees, X. Ding, B. W. Filippone, A. García, P. Geltenbort, S. Hasan, K. P. Hickerson, J. Hoagland, R. Hong, A. T. Holley, T. M. Ito, A. Knecht, C.-Y. Liu, J. Liu, M. Makela, R. Mammei, J. W. Martin, D. Melconian, M. P. Mendenhall, S. D. Moore, C. L. Morris, S. Nepal, N. Nouri, R. W. Pattie, A. Pérez Gálvan, D. G. Phillips, R. Picker, M. L. Pitt, B. Plaster, D. J. Salvat, A. Saunders, E. I. Sharapov, S. Sjue, S. Slutsky, W. Sondheim, C. Swank, E. Tatar, R. B. Vogelaar, B. VornDick, Z. Wang, W. Wei, J. W. Wexler, T. Womack, C. Wrede, A. R. Young, and B. A. Zeck. Improved limits on Fierz interference using asymmetry measurements from the Ultracold Neutron Asymmetry (UCNA) experiment. *Phys. Rev. C*, 101:035503, Mar 2020. URL: <https://link.aps.org/doi/10.1103/PhysRevC.101.035503>, doi:10.1103/PhysRevC.101.035503.
- [106] A. W. Thomas. Cvc in particle physics, 1996. URL: <https://arxiv.org/abs/nucl-th/9609052>, doi:10.48550/ARXIV.NUCL-TH/9609052.
- [107] K. K. Vos, H. W. Wilschut, and R. G. E. Timmermans. Symmetry violations in

- nuclear and neutron β decay. *Rev. Mod. Phys.*, 87:1483–1516, Dec 2015. doi:10.1103/RevModPhys.87.1483.
- [108] K. K. Vos, H. W. Wilschut, and R. G. E. Timmermans. Symmetry violations in nuclear and neutron β decay. *Reviews of Modern Physics*, 87(4):1483–1516, dec 2015. URL: <https://doi.org/10.1103/RevModPhys.87.1483>, doi:10.1103/revmodphys.87.1483.
- [109] P. A. Voytas, E. A. George, G. W. Severin, L. Zhan, and L. D. Knutson. Measurement of the branching ratio for the β decay of ^{14}O . *Phys. Rev. C*, 92:065502, Dec 2015. doi:10.1103/PhysRevC.92.065502.
- [110] D. Vénos, A. Špalek, O. Lebeda, and M. Fišer. $^{83\text{m}}\text{Kr}$ radioactive source based on ^{83}Rb trapped in cation-exchange paper or in zeolite. *Applied Radiation and Isotopes*, 63(3):323–327, 2005. URL: <https://www.sciencedirect.com/science/article/pii/S0969804305001260>, doi:<https://doi.org/10.1016/j.apradiso.2005.04.011>.
- [111] Xiangzun Wang, Carmen Ziener, Hartmut Abele, Stefan Bodmaier, Dirk Dubbers, Jaqueline Erhart, Alexander Hollering, Erwin Jericha, Jens Klenke, Harald Fillunger, Werner Heil, Christine Klauser, Gertrud Konrad, Max Lamparth, Thorsten Lauer, Michael Klopff, Reinhard Maix, Bastian Märkisch, Wilfried Mach, Holger Mest, Daniel Moser, Alexander Pethoukov, Lukas Raffelt, Nataliya Rebrova, Christoph Roick, Heiko Saul, Ulrich Schmidt, Torsten Soldner, Romain Viot, Oliver Zimmer, and (the PERC collaboration). Design of the magnet system of the neutron decay facility PERC. *EPJ Web Conf.*, 219:04007, 2019. doi:10.1051/epjconf/201921904007.
- [112] Adam West. Lepton dipole moments, 2016. URL: <https://arxiv.org/abs/1607.00925>, doi:10.48550/ARXIV.1607.00925.

- [113] Victor Wiley and Thomas Lucas. Computer vision and image processing: A paper review. *International Journal of Artificial Intelligence Research*, 2:22, 02 2018. doi: 10.29099/ijair.v2i1.42.
- [114] D.H. Wilkinson and B.E.F. Macefield. A parametrization of the phase space factor for allowed β -decay. *Nuclear Physics A*, 232(1):58–92, 1974. URL: <https://www.sciencedirect.com/science/article/pii/0375947474906459>, doi:[https://doi.org/10.1016/0375-9474\(74\)90645-9](https://doi.org/10.1016/0375-9474(74)90645-9).
- [115] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105:1413–1415, Feb 1957. URL: <https://link.aps.org/doi/10.1103/PhysRev.105.1413>, doi: 10.1103/PhysRev.105.1413.
- [116] C. S. Wu, B. M. Rustad, V. Perez-Mendez, and L. Lidofsky. The beta-spectrum of He^6 . *Phys. Rev.*, 87:1140–1141, Sep 1952. URL: <https://link.aps.org/doi/10.1103/PhysRev.87.1140>, doi:10.1103/PhysRev.87.1140.
- [117] D. Zumwalt. *Developments for the ${}^6\text{He}$ β, ν angular correlation experiment*. PhD thesis, University of Washington, Sep 2015.