

©Copyright 2023

Ethan T. Muldoon

MuSun: A Complete Analysis of the Precision Measurement of
Nuclear Muon Capture in Deuterium.

Ethan T. Muldoon

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

Doctor of Philosophy

University of Washington

2023

Reading Committee:

Peter Kammel, Chair

David Hertzog

Silas Beane

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

MuSun: A Complete Analysis of the Precision Measurement of Nuclear Muon Capture in Deuterium.

Ethan T. Muldoon

Chair of the Supervisory Committee:
Research Professor Peter Kammel
Physics

The MuSun experiment aims to measure the muon capture rate on deuterium to 1.5% precision. MuSun will be the first experiment to meaningfully constrain the weak axial current in the two-nucleon system, which is relevant for Effective Field Theory approaches to studying light nuclei and is also related to solar proton-proton fusion and to the neutrino-deuterium scattering reactions measured at the Sudbury Neutrino Observatory. The capture rate is determined via a 10 ppm measurement of the negative muon disappearance rate in a deuterium target, which differs from the free muon decay rate by approximately 0.1% due to muon capture providing an additional decay channel. Such precision is achieved by measuring over 10^{10} muons, using a cryogenic time-projection chamber to ensure clean stops in deuterium gas. This thesis describes the experimental design and hardware, and gives an overview of the main production runs in 2014 and 2015 that collected the full statistics needed for the experiment. It then explains the operation of the MuSun analysis software and some aspects of software development. Finally several in-depth analyses constraining specific sources of systematic uncertainty are presented. These focus primarily on characterizing the effects of mis-reconstructed muon stop positions, as well as background signals produced by capture on heavier elements when muons stop outside the target. These analyses establish uncertainty estimates near the target level, and a final unblinded result will be produced in 2023.

TABLE OF CONTENTS

	Page
Chapter 1: Introduction	1
Chapter 2: Theory and Motivation	7
2.1 Weak Interaction	7
2.1.1 Muon Decay	9
2.2 Muon Capture	11
2.2.1 Single Nucleon	13
2.2.2 Light Nuclei	15
2.2.3 Heavy Nuclei	17
2.3 Effective Field Theory	18
2.3.1 QCD Symmetries	18
2.3.2 Chiral Perturbation Theory	19
2.3.3 Pionless EFT	22
2.4 History of Λ_d Calculations	25
2.5 Significance of MuSun	26
2.6 Summary	28
Chapter 3: Experimental Design Considerations	29
3.1 Lifetime Method	29
3.2 Precision Goal	31
3.3 Muon Kinetics	32
3.3.1 Hyperfine States	33
3.3.2 Muon Catalyzed Fusion	34
3.3.3 Optimized Running Conditions	36
3.4 Capture on Other Nuclei	37
3.4.1 Wall Stops	37

3.4.2	Gas Purity	40
3.5	Summary	42
Chapter 4:	Experimental Hardware	43
4.1	Accelerator	44
4.2	Beamline	46
4.3	Entrance Detectors	48
4.3.1	Detectors	49
4.3.2	Muon-On-Request System	50
4.3.3	Pileup Protection	51
4.4	Target	51
4.4.1	Time Projection Chamber	52
4.4.2	Spin Precession	52
4.5	Electron Detectors	53
4.5.1	Electron Trackers (ePCs)	53
4.5.2	Electron Scintillator (eSC)	54
4.6	Neutron Detectors	55
4.7	Gas Purification System	56
4.8	Computer Systems	58
4.8.1	Data Acquisition System	59
4.8.2	Online Analysis	60
4.9	Blinding Procedure	60
4.10	Summary	61
Chapter 5:	Time Projection Chamber	62
5.1	Constraints	62
5.1.1	Drift Field	63
5.1.2	Frisch Grid	64
5.1.3	Gas Amplification	64
5.2	Final Design	65
5.2.1	Geometry	65
5.2.2	Materials	65
5.2.3	Readout	68

5.3	Detector Response	69
5.4	Energy Signatures	71
5.4.1	Row Energies	71
5.5	Fusion Products	73
5.5.1	S-Energy	74
5.5.2	E_9 Energy	75
5.6	Electron Tracks	75
5.7	Summary	76
Chapter 6: Data Collection		77
6.1	Early Runs	78
6.2	R2014	78
6.3	R2015	80
6.3.1	Changes Since R2014	81
6.3.2	Datasets	83
6.4	R2016	83
6.4.1	Neutron Calibration Targets	84
6.4.2	Impurity Injection	84
6.4.3	Beam Background Tests	85
6.5	Simulation	86
6.6	Summary	87
Chapter 7: Low-Level Analysis Software		88
7.1	Analysis Framework	88
7.2	MU	91
7.3	TPC WFD Pulses	91
7.3.1	Pedestal Calibration	92
7.3.2	Dual Gain Processing	93
7.3.3	Pulse Processing	95
7.4	eSC WFD Pulses	97
7.4.1	Pulse Processing	98
7.4.2	Clock Alignment	98
7.5	ePC Clustering	98

7.6	Entrance Detector Analysis	99
7.6.1	μ SC	99
7.6.2	μ SCa	100
7.6.3	μ PC	100
7.6.4	Muon Entrances	101
7.7	Outputs	103
7.7.1	Event Tree	103
7.8	Summary	105
Chapter 8: High-Level Analysis Software		106
8.1	TPC Tracking	106
8.1.1	Clustering	107
8.1.2	Tracking Algorithms	108
8.1.3	Absolute Gain Calibration	111
8.2	Electron Tracking	113
8.2.1	eSC Coincidences	114
8.2.2	eSC Position	116
8.2.3	ePC Tracking	116
8.2.4	Electron Definitions	117
8.3	Neutron Detectors	118
8.3.1	Pulse Finding	119
8.3.2	Interpolation	120
8.3.3	Particle Identification	121
8.4	Outputs	123
8.4.1	Skimming	124
8.4.2	Lifetime Tree	125
8.5	Stage Three	126
8.6	Summary	127
Chapter 9: Muon Tracking		128
9.1	Formalism	128
9.1.1	Fusions Event Decay Times	129
9.1.2	Acceptance Function	130

9.1.3	Migration	134
9.1.4	Electron Interference	135
9.2	TPC Energy Deposition	136
9.2.1	E_0 vs E_1	137
9.2.2	Upstream Energies	142
9.3	Standard Trackers	142
9.3.1	Basic Tracker	142
9.3.2	Road Tracker	143
9.3.3	Upstream Tracker	144
9.4	P-Dir Tracker	144
9.4.1	Proton Direction Cut	144
9.4.2	Double Fusions	147
9.4.3	Helium Fusions	148
9.4.4	Electron Interference	149
9.4.5	Summary	150
9.5	Tracker Performance	152
9.5.1	Longitudinal Migration	153
9.5.2	Transverse Migration	157
9.6	Conclusion	159
Chapter 10: Systematics Overview		160
10.1	Stops on Other Nuclei	160
10.1.1	Wall Stops	161
10.1.2	Chemical Purity	161
10.1.3	Isotopic Purity	164
10.2	TPC Track Interference	165
10.3	Beam Backgrounds	166
10.3.1	Muon On Request	166
10.3.2	Muon Pileup	168
10.3.3	Background Slope	170
10.4	Electron Detectors	170
10.4.1	eSC Afterpulsing	170
10.4.2	ePC Noise	172

10.5 Conclusion	173
Chapter 11: Wall Stops	174
11.1 Neutron Detector Response	175
11.2 Neutron Sources	175
11.2.1 Muon Capture in Deuterium	176
11.2.2 Muon-Catalyzed Fusions	178
11.2.3 Wall Stops	179
11.2.4 Photo-Neutrons	180
11.2.5 Beam Background	181
11.3 Gamma Sources	182
11.3.1 Bremsstrahlung	182
11.3.2 Nuclear Capture	182
11.3.3 Muonic X-Rays	183
11.3.4 Beam Background	184
11.4 Analysis Procedure	184
11.4.1 Signal Modeling	187
11.4.2 Calibrations	190
11.4.3 Lifetime Effect	191
11.5 Results	193
11.5.1 Stop Y	193
11.5.2 Stop X and Z	194
11.6 Summary	195
Chapter 12: Time-Dependent Background	196
12.1 Electron Multiplicity	196
12.1.1 Lifetime Shift	197
12.2 Measurement	199
12.2.1 Slope Model	199
12.2.2 Muon Clock Fits	202
12.2.3 Clean Track Analysis	205
12.2.4 Simultaneous Fit	210
12.2.5 Slope Fit Comparison	212

12.3	Potential Causes	214
12.3.1	Beam Fluctuations	215
12.3.2	Detection Efficiency Changes	215
12.3.3	Kicker Instability	215
12.3.4	Long-Lived Daughter Nuclei	218
12.4	Simulation	220
12.4.1	Simulation Procedure	221
12.4.2	Linear Slope	222
12.4.3	Curvature	222
12.5	Conclusion	224
Chapter 13: Fusion Interference		225
13.1	Migration Modeling	225
13.1.1	Migration Probability	226
13.1.2	Stopping Distribution	228
13.1.3	Limitations	230
13.2	Fusion Fraction	231
13.2.1	Fusion Tagging	232
13.2.2	Lifetime Effect	233
13.2.3	Stop Distribution Shaping	233
13.3	Early Time Deviation	238
13.3.1	Background	242
13.3.2	Wall Stops	244
13.3.3	Lifetime Calibration	245
13.3.4	Road Tracker	249
13.4	Conclusion	251
Chapter 14: Electron Interference		253
14.1	Mechanism	254
14.2	Event Tagging	256
14.3	Time Distributions	257
14.4	Early Time Asymmetry	259
14.5	Downstream Migration Correction	261

14.6 Pileup Electron Asymmetry Correction	265
14.7 Pileup Electron Modeling	267
14.8 Conclusion	269
Chapter 15: Consistency Checks	271
15.1 Dataset Consistency	272
15.2 Fit Range Variations	273
15.2.1 Start Time	274
15.2.2 Stop Time	276
15.3 TPC Position	277
15.3.1 Stop X	277
15.3.2 Stop Y	278
15.3.3 Stop Z	279
15.4 Electron Direction	282
15.5 eSC Segment	282
15.6 eSC Z Position	283
15.7 Summary	284
Chapter 16: Conclusion	285
16.1 Systematics	285
16.2 Limitations	287
16.3 Outlook	287
16.4 Summary	289
Bibliography	290
Appendix A: Electron Tracking	300
A.1 Data Quality	300
A.1.1 Sparks	300
A.1.2 Afterpulsing	303
A.1.3 Bad Sector	305
A.2 Preprocessing	307
A.3 Clusters	308
A.3.1 Clustering	308

A.3.2	Fragmentation	309
A.4	Plane Coincidences	310
A.5	Tracks	314
A.5.1	Quality Cuts	315
A.5.2	High Multiplicity Events	319
A.6	Results	320
A.7	Conclusion	322
Appendix B: Muon On Request		324
B.1	Muon Background	324
B.1.1	Step Shape	324
B.1.2	Long-Range Correlations	325
B.1.3	Clock Events	326
B.1.4	Kicker Re-fire Offset	328
B.2	Secondary Particles	328
B.2.1	Step Shape	329
B.2.2	Transient Spike	330
B.3	TPC Tracks	331
B.4	Conclusion	333
Appendix C: Computer Systems		334
C.0.1	Data Storage	334
C.0.2	Job Submission	335
C.0.3	CENPA Cluster	337
Appendix D: eSC Time Alignment		338
D.1	Problem	338
D.2	Intra-Segment Alignment	339
D.3	Inter-Segment Alignment	340
Appendix E: Unique Pulser		342
Appendix F: Neutron PSD Cut Calibration		345
F.1	Cut Generation	345

F.2 Datasets	348
Appendix G: TPC Tracking Errors	350
G.1 Non-Fusion Migrations	350
G.2 P-Dir Cuts	353
Appendix H: Muon Beam Oscillations	355
Appendix I: Nuclear Isomers	359
I.1 Long-Lived Isomers	359
I.2 Locations of Stopped Muons	362
I.3 Testing	365
I.4 Conclusion	367
Appendix J: Fusion Tag	368
Appendix K: Background Dip	370
Appendix L: Downstream Electron Migration Model	372

ACKNOWLEDGMENTS

The work described in this thesis builds on previous contributions from many members of the MuSun collaboration. Since I joined near the end of the experiment I only worked directly with most of you during the last two data collection periods, but I would have liked the opportunity to participate in more runs and get to know all of you better. Thanks to Claude Petitjean, Fred Gray, Rob Carey, Tim Gorringer, and Wes Gohn for teaching me about the beam and detector hardware. To the PNPI team, including Alexei Vorobyov, Alexander Vasilyev, Peter Kravtsov, Marat Vznuzdaev, Nikolai Voropaev, Polina Kravchenko, Alexander Nadtochy, Volodya Ganzha, and Kuzma Ivshin, the experiment could never have succeeded without your expertise. And to all the other collaboration members, particularly the former students and postdocs who shared the hectic experience of learning about the experiment and troubleshooting problems during long night shifts at PSI and whose analyses laid the foundation for everything I have worked on.

I am grateful to the whole CENPA community, who have created a great learning environment for students. All of the faculty, staff, and other students contribute, but in particular I must thank David Hertzog for his insight and practical advice. I would also like to give special thanks to the MuSun group at UW, you were a really fun and creative group to work with. To Frederik Wauters and Michael Murray, thank you for showing me the ropes and helping me get up to speed early on. To Rachel Ryan, Daniel Salvat, and Duncan Prindle, I learned so much from our discussions and from struggling with difficult analysis problems together. And of course to my advisor Peter Kammel, thank you for all of your teaching and guidance throughout my time here.

Finally, I could never have completed this analysis without the unwavering support of my friends and family. Will, Dan, Casey, and Jeff, the fun we've had together over the years was invaluable for reducing my stress and keeping me optimistic. And Mom and Dad, thank you so much for instilling in me a passion for science, encouraging me to pursue an advanced physics degree, and for always being there whenever I needed help.

This work was supported in part by the National Science Foundation (grant number PHY-1206039) and the U.S. Department of Energy (grant number DE-FG02-97ER41020). It also made use of the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by NSF grant number ACI-1548562.

Chapter 1

INTRODUCTION

The visible matter in the universe is composed almost exclusively of only three fundamental particles: the up and down quarks and the electron. The up and down quarks are bound together into protons and neutrons by the strong force, and together with the electron these form all of the familiar atoms and molecules. However, at higher energy scales there exist two additional groups of matter particles, known as generations, which are identical to the originals in every way except for mass. The reason why there are three generations of matter particles is still a mystery, and may provide a clue to some deeper unified theory of particle physics. Because of their high mass, these second and third generation particles are unstable and rapidly decay into their first generation counterparts. Observation of the formation and decay of these particles has been invaluable for developing the modern standard model of particle physics.

The second generation counterpart to the electron, the muon, has long been of particular interest for experimental physics. Like the electron it has no color charge and is not directly affected by the strong force, so only electroweak interactions must be considered at first order allowing for clean calculations. The primary decay mode of the muon is to an electron and two neutrinos,

$$\mu^- \rightarrow e^- + \nu_\mu + \bar{\nu}_e. \tag{1.1}$$

with a lifetime of $2.1969811(22) \mu\text{s}$ [68]. Observation of muons produced in the upper atmosphere reaching ground detectors provided early evidence for Einstein's theory of special relativity, as this short lifetime would prevent such muons from reaching the ground without the effects of time dilation [60]. However, in the context of particle physics this lifetime is

actually quite long and provides ample opportunity for the manipulation and study of these particles. This makes the muon ideal for many experimental uses, as precision experiments may be performed while still allowing many muon decays to be observed in a short amount of time from a human perspective.

Another useful property of the muon is its mass, approximately 207 times larger than that of the electron. Heavier particles are valuable for studying the Higgs field and other hypothetical beyond the standard model processes whose interaction strengths depend strongly on mass. The third generation counterpart, the tau, has a mass 17 times higher than the muon, but with a lifetime 7 orders of magnitude shorter it is difficult to study. Heavy quarks and hadrons also introduce many complications due to their interactions with the strong force. In contrast the muon only interacts with the electroweak force at tree level, allowing for outstanding precision in both calculations and measurements as shown in table 1.1. Thus, the muon is used in several searches for new physics including searches for lepton flavor violation [10] or the Muon $g-2$ experiment measuring the anomalous magnetic moment [5].

Property	Symbol	Value	Precision
Mass	m_μ	105.6583755(23) MeV	22 ppb
Mean Lifetime	τ_μ	2.1969811(22) μ s	1 ppm
Magnetic Moment Anomaly	$a_\mu = \frac{g-2}{2}$	$116592061(41) \times 10^{-11}$	350 ppb

Branching Ratio	Value	LFV Mode	90% CL
$\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$	≈ 1	$\mu^- \rightarrow e^- \gamma$	$< 4.2 \times 10^{-13}$
$\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu \gamma$	$6.0(5) \times 10^{-8}$	$\mu^- \rightarrow e^- e^+ e^-$	$< 1.0 \times 10^{-12}$
$\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu e^+ e^-$	$3.4(4) \times 10^{-5}$	$\mu^- \rightarrow e^- 2\gamma$	$< 7.2 \times 10^{-11}$

Table 1.1: Selected properties, decay modes, and lepton flavor violating (LFV) limits of the muon, taken from the particle data group [69].

The muon mass is also of interest for nuclear physics. Thanks to their similarity to

electrons, muons may replace electrons in an atom in a process known as atomic capture. The resulting muonic atoms have rather unique properties, and in particular their orbital radius is reduced by a factor of roughly m_μ/m_e relative to their electronic counterparts. This produces a nearly 10^7 enhancement in the overlap between the muon and the nucleus, making muonic atoms a promising avenue for studying nuclear properties. For example, recent research using muon spectroscopy measurements have taken advantage of the enhancement of nuclear size effects to achieve a 0.1% measurement of the proton radius [22]. Interestingly, these muon spectroscopy results differ from the previous electron measurements by nearly 7 sigma, which has spurred the development of several experiments [58, 59].

The overlap of the muon with the nucleus also boosts the chance of a weak interaction between the muon and the nucleus. In particular we are interested in the process of nuclear muon capture, whereby the proton is converted to a neutron

$$\mu^- + p \rightarrow n + \nu_\mu. \quad (1.2)$$

Nuclear muon capture provides a probe of low energy nuclear physics and quantum chromodynamics (QCD), which are notoriously difficult to calculate from first principles alone. Effective field theory (EFT) approaches are drastically simplified in terms of particle content, at the cost of introducing a number of new possible interactions parameterized by so-called low energy constants (LECs). These constants are free parameters in the theory which must be constrained by external means, either by lattice calculations or directly by experiment. A program of muon experiments has therefore been developed to use muon capture measurements to constrain some of these low energy constants.

The MuCap experiment was the first in this program, and measured the muon capture rate on the proton to the 1% level. The MuCap result allows for the extraction of the pseudoscalar coupling constant, g_p , which agrees well with theoretical predictions [9]. MuSun is the next step in this program, substituting the hydrogen with a new deuterium target to study the two nucleon sector via the muon capture reaction

$$\mu^- + d \rightarrow n + n + \nu_\mu. \quad (1.3)$$

The theory of the two nucleon sector is currently well known except for a single LEC specifying the coupling of the axial two body current. Depending on the specific effective field theory construction, this constant is known as either L_{1A} or d_R . Extraction of this parameter by MuSun will enable a matching degree of precision in calculations of certain related reactions of astrophysical significance. These include the charged and neutral deuterium breakup reactions as well as solar proton-proton fusion from which the MuSun experiment derives its name. Furthermore, as this LEC is currently calculated from more complex three body processes, the MuSun measurement is cleaner and will provide an important consistency check for the theory. Recently there have been several important theoretical advancements, including a first lattice calculation of L_{1A} [63] and a new Bayesian analysis of muon capture on deuterium using chiral EFT[32], so it will be interesting to see how the MuSun result compares to these new predictions.

Experiment	Year	Target	T(K)	Detector	$\Lambda_d(s^{-1})$
Wang [67]	1965	liquid H_2+D_2	20	neutron	365 ± 96
Bertin [16]	1973	gaseous H_2+D_2	300	neutron	-
Bardin [13]	1986	liquid D_2	20	electron	470 ± 29
Cargnelli [21]	1989	gaseous D_2	45	neutron	409 ± 40

Table 1.2: Previously measured rates of muon capture on deuterium from the doublet hyperfine state. Interpretation of the 1973 measurement was complicated by unknown hyperfine populations.

Several previous experiments have attempted to determine the rate of muon capture on deuterium, but with inconclusive results as shown in table 1.2. All attempts have yielded large uncertainties, while the most precise measurement differs by nearly three standard deviations from more modern theoretical estimates which predict capture rates just under $400 s^{-1}$ [3, 48]. There are two main experimental effects complicating the precise determination

of the muon capture rate. First, the capture rate from the quartet hyperfine state is much lower than that of the doublet state due to suppression from the V-A (vector minus axial vector) structure of the weak interaction. We are specifically interested in the doublet capture rate Λ_d , and the hyperfine state populations must be well understood to correctly interpret the results. Second, the reduced size of muonic atoms also allows for the formation of muonic deuterium molecules with very closely spaced nuclei, which can spontaneously undergo fusion reactions. This muon-catalyzed fusion process produces a 2.45 MeV neutron background, complicating experiments attempting precision counting of capture neutrons.

Early experiments attempted to suppress muon-catalyzed fusion by using a primarily hydrogen target with a small concentration of deuterium added. However, this still allows $pd\mu$ fusions resulting in a buildup of μ^3He , so later experiments changed to pure deuterium targets. The 1973 value is omitted from the table because a temperature dependence of the $dd\mu$ formation rate was discovered in 1983 which called this result into question [37]. The 1986 measurement differs from the others by using the "lifetime method", where muon decay electrons are used to estimate the total muon disappearance rate rather than attempting to detect the capture neutrons directly. The contribution due to muon capture may then be estimated by comparing the measured disappearance rate to the free muon decay rate. Finally, the 1989 experiment used a low density target resulting in a high fraction of muon stops outside the gas and high backgrounds which dominated the signal.

The MuSun experiment aims to measure the capture rate to 1.5% precision by incorporating several advances over previous deuterium capture experiments. The target uses deuterium gas which has been purified to the ppb level, suppressing backgrounds due to capture on other atoms. An active target incorporating a time projection chamber (TPC) provides 3D tracking for each incoming muon, allowing selection of clean stops in the deuterium gas. The target conditions have also been carefully adjusted to minimize the rate of muon-catalyzed fusion and to allow clean interpretation of the hyperfine state populations. Finally, the lifetime method is used to avoid the challenges of directly measuring the absolute rate of capture neutron emission.

The full statistics necessary to achieve the precision target have been collected over two separate beam periods in 2014 (R2014) and 2015 (R2015). A followup beam period dedicated to systematics studies occurred in 2016 (R2016). The R2014 dataset has already been discussed in detail in Rachel Ryan's PhD thesis[62]. This thesis will focus on the data collection for the R2015 and R2016 datasets, as well as a variety of analysis improvements working towards a final result.

This thesis consists of three parts, the first of which focuses on the design of the experiment and collection of the data. Chapter 2 will discuss the theoretical motivation for the MuSun measurement, followed by an overview of the important design considerations for the experiment in chapter 3. Chapter 4 will then cover the final experimental setup in more detail, with chapter 5 focusing on the design of the TPC. Chapter 6 summarizes the operation of the experiment and the MuSun data collection campaign.

The second part of the thesis will cover the MuSun analysis software and procedures. The analysis framework consists of two main stages, which will be described in chapters 7 and 8. Chapter 9 will then discuss the software for muon tracking in the TPC, focusing on upgrades developed by the author.

The last portion focuses on resolving the major remaining sources of systematic error, starting with a summary in chapter 10. Chapters 11 through 14 contain detailed discussion of specific effects analyzed by the author. Chapter 15 will then present a variety of consistency checks comparing different datasets and analysis cuts. Chapter 16 concludes with a review of the current status and outlook of the experiment.

Chapter 2

THEORY AND MOTIVATION

Muon capture on atomic nuclei is fundamentally a weak interaction process, but an accurate description must also include corrections arising from strong interactions within the nucleus. It therefore serves as a valuable probe of the properties of strongly bound nuclear QCD states. This chapter will give an overview of the weak interaction and its well-understood application to the purely leptonic process of muon decay, before describing the properties of muon capture on nuclei of various masses. It then discusses the theoretical approaches relevant to muon capture on deuterium specifically in more detail, and introduces some of the recent theory developments. Finally, the implications of the MuSun measurement for these theories will be discussed.

2.1 Weak Interaction

A first theoretical description of the weak interaction was proposed by Enrico Fermi in the context of the neutron beta decay process

$$n \rightarrow p + e^- + \bar{\nu}_e. \quad (2.1)$$

His initial formulation described the weak interaction as a point-like four-fermion interaction, with the Lagrangian expressed as the product of two currents

$$\mathcal{L} = G_F J^\alpha L_\alpha. \quad (2.2)$$

In terms of the field creation and annihilation operators Ψ , the hadronic (J) and leptonic (L) currents are

$$J^\alpha = \bar{\Psi}_p \gamma^\alpha \Psi_n \quad L^\alpha = \bar{\Psi}_e \gamma^\alpha \Psi_{\nu_e}. \quad (2.3)$$

These have the effect of converting a neutron into a proton and a neutrino into an electron, with the Dirac matrices γ causing the currents to transform as vectors under Lorentz boosts.

Although this description was somewhat successful, it did not explain the observation of Gamow-Teller transitions in which the spin of the nucleus changes by 1. In 1956 Lee and Yang proposed the possibility of parity violation, implying that the currents in equation 2.3 may in fact include several other Lorentz invariant combinations of gamma matrices:

$$\begin{aligned}
 \text{Scalar (S)} & \quad \mathbb{1} \\
 \text{Pseudoscalar (P)} & \quad \gamma^5 \\
 \text{Vector (V)} & \quad \gamma^\alpha \\
 \text{Axial Vector (A)} & \quad \gamma^\alpha \gamma^5 \\
 \text{Tensor (T)} & \quad \sigma^{\alpha\beta} \equiv \gamma^\alpha \gamma^\beta - \gamma^\beta \gamma^\alpha.
 \end{aligned} \tag{2.4}$$

Because the gamma matrices satisfy the Clifford algebra anti-commutation relation, adding any additional matrices simplifies to one of the above combinations and no further terms are possible. A series of experiments followed which identified the correct V-A structure of the weak interaction[25]. This is equivalent to adding the left-handed projection operator

$$P_L = \frac{1 - \gamma^5}{2} \tag{2.5}$$

to the currents, yielding the surprising fact that the weak interaction only couples to particles with left-handed chirality.

The corrected Lagrangian then becomes

$$\mathcal{L} = \frac{4}{\sqrt{2}} G_F J^\alpha L_\alpha \quad L^\alpha = \bar{\Psi}_e \gamma^\alpha P_L \Psi_{\nu_e}. \tag{2.6}$$

The numerical prefactor has been added to retain the original value of G_F , with the extra factor of $\sqrt{2}$ arising because Fermi assumed both spin states were possible for the neutrino and thus double-counted the contribution to the decay rate. The hadronic current is similar to the leptonic current when expressed in terms of individual free quarks, but becomes complicated for strongly bound states as we will see in the next section.

In 1961 Glasgow proposed a gauge theory that unified the electromagnetic and weak forces. In this theory the weak interaction is no longer a four-fermion contact interaction, instead being mediated by the massive W vector boson as shown in figure 2.1 below. The W propagator is given by

$$-\frac{g_{\alpha\beta} - q_\alpha q_\beta / M_W^2}{q^2 - M_W^2}, \quad (2.7)$$

where $g_{\alpha\beta}$ is the metric tensor and q and M_W are the momentum and mass of the W boson, respectively. At low momenta $q \ll M_W$, the propagator reduces to $1/M_W^2$. The vertex factor is conventionally written as $g/\sqrt{2}$, resulting in the relationship

$$\frac{4}{\sqrt{2}}G_F = \frac{g^2}{2M_W^2}. \quad (2.8)$$

This theory was not gauge invariant due to the presence of the massive vector boson, a problem that was remedied by Weinberg and Salam several years later via the mechanism of spontaneous symmetry breaking. The symmetry of the original massless vector bosons is broken by the ground state of the system, which selects a preferred basis rotated by the weak mixing angle θ_W . The coupling constants e and g of the resulting electromagnetic and weak interactions are then related by

$$e^2 = g^2 \sin^2 \theta_W \approx 0.22290(30)g^2. \quad (2.9)$$

Thus the weak interaction is in principle stronger than the electromagnetic interaction, but is heavily suppressed at low momentum due to the W mass of $80,433.5(9.4) \text{ MeV}/c^2$ [24]. This standard model picture of the electroweak interaction is then completed with the Higgs mechanism, which results in massive gauge bosons and an additional massive scalar Higgs boson. The Higgs boson was finally observed at the large hadron collider (LHC) in 2012.

2.1.1 Muon Decay

Muon decay, illustrated in figure 2.1, is ideal for characterizing the strength of the weak interaction. Because this is a purely leptonic interaction it does not suffer from the complications introduced by strongly interacting hadronic states, and since the W mass is




Figure 2.1: Tree level Feynman diagrams of the weak interaction in either the Fermi theory (left) or the standard model (right).

over 750 times larger than the muon mass of 106 MeV the contact interaction formulation is extremely accurate. The Lagrangian contains two copies of the leptonic current:

$$\mathcal{L} = \frac{4}{\sqrt{2}} G_F L^\alpha L_\alpha^\dagger \quad (2.10)$$

$$L^\alpha = \bar{\Psi}_e \gamma^\alpha P_L \Psi_{\nu_e} + \bar{\Psi}_\mu \gamma^\alpha P_L \Psi_{\nu_\mu},$$

where a second term has been added to account for the muon.

Determination of the muon decay rate requires calculation of the invariant amplitude,

$$M = \frac{G_F}{\sqrt{2}} \bar{u}_{\nu_\mu} \gamma_\alpha (1 - \gamma^5) u_\mu \bar{u}_e \gamma^\alpha (1 - \gamma^5) u_{\nu_e}, \quad (2.11)$$

written in terms of the Dirac spinors u_i arising from the field operators acting on the initial and final states. Fermi's golden rule states that the transition probability is proportional to the coupling and the density of states, yielding a differential transition probability

$$d\Gamma = |M|^2 (2\pi)^2 \delta^4(p_e + p_{\nu_e} + p_{\nu_\mu} - p_\mu) \left(\prod_f \frac{m_f}{E_f} \right) \left(\prod_f \frac{d^3 p_f}{(2\pi)^3} \right), \quad (2.12)$$

where the sums run over the four fermions $f = \mu, e, \nu_\mu,$ and ν_e . Plugging in equation 2.11 and integrating over the outgoing particle momenta yields an expression for the muon lifetime τ_μ :

$$\tau_\mu = \frac{192\pi^3}{G_F^2 m_\mu^5}. \quad (2.13)$$

This calculation only considered the fundamental particle interactions (tree level), and higher-order terms must be added to account for Feynman diagrams containing internal loops. However, the benefit of the purely leptonic system is that the only higher-order terms are due to radiative corrections, many of which were calculated by van Ritbergen and Stuart in 1999 to calculate the relation between τ_μ and G_F with sub-ppm precision[66]. In 2013 the MuLan experiment measured the positive muon lifetime to the ppm level as well, finding $\tau_\mu = 2,196,980.3(2,2)$ ps[68]. Together these results constrain the Fermi coupling constant G_F to 0.5 ppm. In addition to being an important test of the standard model in its own right, this result is very useful for the MuSun experiment where it will be compared against the negative muon disappearance rate in matter to ascertain the additional capture rate contribution.

2.2 Muon Capture

The characteristic radius of a hydrogen atom is approximately equal to the Bohr radius

$$a_0 = \frac{\hbar}{m_e c \alpha}. \quad (2.14)$$

In a more general hydrogen-like atom the mass of the electron m_e must be replaced by the reduced mass, while the radius is reduced by an additional factor of the nuclear charge Z . The resulting radius becomes

$$a_{Z,m} = \frac{m_e a_0}{m Z}. \quad (2.15)$$

A muon thus has an orbital radius approximately 200 times smaller than an electron for a given atom. This small radius also means the approximation of a hydrogen-like atom is valid even if there are other electrons present, because they typically be well outside the orbit of the muon and have little effect.

The muon capture reactions discussed in this thesis depend on the overlap between the muon wavefunction and the nucleus. The radial component of a hydrogenic 1S wavefunction takes the form

$$\psi(r) = 2 \left(\frac{1}{a_{Z,m}} \right)^{3/2} e^{-(r/a_{Z,m})}, \quad (2.16)$$

where r is the radial distance. Approximating the nucleus as a point particle, the probability of the finding the muon at the nucleus is thus proportional to the cube of the reduced mass and atomic number:

$$|\psi(0)|^2 \propto \left(\frac{1}{a_{Z,m}}\right)^3 \propto (mZ)^3. \quad (2.17)$$

The overlap between the muon and the nucleus is thus enhanced by a factor of nearly 10^7 compared to that of the electron, making it an attractive choice for this and other measurements that depend on this overlap. For capture specifically the muon may also interact with any of the individual protons in the nucleus, producing another factor of Z . For light nuclei the muon capture rate therefore scales roughly as

$$\Lambda \propto m^3 Z^4. \quad (2.18)$$

The capture reaction itself is essentially identical to beta plus decay, except with a muon being adsorbed rather than a positron being emitted. Expressed in terms of individual free quarks the reaction is simply

$$u + \mu \rightarrow d + \nu_\mu. \quad (2.19)$$

The interaction is characterized by the same Lagrangian as in equation 2.6 above, but we have not yet discussed the hadronic current. Written in terms of the field operators for individual free quarks, it is similar to the leptonic current apart from the addition of the CKM matrix element V_{ud} representing the quark flavor mixing:

$$J^\alpha = \bar{\Psi}_d V_{ud} \gamma^\alpha P_L \Psi_{\nu_\mu}. \quad (2.20)$$

To determine the transition probability it is necessary to calculate the invariant amplitude, but while the leptonic contribution is the same as in muon decay

$$\langle \bar{\nu}_\mu | L_\alpha | \mu \rangle = \bar{u}_{\nu_\mu} \gamma_\alpha P_L u_\mu \quad (2.21)$$

the result of applying the quark field operators to a bound state of quarks is less obvious. The details of the hadronic contribution and the theoretical tools used to analyze it vary

depending on the size of the nucleus involved. Particularly for light nuclei, the muon capture process provides a unique environment where observables related to strongly bound QCD states may be both measured and predicted with high precision.

2.2.1 Single Nucleon




Figure 2.2: Tree level Feynman diagram of muon capture on the proton.

Muon capture on the proton, shown in figure 2.2, is an important probe for testing the theory of the weak interaction in the single nucleon sector. With a single nucleon involved we can still approximate this as a point-like interaction, but the hadronic contribution will differ from the standard V-A form and is difficult to calculate from first principles. Instead the matrix elements may be written in a general form including all possible Lorentz terms, often written with vector terms V^α and axial terms A^α grouped together:

$$\langle n | J^\alpha | p \rangle = \bar{u}_n \frac{1}{2} [V^\alpha - A^\alpha] P_L u_p \quad (2.22)$$

$$V^\alpha = F_V \gamma^\alpha + i \frac{F_M}{2m_N} \sigma^{\alpha\beta} q_\beta + \frac{F_S}{m_\mu} q^\alpha \quad (2.23)$$

$$A^\alpha = \left[F_A \gamma^\alpha + i \frac{F_T}{2m_N} \sigma^{\alpha\beta} q_\beta + \frac{F_P}{m_\mu} q^\alpha \right] \gamma^5. \quad (2.24)$$

The various F_i parameters are known as the weak nuclear form factors, and encapsulate the effects of the strong interactions between quark and gluon constituents within the nucleon[34]. These form factors have a small dependence on the energy scale, $F \rightarrow F(q^2)$.

The scalar (F_S) and pseudo-tensor (F_T) terms are so-called "second-class" currents and are expected to be nearly zero. This is because of their transformation under G-parity, a generalization of charge conjugation symmetry involving an additional rotation in isospin space. The vector terms V^α generally have positive G-parity and the axial terms A^α have negative G-parity, but the second-class currents have opposite G-parity to the rest of their groups and are thus heavily suppressed.

The conserved vector current (CVC) hypothesis states that the vector current $\bar{u}_n V^\alpha u_p$ forms an isospin triplet along with its complex conjugate and the isovector electromagnetic current j_{EM}^α [25]. The conservation of the electromagnetic current then implies the conservation of the weak vector current, and the vector (F_V) and magnetic (F_M) form factors related to the well studied electromagnetic form factors. This yields the expressions

$$F_V(q^2 = 0) = 1 \tag{2.25}$$

$$F_M(q^2 = 0) = \frac{1}{2}(\mu_p - \mu_n - 1), \tag{2.26}$$

where $\mu_p = 2.79$ and $\mu_n = -1.91$ are the magnetic moments of the proton and neutron in units of the nuclear Bohr magneton[52]. The CVC hypothesis also implies that the scalar current should be zero independent of the previous argument about it being a second-class current.

Finally, the axial (F_A) and pseudoscalar (F_P) form factors are unique to the weak interaction, and must be determined from experimental measurements. The axial form factor has been determined to have a value of $F_A = -1.2723 \pm 0.0023$ via measurements of the asymmetry in the neutron beta decay rate[57, 51, 53]. The factor of momentum in the pseudoscalar current means that it has a small effect on these low-energy beta decay experiments.

This left the pseudoscalar form factor as the only unknown in equation 2.24. Initial estimates of F_P were calculated using the partially conserved axial current (PCAC) hypothesis, which states that in the limit of massless quarks the axial current would be approximately conserved in QCD. With massive quarks the axial current is no longer conserved, but its

divergence is related to the pion field

$$\partial_\alpha A^\alpha(x) \propto \phi(x). \quad (2.27)$$

Expanding the divergence around zero momentum transfer yields the first order term

$$F_P(q^2) = \frac{2m_\mu m_N}{m_\pi^2 - q^2} F_A(0). \quad (2.28)$$

Chiral perturbation theory, discussed in a later section, provides a way to systematically improve[15] this estimate, giving a precise QCD prediction

$$F_P^{chiral}(q^2) = \frac{2m_\mu g_{\pi NN} F_\pi}{m_\pi^2 - q^2} - \frac{1}{3} F_A(0) m_\mu m_N r_A^2. \quad (2.29)$$

The PCAC hypothesis has been very successful, but a precision measurement of F_P was required to confirm the validity of this assumption. Muon capture on the proton was uniquely placed to measure F_P , as in this case the momentum transfer is nearly the entire muon mass energy. In 2013 the MuCap collaboration achieved a 1% measurement of the capture rate from the singlet state of the muonic hydrogen atom $\Lambda_S = (715.6 \pm 5.4^{\text{stat}} \pm 5.1^{\text{sys}}) s^{-1}$. Even when the uncertainty in the axial radius of the nucleon is taken into account, the MuCap result confirms the theoretical prediction of F_P to 8%[36].

2.2.2 Light Nuclei

With multiple nucleons in the nucleus several new processes become possible in addition to the one-body weak nuclear current from equation 2.24. For capture on the deuteron there are also two-body currents involving both nucleons, and for tritium or helium-3 three-body currents must also be considered. Historically a phenomenological approach called the standard nuclear physics approach (SNPA) was successful in calculating observables for few-nucleon systems, including the muon capture rate on the deuteron. Waveforms are derived from approximate potentials tuned to match existing nuclear physics measurements and scattering data. Transition operators are then given by a combination of the one-body

current with two- and three-body terms derived from meson exchange currents (MECs), including the possibility of exchanging multiple pions.

There were several deficiencies to the SNPA approach. It was explicitly model-dependent, with several possible potentials to choose from and relatively little constraint on short-range behavior. It is also difficult to systematically identify the most significant terms, with no clear expansion parameter or power-counting hierarchy. Finally, the calculations have little relation to the fundamental QCD theory, which should be the underlying description of the full nuclear dynamics.

Modern calculations of nuclear properties typically use a more systematic effective field theory (EFT) approach, in which the short-range and high-energy behavior has been essentially integrated out. In this approach one first creates an effective Lagrangian including the relevant degrees of freedom at the desired energy scale, and including all interaction terms permitted by the symmetries of the underlying theory. This typically results in a large number of complicated many-body interaction terms, but these may be organized into powers of the momentum divided by the cutoff energy. If the EFT is based on a large energy gap present in the full theory then this expansion parameter can be quite small, and only a few low-order terms in the expansion need be considered. Each of these terms represents the integrated effects of the high-energy dynamics, which is parameterized by a free parameter known as a low energy constant (LEC) that must be determined empirically.

Actually, most of the discussion so far has been about EFTs, with the Fermi theory of the weak interaction being perhaps the most well-known example of such a theory. In that case the energy scale involved is so much smaller than the high energy degrees of freedom that the description is extremely accurate even at first order. For studying light nuclei we typically use theories with high energy cutoffs at either the pion mass or an energy scale of approximately 1 GeV related to the approximate chiral symmetry of QCD, these will be discussed in more detail in the next section. In these cases the cutoff is not particularly large relative to the energies involved in muon capture, and the theory based on the pion mass in particular is only usable if the majority of the muon mass energy is carried away by the

neutrino. However, they still provide the benefit of a model-independent and systematic way to organize the calculation and estimate the uncertainties arising from higher-order terms.

Finally, for light nuclei it is also possible to calculate some results from first principles using lattice QCD. These calculations yield interesting results in their own right, but may also be used as input data to fix various LECs that arise when using the EFT approach. In particular, a first lattice calculation of the matrix element determining the $pp \rightarrow de^+\nu$ fusion cross section and the Gamow-Teller matrix element contributing to tritium beta decay were calculated for the first time in 2017[63]. These calculations and the MuSun measurement of muon capture on deuterium both constrain the same LEC, and should provide independent confirmation of the result.

2.2.3 Heavy Nuclei

As the number of nuclei in the system increases the theory must deal with more complicated many-body interactions and long-range correlations. As such it rapidly becomes difficult to apply the EFT and lattice QCD approaches discussed above. For heavy nuclei other theoretical tools are commonly used, such as nuclear shell models, mean-field theories, or other EFTs designed specifically for large nuclei. These theories are beyond the scope of this thesis and will not be discussed in detail here, but it will be useful to cover some general properties of muon capture on heavy nuclei as it will be an important source of potential background for the MuSun experiment.

The point nucleus approximation predicts the capture rate reasonably well below about $Z = 10$, but for heavier nuclei the finite size of the nucleus must be accounted for. The factor of muon mass in the orbital radius means that it is already only 284 fm for muonic hydrogen, and for higher Z the orbital radius quickly becomes comparable to the nuclear radius of a few fm. To correct for the finite size Z may be replaced with Z_{eff} , which may be less than half the true Z for heavy elements above $Z = 60$. Another consideration is that heavy nuclei typically have more neutrons than protons, so the Pauli exclusion principle makes it more difficult for protons to convert to neutrons. Primakoff proposed a more accurate formula for

the capture rate

$$\Lambda(A, Z) = Z_{eff}^4 X_1 \left[1 - X_2 \left(\frac{A - Z}{2A} \right) \right], \quad (2.30)$$

with $X_1 = 170s^{-1}$ and $X_2 = 3.125$ being determined from a fit to the data[50].

Muon capture on a given nucleus will convert one proton into a neutron, reducing the atomic number by one



The neutrino typically takes away the majority of the muon mass energy, but still leaves the nucleus in a highly excited state with around 5-50 MeV excitation energy. The nucleus usually de-excites by emitting up to about 5 neutrons, and also has a smaller chance to eject protons or even alpha particles[50]. Lighter nuclei generally have a higher chance of emitting charged particles, while heavier nuclei instead tend to eject larger numbers of neutrons. The isotopes produced by muon capture reactions have been studied for a variety of target elements[35]. The de-excitation may also occur via emission of gamma rays alone, without changing the mass of the nucleus.

2.3 Effective Field Theory

2.3.1 QCD Symmetries

Two important symmetries must be considered when discussing QCD: isospin and chiral symmetries. Consider the simplified case of only one generation of quarks, with the two flavors comprising a doublet field

$$\Psi = \begin{pmatrix} u \\ d \end{pmatrix}. \quad (2.32)$$

An isospin rotation is described by a unitary 2x2 matrix Ω

$$\Psi \rightarrow \Psi' = \Omega \Psi. \quad (2.33)$$

This results in three vector currents, with the divergence of the charged current given by

$$\partial_\alpha V^\alpha = \partial_\alpha (\bar{u} \gamma^\alpha d) = i(m_u - m_d) \bar{u} d. \quad (2.34)$$

The current would be conserved in the limit where the quark masses m_u and m_d are equal, and this is nearly the case in QCD. This observation is what motivated the previously discussed CVC hypothesis used to predict the nucleon vector and magnetic form factors.

Such a rotation may also be applied to individual chiral components of the field,

$$\begin{aligned}\Psi_L &\rightarrow \Psi'_L = \Omega_L \Psi_L \\ \Psi_R &\rightarrow \Psi'_R = \Omega_R \Psi_R.\end{aligned}\tag{2.35}$$

This similarly generates three currents, with the divergence of the charged axial current being given by

$$\partial_\alpha A^\alpha = \partial_\alpha (\bar{u} \gamma^\alpha \gamma_5 d) = i(m_u + m_d) \bar{u} \gamma_5 d.\tag{2.36}$$

This current is conserved if both quark masses are zero, giving rise to the PCAC hypothesis.

Previously the PCAC hypothesis was discussed in terms of relating the divergence to the pion field. This is because the pions may be identified as the Goldstone bosons arising from the chiral symmetry being spontaneously broken by QCD. As described by Nambu[55] and Goldstone[33] in 1960, this spontaneous symmetry breaking should produce massless bosons. However, in the case of the pions the chiral symmetry is also broken explicitly, albeit by a small amount, and so the pion mass becomes related to the sum of the quark masses

$$m_\pi^2 = (m_u + m_d) \frac{\langle \bar{q}q \rangle}{f_\pi^2},\tag{2.37}$$

where $\langle \bar{q}q \rangle$ is the so-called condensate parameter and f_π is the pion decay constant.

2.3.2 Chiral Perturbation Theory

Exploiting the approximate chiral symmetry of QCD, the Hamiltonian may be written

$$H = H_0 + H_{sb},\tag{2.38}$$

where H_0 includes the terms that respect chiral symmetry while H_{sb} contains the symmetry breaking quark mass terms. The chiral symmetry is spontaneously broken at an energy scale of $\Lambda_\chi \approx 1$ GeV[38].

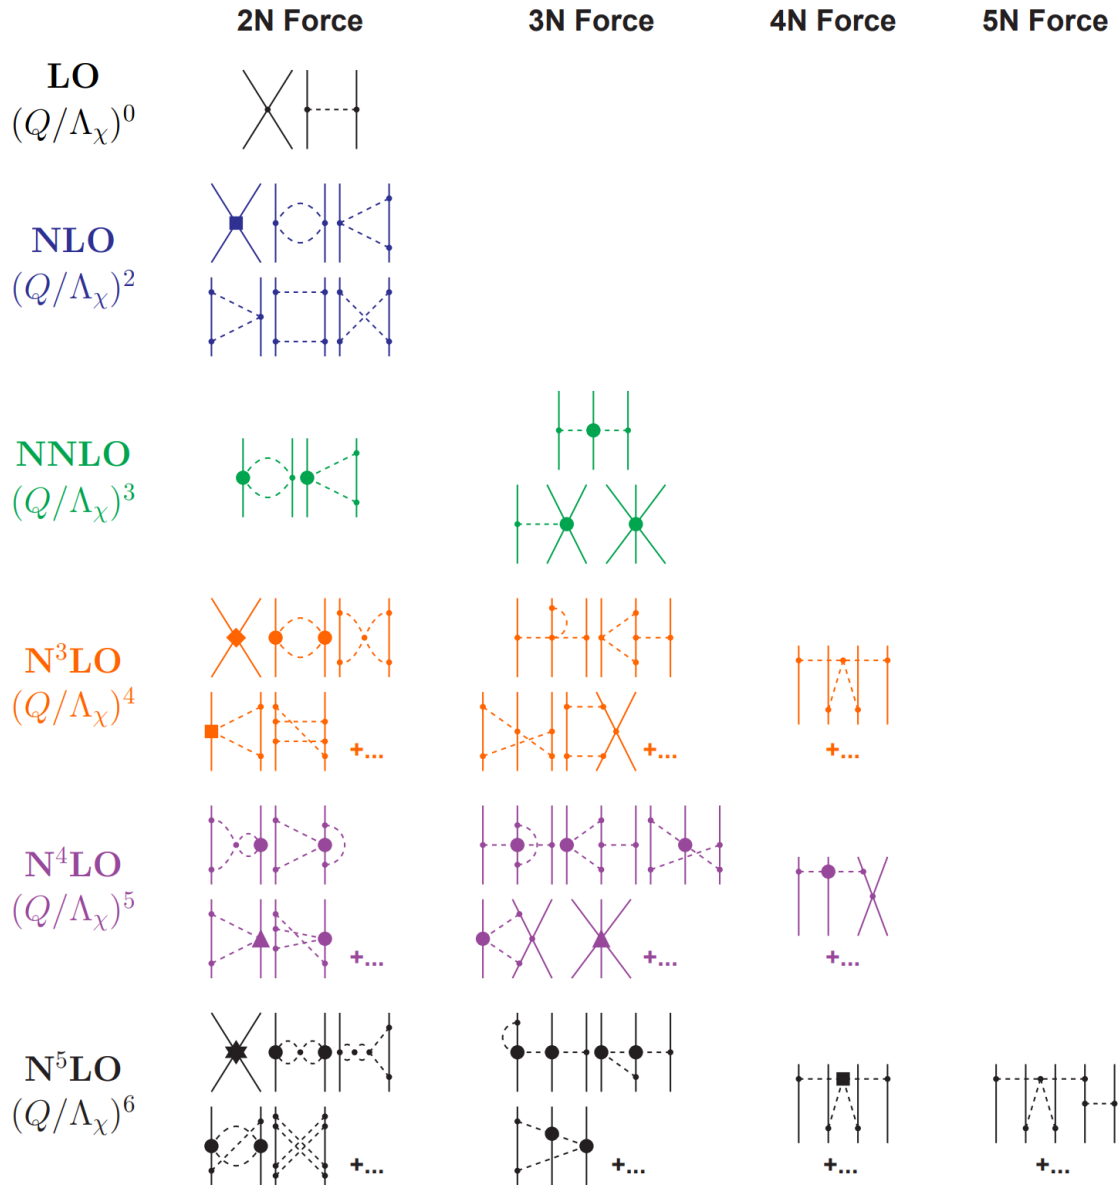


Figure 2.3: Hierarchy of nuclear forces in chiral perturbation theory (χ PT). Solid lines represent nucleons and dashed lines pions. Small dots, large solid dots, solid squares, triangles, diamonds, and stars denote vertices of index $\delta = 0, 1, 2, 3, 4,$ and $6,$ respectively. Reproduced from [45].

Chiral perturbation theory (χ PT) uses this chiral symmetry breaking scale as its high-energy cutoff, integrating out all physics above this energy. The nucleon mass is not small compared to Λ_χ , so χ PT would seem to be useful mainly for processes involving only light mesons. To include the nucleons without introducing terms proportional to the nucleon mass, the baryons are treated non-relativistically with small momentum transfer to and from the pions. This modification is known as heavy baryon chiral perturbation theory (HB χ PT), and has become one of the primary frameworks for analyzing light nuclei[44]. The name chiral effective field theory (χ EFT) may also be used to refer to the more complicated multi-nucleon case, with χ PT then being reserved for single-nucleon or pion-only cases.




Figure 2.4: Feynman diagrams that contribute to the two-nucleon axial vector coupling (top) and to the 3N forces at N²LO (bottom). Solid lines represent nucleons, dashed lines pions, and the wavy line is the W boson.

Because the expansion parameter p/Λ_χ is not particularly small, it is necessary to include the first several orders of the expansion to achieve precise results. These orders are named leading order (LO), next-to leading order (NLO), and then next-to-next-to leading order and

so on, denoted $N^k\text{LO}$. In χPT there is a minor complication that the $(p/\Lambda_\chi)^1$ terms vanish, so there is disagreement as to whether the $(p/\Lambda_\chi)^2$ should be referred to as NLO or $N^2\text{LO}$. This thesis will refer to it as NLO, following the convention used in figure 2.3 where Feynman diagrams for interactions up to $N^5\text{LO}$ are shown. As the order increases interactions appear involving increasing numbers of nucleons, but there are also new two- and three-nucleon interactions associated with new LECs.

The focus of this thesis is primarily on the coupling of the weak current to the two-nucleon system, which is described by the LO diagrams shown in figure 2.4. The contact interaction is characterized by a single LEC, d_R , which is what the MuSun experiment aims to determine. When calculating up to $N^2\text{LO}$ the two-nucleon couplings may be related to the forces in the three-nucleon system, also shown in figure 2.4. The LECs c_1 , c_3 , and c_4 enter at NLO, while c_D and c_E are $N^2\text{LO}$ terms. These can be used to determine the value of d_R via the relationship

$$d_R = -\frac{1}{4} \frac{m_N}{g_a \Lambda_\chi} c_D + \frac{m_N}{3} (c_3 + 2c_4) + \frac{1}{6}. \quad (2.39)$$

The four LECs that characterize πNN vertices at NLO, labeled c_i , are determined from measurements of elastic pion-nucleon scattering[20, 29]. The parameters c_D and c_E are most precisely measured with a combination of either ^3He or triton binding energies and the tritium beta decay rate Γ_β^t . Γ_β^t is strongly affected by d_R but is relatively insensitive to the individual contributions from either c_D or c_E , so the binding energy is used to fix the relation between c_D and c_E and Γ_β^t is calculated in terms of this relation and compared to experiment. Because tritium beta decay has been measured very precisely this procedure currently provides the best estimate of d_R , but the MuSun result will offer a more direct determination of d_R as well as a check on the overall consistency of the current theory.

2.3.3 Pionless EFT

The $\text{HB}\chi\text{PT}$ formulation is complicated by the existence of additional length scales such as the large S-wave NN scattering length. Kaplan, Savage, and Wise (KSW) argued that




Figure 2.5: Leading order diagram contributing to NN axial vector coupling in pionless EFT.

this is a sign that the short-range physics is "finely-tuned" and non-perturbative at low energy[39]. In particular there may be significant cancellation between terms at nominally different orders in the expansion, indicating that it may not be safe to neglect higher-order terms. They instead proposed a new pionless effective field theory (π EFT), which treated the pion as a high-energy degree of freedom to be integrated out[40]. This leaves only nucleons as the relevant constituents of the theory, and with no complicated dynamics arising from the exchange of pions it produces a well-defined explicit power counting scheme in terms of the momentum divided by the pion mass, p/m_π . The resulting theory is very concise, with all interactions being contact interactions between the nucleons. At first order the axial vector coupling in the two-nucleon sector gives rise to the diagram in figure 2.5, parameterized by a single LEC, L_{1A} , similar to the d_R parameter from HB χ PT.

Although π EFT has benefits for low-energy processes, its relevance to muon capture reactions is questionable due to the final state energy including most of the muon mass energy $m_\mu \approx 0.75m_\pi$. This energy is shared among the nucleons and crucially the neutrino, which does not play a role in the EFT calculations and may carry away a significant fraction of the energy. Figure 2.6 shows the possible final state energies of the two neutrons and the neutrino produced in muon capture on the deuteron, with π EFT only applying to a small region of the plot where over 90% of the energy is lost to the neutrino. However, the high density of states in this region of phase space means that it contributes to roughly 90% of the total capture rate.

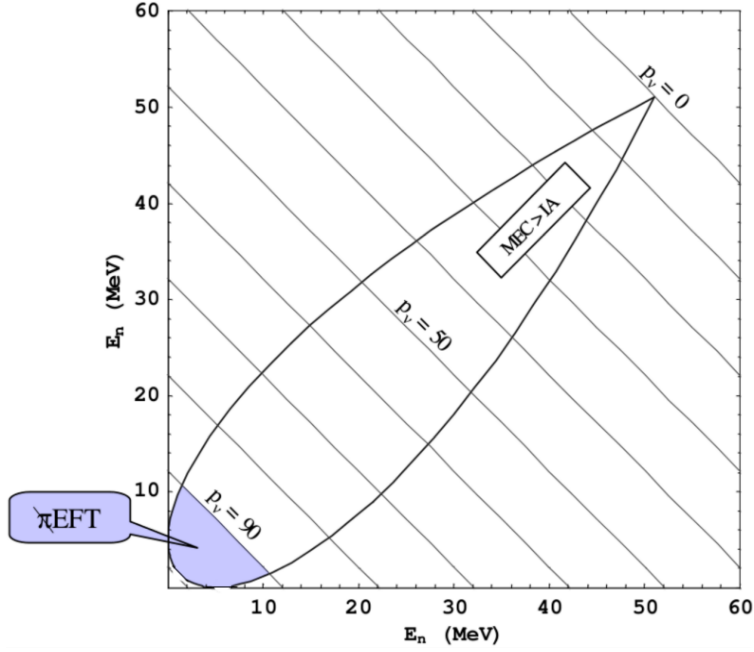


Figure 2.6: Final state energy of the two neutrons and neutrino produced in muon capture on the deuteron. The energy of each neutron (E_n) is plotted on the X and Y axes, with diagonal lines indicating constant neutrino momentum (p_ν) in units of MeV/c. π EFT is only applicable in the bottom left region where $p_\nu > 90$ MeV/c.

To complete the theoretical calculation one could match the π EFT result to the HB χ PT result at low energies and then extend the partial capture rate $d\Lambda_d(E)/dE$ to the full range of energies, which would suppress any uncertainty from the HB χ PT theory. The μ d capture rate has been calculated up to N²LO in π EFT[23], with some more recent work making an effort to consistently study three-nucleon processes such as tritium beta decay in the π EFT framework[27]. Alternatively an experiment could aim to measure the partial capture rate for only low neutron energies where π EFT is applicable, which would not represent a large loss in statistics but would require accurate determination of the neutron energy and prevent the use of the lifetime method.

2.4 History of Λ_d Calculations

As mentioned above, historically the phenomenological SNPA approach was used to calculate observables for few-nucleon systems, including muon capture on deuterium[46]. Beginning in the 2000s there was a gradual transition to the more systematic χ EFT approach, starting with an intermediate theory known as hybrid chiral effective field theory (χ EFT*). Accurately deriving the nuclear wavefunctions in the χ EFT approach was quite challenging, and for some time the older empirical wavefunctions were more precise. In χ EFT* one therefore derives transition operators based on χ EFT, but then applies those operators to the phenomenological wavefunctions for the initial and final states[49, 8]. Interestingly, because the c_D and c_E LECs discussed previously are mainly responsible for defining the nuclear potentials, in the hybrid approach they are unused and d_R is the only unknown LEC describing both muon capture and tritium beta decay.

The first fully consistent χ EFT calculations were performed in 2012 By Marcucci and Adam[47, 3]. It is worth noting that these early results suffered from a widely-spread error in the calculation of equation 2.39 in which the factor of -1/4 was omitted, and the original results must be modified accordingly[48]. Another issue was that the uncertainty in F_A had been underestimated. The momentum dependence of F_A is governed by the axial radius r_A in the expansion at small q^2 ,

$$F_A(q^2) = F_A(0) \left(1 + \frac{1}{6} r_A^2 q^2 + \dots \right). \quad (2.40)$$

This had previously been fit with a one-parameter dipole model, in which case the higher-order terms denoted by the ellipsis are completely determined. The true functional form of $F_A(q^2)$ is unknown, and a more general analysis leads to much larger uncertainties[36]. In 2018 an improved analysis was produced addressing both of these issues, but focusing solely on the 1S_0 channel[1].

Finally, in 2023 two new state of the art calculations were released. Bonilla et al. performed a complete partial wave expansion to N²LO, as well as an analysis of the theoretical uncertainties, and produced estimates for the capture rate from both the doublet and quar-




Figure 2.7: Theoretical predictions of the capture rate on deuterium are shown in blue[3, 47, 18, 32], with experimental results shown in red for comparison[67, 16, 13, 21]. Note that the 2012 predictions do not include the uncertainty related to r_A , while interpretation of the 1973 measurement was complicated by unknown hyperfine populations. MuSun will bring the experimental uncertainty in line with current theory predictions, although the results are still blinded so the vertical position of this point is arbitrary.

tet hyperfine states[18]. Gnech et al. then performed a similar partial wave expansion, with the aim of producing a robust theoretical error estimate based on a Bayesian analysis of the truncation errors, the model dependence, and the propagation of LEC uncertainties[32]. Figure 2.7 shows a comparison of some of the most recent theory results.

2.5 Significance of MuSun

The MuSun experiment will be the first precision measurement of the rate of a weak process in the two-nucleon sector, and will provide a clean 2N measurement of the χ PT LEC d_R . This parameter has already been estimated based on 3N measurements such as tritium beta decay, but there the determination is less clean due to additional LECs that play a role in

the 3N system. The MuSun experiment also remains one of the best options for producing an accurate estimate of the π EFT LEC L_{1A} .

In addition to d_R , the MuSun measurement may be used to estimate both c_D and the axial radius r_A . In the case of c_D this will be much less accurate than the existing value, but for r_A the MuSun uncertainty will be comparable to the existing uncertainty. The MuSun experiment will not significantly improve the current value, but it will still be a fundamental test of the χ EFT system and a significant disagreement from the current theoretical predictions could require a serious revision of the chiral electroweak currents.




Figure 2.8: Relevant astrophysical reactions also characterized by d_R .

Another more concrete application of the MuSun measurement is to some other 2N interactions with a similar dependence on d_R . Figure 2.8 shows several of these that are of interest to astrophysics, including the solar proton-proton (p-p) fusion reaction and the charged current (CC) and neutral current (NC) neutrino scattering reactions measured by the SNO collaboration. All three reactions are difficult to measure experimentally at terrestrial conditions, with the neutrino cross sections being extremely small and the fusion reaction requiring extreme temperatures or pressures to overcome the Coulomb barrier. In contrast the muon capture reaction is relatively easy to observe with high statistics, although achieving commensurate systematic uncertainty is quite difficult as will be shown in the remainder

of this thesis.

The p-p fusion reaction is the first step in the process that fuels all main sequence stars, and as such it is critical for characterizing the evolution of stars and to determine their ages. The rate of p-p fusion is normally characterized by the S-Factor, which is given by rescaling the total cross section to account for the Coulomb repulsion between the protons. Recent calculations of the S-Factor using χ PT obtain values about 3% higher than expected, corresponding to less than 1% change in age determination and 8% to neutrino flux predictions[65]. The SNO experiment measures both the CC and NC reactions in a heavy water target, using them to extract the electron neutrino and total neutrino fluxes, respectively[4]. This measurement gives direct evidence for neutrino oscillations, but relies on theoretical calculations of the cross section using χ PT. Once again, if the clean measurement of d_R from MuSun confirms the existing estimate from the 3N system then there will be little effect, but a significant deviation could have major implications for these other 2N systems. This connection to the solar p+p fusion reaction and to solar neutrino measurements is the reason this experiment was named MuSun.

2.6 Summary

Muon physics plays an important role in understanding the weak interaction, whether by characterizing the strength of the weak interaction with the muon decay rate or by revealing more subtle details of the hadronic currents in muon capture reactions. The MuCap has previously measured muon capture on the proton to high precision, constraining the pseudoscalar nuclear form factor F_P . Now MuSun aims to perform a similar measurement of muon capture on deuterium, which will constrain the d_R parameter of χ EFT to 20% rather than the 100% of the current experimental 2N estimate. This will provide an important test of the theory predictions from 3N measurements, and may have serious theoretical implications if a significant discrepancy is discovered.

Chapter 3

EXPERIMENTAL DESIGN CONSIDERATIONS

Previous experiments aiming to determine the rate of muon capture on deuterium have struggled with several challenging systematic effects that prevented precise and unambiguous interpretation of their results. The MuSun experiment incorporates several new concepts to address these challenges. This chapter will provide an overview of the experimental strategy as well as design considerations to mitigate the most significant obstacles. A more detailed description of the various components of the experiment will follow in chapter 4.

3.1 *Lifetime Method*

Nuclear muon capture on deuterium is a rare process given by



The most obvious approach to measuring the capture rate would be to directly count these reaction products, but with only neutral particles in the final state this is quite challenging. Neutron detectors are difficult to precisely calibrate, and the three-body final state produces neutrons smoothly distributed over a wide range of energies as discussed in chapter 2. The direct neutron detection method is further complicated by a background of 2.45 MeV neutrons produced in muon-catalyzed fusion reactions. The large uncertainties of most previous muon capture experiments were due in part to the challenges of this direct detection approach.

The MuSun experiment instead uses a technique known as the "lifetime method", first introduced by the 1981 experiment at Saclay[12], which focuses on measuring the Michel

electrons produced by normal muon decays:

$$\begin{aligned}\mu^- &\rightarrow e^- + \bar{\nu}_e + \nu_\mu. \\ \mu^+ &\rightarrow e^+ + \nu_e + \bar{\nu}_\mu.\end{aligned}\tag{3.2}$$

The decay rate for positively and negatively charged muons is the same according to CPT (charge, parity, and time) symmetry¹. However, negative muons form muonic deuterium atoms and may undergo nuclear capture, while positive muons are repelled by the nucleus and would need to overcome the Coulomb potential barrier before having a chance to interact. Thus, while positive muons decay at the free muon decay rate λ_{μ^+} independent of the target material, negative muons disappear slightly faster due to the additional capture rate Λ_d .

The free muon decay rate is $455,170.2 \text{ s}^{-1}$ while the capture rate in deuterium is approximately 400 s^{-1} , so the large majority of muons stopping in the target will still decay into electrons. The number of muons remaining in the target after a given time, $N(t)$, approximately follows an exponential decay curve with total disappearance rate λ_{μ^-} :

$$N(t) \approx N_0 e^{-\lambda_{\mu^-} t} \approx N_0 e^{-(\lambda_{\mu^+} + \Lambda_d)t}.\tag{3.3}$$

The number of decay electrons observed at a given time is proportional to the number of muons remaining in the target, so λ_{μ^-} may be extracted by fitting the muon lifetime distribution. The extra contribution due to muon capture may then be isolated by comparing λ_{μ^-} to λ_{μ^+} .

The MuSun experiment consists of three primary detector systems, shown in figure 3.1. A set of entrance detectors monitor the muon beam, and start the lifetime measurement when a muon is detected. The muons then stop in a pressure vessel containing deuterium gas, with the TPC monitoring the muon track to ensure clean stops in the center of the target. The muons remain in the TPC until they decay, at which point the decay electron passes through the cylindrical electron detectors surrounding the target. The electron detection

¹For negative muons in muonic atoms there is actually an $\approx 0.5\%$ reduction in the decay rate due to relativistic time dilation [26].




Figure 3.1: Rendering of the MuSun detector arrangement. Muons pass through a set of entrance detectors before stopping in the active target while being tracked by the TPC. The decay electrons then exit through a series of concentric cylindrical detectors.

time marks the end of the lifetime measurement, and plotting the time difference for many events produces the decay time distributions needed for the lifetime method.

The free muon decay rate has already been measured to the ppm level by the MuLan experiment[68]. The MuSun experiment therefore focuses primarily on measuring the negative muon decay rate, with only limited positive muon data collected for the purposes of systematics studies.

3.2 Precision Goal

While the lifetime method eliminates difficult systematic uncertainties related to counting the capture neutrons, it does require very high precision to measure the difference in disappearance rates. The MuSun precision goal of 1.5% of the capture rate requires a 6 s^{-1}

uncertainty, with an equal division into statistical and systematic uncertainties giving an error budget of 4 s^{-1} for each. This corresponds to a 10 ppm level measurement of the total disappearance rate, requiring both high statistics and detailed understanding of the remaining systematic effects.

Approximately 10^{10} muon decay events are required to achieve the statistical uncertainty goal. Reaching that number of events in a reasonable amount of time is challenging, because the lifetime method relies on having a single muon in the target at a time. Events with multiple muons introduce an ambiguity regarding which entrance time to compare to, and so they must be rejected. To avoid muon pileup the experiment uses a muon-on-request (MORE) system, which diverts the beam after a muon entrance is detected. This system allows for a three times higher data collection rate, but introduces several complications which are discussed in appendix B. The required statistics took approximately six months to collect, split over the R2014 and R2015 production runs at the Paul Scherrer Institut (PSI) in Switzerland. The MuSun data collection campaign will be discussed in more detail in chapter 6.

Reaching the target systematic uncertainty requires a detailed understanding of all potential systematic effects. Because the lifetime method measures the total muon disappearance rate, it is sensitive to any effects that might distort the observed muon lifetimes. Chapter 10 will cover the main sources of systematic uncertainty, with later chapters focusing on the analysis of specific effects. However, several key systematics can be mitigated from the beginning with careful experimental design. The remainder of this chapter will focus on the requirements for the MuSun target to minimize any undesired effects.

3.3 Muon Kinetics

In addition to muon capture on the nucleus, a variety of other processes are possible for muonic deuterium atoms as shown in the kinetics diagram in figure 3.2. MuSun is specifically interested in only the capture rate from the doublet state Λ_d , and the presence of other reactions and decay channels may significantly interfere with the desired measurement. To

obtain an unambiguous determination of the capture rate, the experimental running conditions are optimized to suppress these interference effects.

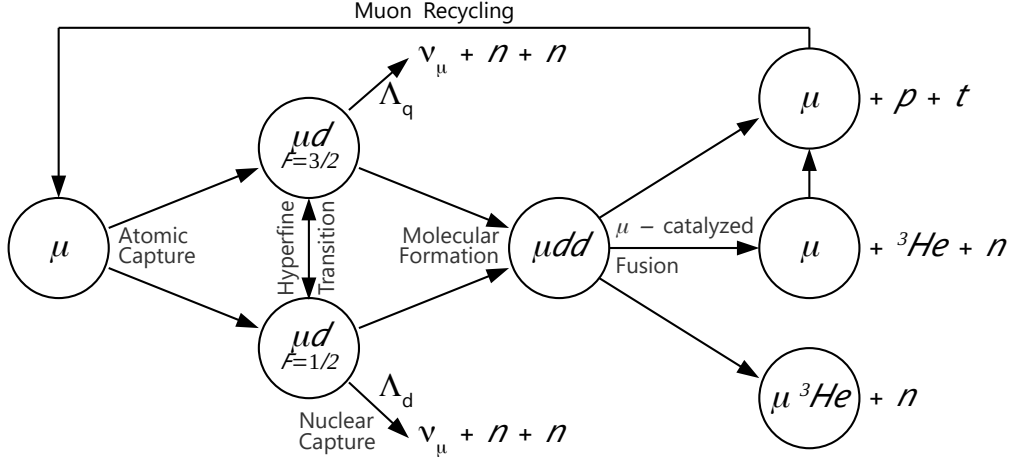


Figure 3.2: Simplified diagram of muon kinetics in deuterium gas. After forming muonic atoms the muons may either decay, capture, or participate in a molecular formation and subsequent muon-catalyzed fusion event. Fusions typically leave the muon unbound, allowing it to undergo atomic capture on a new atom and restart the process.


3.3.1 Hyperfine States

A muon passing through the target gas will gradually slow until it has a low enough energy to undergo atomic capture on deuterium. The resulting muonic deuterium atom will initially be in a highly excited state, which is then quickly de-excited through the Auger effect, radiative transitions, and Coulomb collisions. Once the atom reaches the 1S ground state it will populate one of the two hyperfine states, the doublet ($F=1/2$) and the quartet ($F=3/2$), according to their statistical weights. Due to the helicity structure of the weak interaction, capture from the quartet state is suppressed by roughly 40 times relative to the doublet state. The relative population of the quartet state must be well understood to avoid a systematic bias in the measurement of Λ_d .

Process	Symbol	Units	Value		Ref
			300K	30K	
free muon decay rate		s^{-1}	455,170.2		[68]
μd quartet capture rate	Λ_q	s^{-1}	~ 10		[18]
μd doublet capture rate	Λ_d	s^{-1}	~ 400		[32]
^3He capture rate	Λ_{He}	s^{-1}	2216(70)		[2]
hf transition $q \rightarrow d$	λ_{qd}	μs^{-1}	35(5)	37.0(4)	[37]
hf transition $d \rightarrow q$	λ_{dq}	μs^{-1}	10.5	0	
quartet $dd\mu$ formation	λ_q	μs^{-1}	3.75	3.98(5)	[11]
doublet $dd\mu$ formation	λ_d	μs^{-1}	2.549(23)	0.053(3)	
^3He fusion fraction	β	-	0.590(6)	0.517(15)	
^3He sticking probability	$\bar{\omega}$	-	0.1206(6)		

Table 3.1: Muon kinetics transition rates and branching ratios.

The quartet state may transition to the doublet state via a collisional charge exchange reaction in which the muon is transferred to a second deuterium atom with opposite spin:



The reverse reaction is suppressed by the 49 meV energy difference between the hyperfine states, making this a good way to depopulate the unwanted quartet state. Because the reaction requires collision with a second deuterium atom, the reaction rate is proportional to the target gas density.

3.3.2 Muon Catalyzed Fusion

Another major complication for the measurement of the capture rate is a process known as muon-catalyzed fusion, shown on the right side of figure 3.2, which has been studied

extensively by Balin et. al.[11]. Starting from either hyperfine state of muonic deuterium, a $dd\mu$ molecule may be formed after collision with another deuterium molecule:



This is a resonant process relying on a matching of the kinetic energy to a loosely bound state, from which the released energy is absorbed into vibrations and rotations of the newly formed molecule. As such it has a strong temperature dependence, as shown in figure 3.3.




Figure 3.3: Molecular formation rates λ_q and λ_d from the two hyperfine states (left) and the hyperfine transition rate λ_{qd} (right) as a function of temperature. Figure reproduced from [11] with modifications.

Because the muon is 200 times more massive than the electron its orbital radius is greatly reduced, and the radius of the $dd\mu$ molecule is similarly reduced relative to an ordinary deuterium molecule. At this distance there is a significant probability of overcoming the Coulomb potential barrier, allowing the two nuclei to rapidly undergo nuclear fusion. There are two possible reaction pathways, producing either a proton and a triton (${}^3\text{H}$) or a neutron and a ${}^3\text{He}$. After the reaction the muon typically becomes unbound, remaining in roughly the same location in the target. The muon then quickly binds to another deuterium atom,

which is called recycling since it returns the muon to the initial state in the kinetics diagram. However, in the case of a ${}^3\text{He}$ fusion the muon may remain bound creating a $\mu^3\text{He}$ atom. The chance to remain bound is called the sticking probability, and is about 12% as shown in table 3.1.

Because the muon only acts as a catalyst and is not consumed in these fusion reactions, they do not directly affect the disappearance rate. However, it is still important to suppress these reactions as they can indirectly interfere with the measurement in multiple ways. The recycled muons will once again statistically populate both hyperfine states, preventing the complete elimination of the quartet state. The $\mu^3\text{He}$ atoms also a concern, as they have a higher capture rate than in deuterium. Finally, any analysis cuts that are biased with respect to fusion events will distort the observed muon lifetime, as the chance of a fusion occurring is higher for longer-lived muons. This is a very important effect and much work has gone towards mitigating it, which will be discussed in detail in chapters 9 and 13.

3.3.3 *Optimized Running Conditions*

The MuSun target conditions are optimized to satisfy three requirements:

1. Maximize population of doublet hyperfine state
2. Minimize population of quartet hyperfine state
3. Minimize production of $\mu^3\text{He}$

These objectives may be satisfied simultaneously by choosing a high gas density and a low operating temperature. High density promotes the collisional de-excitation between the quartet and doublet hyperfine states, satisfying the first two requirements. A low operating temperature reduces the molecular formation rate, preventing the subsequent fusion reactions and reducing both the $\mu^3\text{He}$ population and muon recycling back into the quartet state. As shown in figure 3.3 the molecular formation rate is only suppressed for the doublet state while

remaining high for the quartet state. Direct measurement of the fusion time distribution, such as the detection of neutrons from $n+{}^3\text{He}$ fusions, may therefore be used to infer the relative populations of the hyperfine states as a function of time.

The kinetics diagram in figure 3.2 may be translated into a system of differential equations for the population of each state. Solving these equations gives the populations as a function of time, and the results for various temperature and density combinations are shown in figure 3.4. The gas density is expressed as a fraction of the density of liquid hydrogen, $N_{LH_2} = 4.25 \times 10^{22}$ atoms/cm³. The optimized MuSun running conditions are a temperature of 31 K and a density of 6.5% LH2, corresponding to a pressure of 5.1 bar. This configuration suppresses the quartet population fraction to the 10^{-3} level within the first 3 μs , and keeps the $\mu^3\text{He}$ population below the 10^{-2} level.

3.4 Capture on Other Nuclei

Recall from section 2.2.3 that the capture rate scales roughly with the fourth power of the atomic number. Because the lifetime method is sensitive to any deviation from the free muon decay rate, captures on any nuclei other than deuterium will bias the measurement. Table 3.2 shows measured capture rates for various relevant materials.

Figure 3.5 shows the stop fraction required to produce a 2 s^{-1} disappearance rate shift as a function of atomic number. The sensitivity initially increases sharply with Z , but for high- Z atoms the sensitivity falls again. This occurs because the high capture rate means that as few as 5% of the stops emit decay electrons, and the signal that is observed falls off rapidly and is largely gone before the 1 μs fit start time. The experiment is therefore most sensitive to materials with $Z \approx 10$.

3.4.1 Wall Stops

One obvious way for muons to capture on atoms other than deuterium is if they directly impact the pressure vessel or other parts of the apparatus rather than stopping cleanly in the center of the gas volume. These events are collectively referred to as wall stops, and were

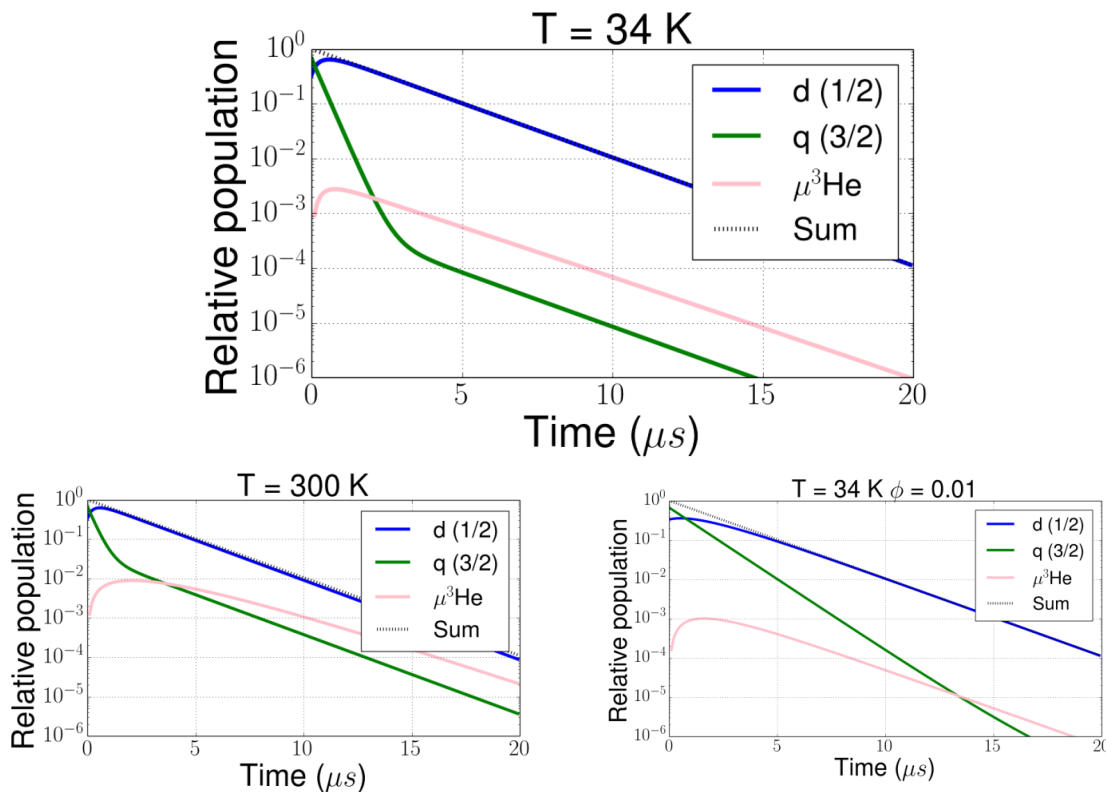


Figure 3.4: Populations of both hyperfine states and $\mu^3\text{He}$ as a function of time. The top plot shows the MuSun target conditions, while the bottom plots show the effects of increasing temperature or decreasing gas density

a major concern in previous experiments. Wall stops are difficult to eliminate even with a tightly focused muon beam and a large target gas volume, as there is always the possibility of the muon scattering in the entrance detector stack or the front flange of the pressure vessel.

The MuSun experiment includes a cryogenic time projection chamber (TPC) inside the gas vessel, to ensure the muons have properly entered the target and stop cleanly in the deuterium. The incoming muons ionize the gas in their path, and the TPC uses an electric field to collect the ionization electrons in an array of anode pads at the bottom of the TPC. This allows a full three-dimensional reconstruction of the muon track from the pad

Element	Capture Rate Λ_Z ($10^3 s^{-1}$)
Deuterium	~ 0.4
Helium-3	2.1
Nitrogen	69
Oxygen	102
Aluminum	705
Iron	4,411
Silver	11,070
Tungsten	12,360

Table 3.2: Capture rates on materials present in the MuSun experiment[64].

coordinates and the charge drift time, while keeping the target volume free of any foreign material. A fiducial volume cut selects muons stopping near the center of the TPC and eliminates events where the muon could have potentially hit the walls.

Although the TPC is very effective, it cannot reconstruct tracks with perfect accuracy. In particular, the possibility of muon catalyzed fusion products necessitates specialized tracking algorithms that will be described in chapter 9. These algorithms require trade-offs that reduce the tracking accuracy for non-fusion muon tracks, allowing some muons to escape the active volume. Thus, while the TPC largely eliminates the original problem of wall stops in the pressure vessel, there is now a new possibility of wall stops in the structure of the TPC itself. The TPC is therefore constructed using primarily silver and tungsten on all exposed inner surfaces, so that any muons that reach these surfaces will have minimal effect.

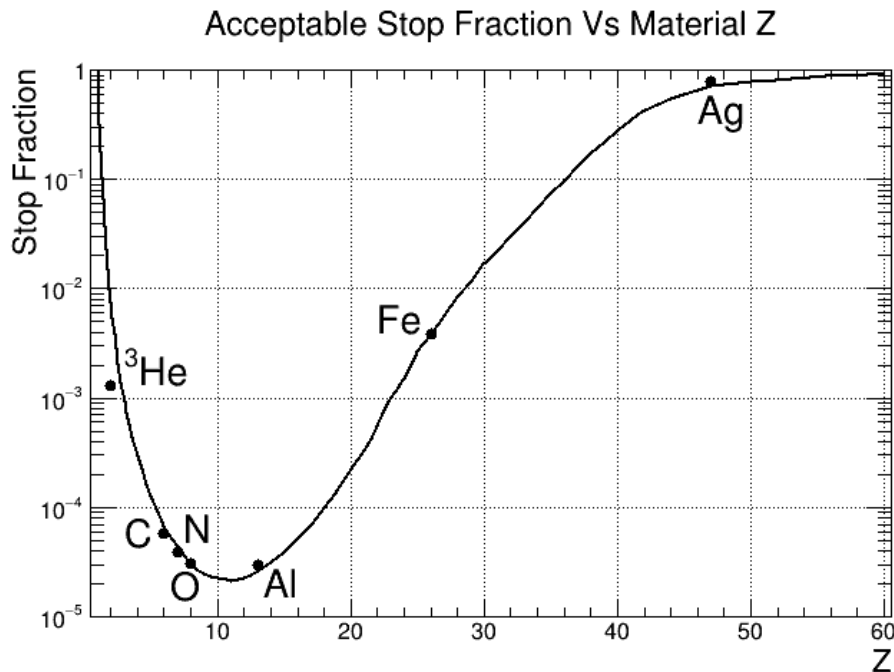


Figure 3.5: Stopping fraction required to produce a 2 s^{-1} disappearance rate shift when fitting starting at $1 \mu\text{s}$. The curve is based on an approximate model of the capture rate as a function of Z while dots use real measurements, both of which are taken from [64].

3.4.2 Gas Purity

Chemical Purity

Another source of non-deuterium captures is from impurities in the target gas, which have an inordinately large effect due to muon transfer onto these impurities. The small orbital radius of muonic deuterium atoms allows them to pass through the electron clouds of other atoms as if they were neutral particles. Once inside the electron cloud of a high- z atom, the muon is exposed to the full charge of the nucleus and is stripped away from the deuteron with transfer rates λ_{dz} shown in table 3.3. Because this process requires a collision the effect scales as $c_Z \lambda_{dz}$, where c_Z is the impurity concentration. The concentration must be constrained to

the ppb level to counteract the large transfer rates, well below the levels indicated in figure 3.5.

Element	Transfer Rate λ_{dz} ($10^9 s^{-1}$)	Capture Rate Λ_Z ($10^3 s^{-1}$)
Nitrogen	145 ± 2	69
Oxygen	63 ± 5	102


Table 3.3: Muon transfer rates from deuterium to impurities, assuming thermalization is much faster than transfer which should be the case for the MuSun experiment [28].

The most likely impurities are the constituents of air, mainly water vapor, nitrogen, and oxygen. The target operating temperature of 31 K is below the freezing point of any of these, but for nitrogen the equilibrium vapor pressure is still appreciable at about 200 ppb. During the R2013 commissioning run a high statistics dataset was taken with doped gas containing roughly 16 ppb of nitrogen, with additional tests performed with approximately 0.5 ppm and 2 ppm impurity at a higher operating temperature. Analyzing this data, it was determined that a 1 ppb nitrogen impurity produces a $4 s^{-1}$ shift in the observed disappearance rate. It is therefore imperative to keep the impurity level extremely low and also to monitor it to within the ppb level.

The purity target is achieved using the circulating hydrogen ultra-high purification system (CHUPS) developed by the Petersburg Nuclear Physics Institute (PNPI) group, which continuously filters the target gas. Gas samples were periodically extracted from the target and measured using a specially calibrated gas chromatography system capable of measuring concentrations down to 1 ppb [30]. In addition, an in-situ method for monitoring the impurity was developed which uses the TPC to detect nuclear recoil energy from nitrogen capture events. This impurity analysis will be discussed in section 10.1.2.

Isotopic Purity

A subtle effect may occur if the target gas contains both deuterium and protium (^1H). This mixture enables an additional form of muon catalyzed fusion involving a $pd\mu$ molecule formed from one of each isotope:



Unlike the muon catalyzed fusion reactions in pure deuterium, in this case the muons are ejected with an energy of 5.3 MeV. This high energy corresponds to a range of approximately 26 cm at the MuSun target gas density, allowing these muons to escape the TPC and potentially reach the walls. To mitigate this effect, the concentration of protium in the target gas must be kept below 100 ppm. This is well below the levels found in commercially available gas, so ultra-pure deuterium must be distilled on site.

3.5 Summary

The MuSun experiment improves on previous experiments in several ways. The lifetime method replaces the challenging problem of accurately measuring capture neutrons with the much easier task of counting muon decay electrons, at the cost of requiring high statistics and precise determination of the remaining systematic errors. Experimental design choices have addressed the most important systematic issues, with the TPC and gas purification system ensuring clean stops in deuterium gas, the conditions of which have been carefully optimized to suppress the unwanted quartet hyperfine state and minimize the rate of muon-catalyzed fusion events. The statistics goal of 10^{10} muon decay events has been collected over the course of about 6 months of operation using the PSI muon beam.

Chapter 4

EXPERIMENTAL HARDWARE

Measuring the rate of muon capture on deuterium to a precision of 1.5% while minimizing or correcting for any potential systematic issues required the development of multiple detector systems and other experimental components. As mentioned in chapter 3, the MuSun experiment consists of three main detector subsystems. Muons first pass through a series of entrance detectors, which ensure proper beam alignment with the target and provide a precise start time for each event. The muons then come to rest in the target, while being tracked by the time projection chamber (TPC) to ensure clean stops in the gas. Finally, the electrons produced as the muons decay are tracked by the concentric cylinders of electron detectors as they exit. The time differences between each muon entrance and the subsequent decay electron detection are used to produce the decay time histograms needed for the lifetime method.

In addition to the main detectors, there are several upstream beamline elements responsible for reducing backgrounds and producing the muon-on-request (MORE) beam structure mentioned in section 3.2. There is also a set of auxiliary liquid scintillator detectors which are used to monitor the neutrons produced by both muon capture and muon catalyzed fusion events. A sophisticated cryogenic and gas system is also required to maintain the deuterium target at the optimal conditions and to remove impurities. Finally, several different computer systems are needed to operate the experiment and monitor its status, as well as to process and record the output data. This chapter will describe each of these systems, starting with the production of the muon beam and following its path through the experiment.

4.1 Accelerator

The MuSun experiment makes use of the Swiss muon source ($S\mu S$) at the Paul Scherrer Institute (PSI), located in Villigen, Switzerland. The muon source is part of the high intensity proton accelerator complex, shown in figure 4.1. Starting with a hydrogen source, protons are stripped and accelerated to 870 keV by a Cockroft-Walton accelerator. Two cyclotron stages accelerate the protons to 72 MeV and then again to 590 MeV, with a beam current of 2.1 mA. This beam is then sent to several targets which produce secondary beams of muons, pions, or neutrons for various experiments.

The $S\mu S$ uses a graphite target that emits a variety of particles, including charged pions which decay to muons with a branching ratio over 99.99%. Pions stopped in the target produce mono-energetic positive "surface muons" with a momentum of 28 MeV/c. However, many pions instead decay outside the target, producing "cloud muons" that may be either positive or negative and have a wide momentum range from about 10 to 280 MeV/c. A series of quadrupole magnets and collimators focus the muons into beams, which are then directed to the various experimental areas. For the MuSun experiment the beam magnets are configured to select cloud muons with a momentum of 40 MeV/c, with a spread of approximately 3% FWHM.

The parity-violating structure of the weak interaction means that in the rest frame of the pion the decay muons must have their spin aligned with their momentum direction. For cloud muons a large pion momentum can allow a muon emitted backwards relative to the pion to still travel forward in the lab frame, but the beam is still expected to be polarized to some extent. This is not a major concern for negative muons, which lose most of their polarization after forming muonic atoms. However, positive muons will retain significant polarization even after coming to rest in the target.

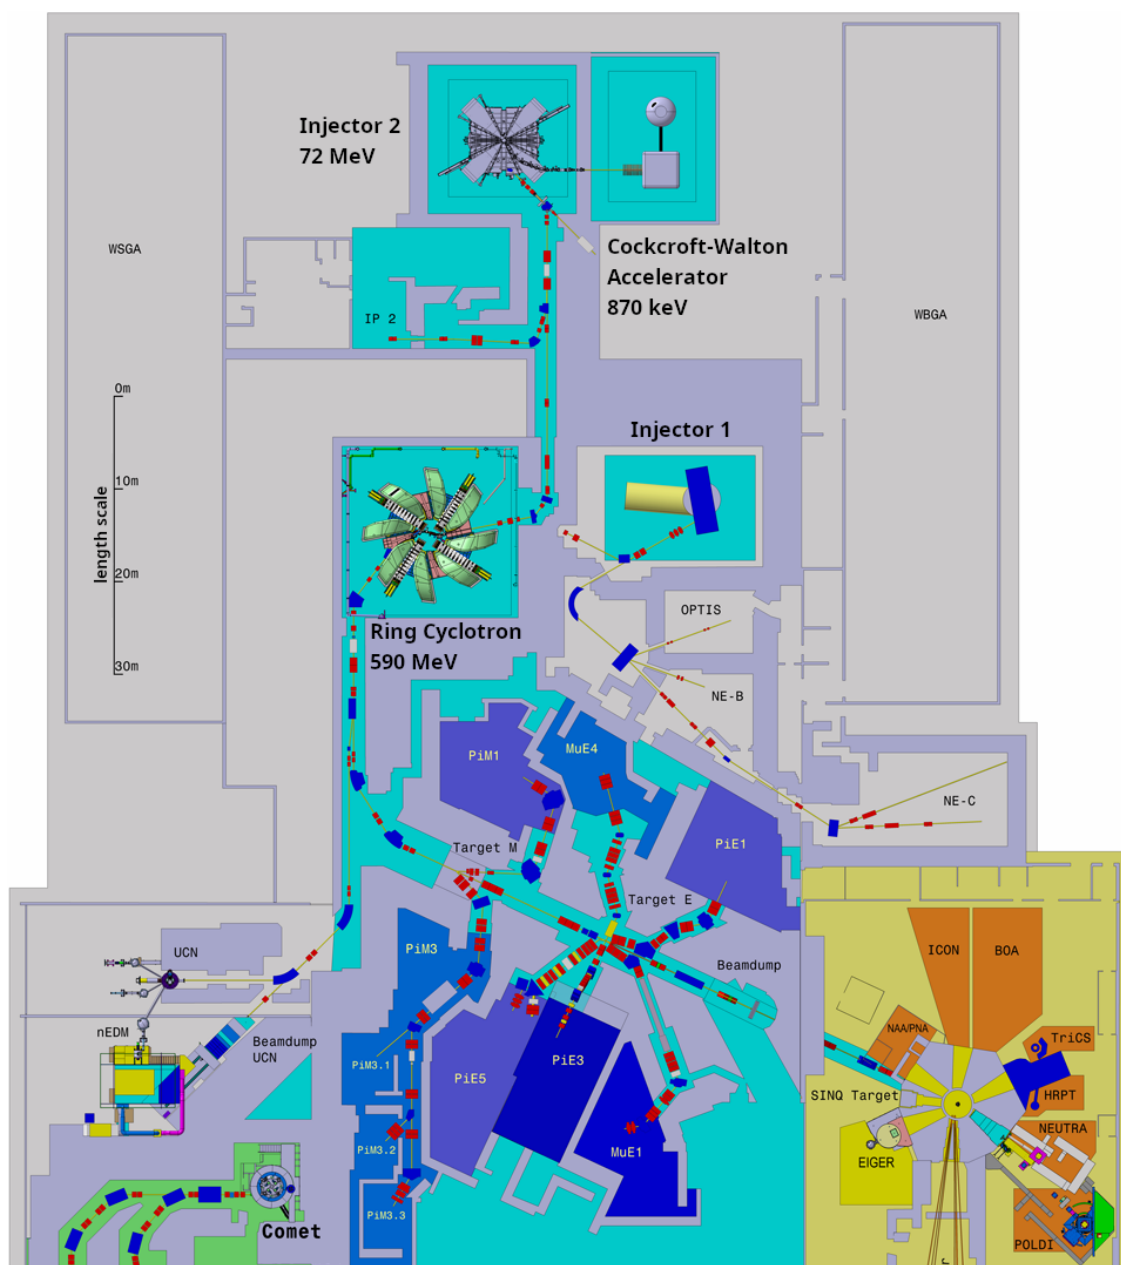


Figure 4.1: Floor plan of the PSI high intensity proton accelerator complex. Blue and red boxes along the beamlines indicate bending and focusing magnets, respectively. The MuSun experiment was located in the piE3 experimental area.

4.2 Beamline

The MuSun experiment includes some specialized beamline elements upstream of the main detector systems which are critical to reaching the precision target. Figure 4.2 shows a photo of the beamline, with the important elements labeled. The kicker is the key component of the MORE system, diverting the beam after a muon entrance to prevent additional pileup muons from interfering with the measurement. The separator filters electrons from the beam and diverts them without affecting the muons, greatly reducing the background signal in the detectors. These devices are positioned in front of a collimator such that diverted particles are stopped well upstream of the target.




Figure 4.2: Upstream beamline elements for the MuSun experiment. The beam enters from the left and passes through the kicker and separator, which deflect background particles before the beam is sent to the main detector.

The kicker was originally built for the MuLan experiment and was also used in the MuCap experiment [14]. It consists of a pair of parallel plates driven by two cabinets each containing

a stack of 17 MOSFET¹ cards. When the muon entrance signal is received from the target, these MOSFETs can charge the plates to a potential difference of 25 kV within 60 ns. At the MuSun beam momentum, this produces a deflection of 22 mrad. The beam is reverted to the normal configuration after a 25 μ s event window, allowing each muon to remain in the target undisturbed for up to 11 muon lifetimes.

The ratio of the normal beam rate to the kicked rate is called the kicker extinction factor. The kicker is unable to completely stop all muons due to tails in the transverse beam profile, but extinction factors of 70-80 are typical. Because of the upstream positioning of the kicker, there is a lag between when the kick signal is sent and when it is received, and an additional delay for in-flight muons to clear the beam pipe after the kicker activates. The result is an 800 ns delay after the muon entrance during which the beam remains at full rate, before dropping to the kicked rate.

The separator uses crossed electric and magnetic fields to act as a velocity selector, allowing only particles for which $qE = qvB$ to pass. The upstream bending and focusing magnets produce a beam with uniform momentum, but the muons have much lower velocities than the electrons due to their high mass. Thus, the separator can be tuned to let the slow moving muons pass straight through, while deflecting the much faster beam electrons. This filtering step is critical, as without it the electron contamination in the beam would exceed the number of muons by two orders of magnitude. The separator was held at a steady 180 kV during production, and required slow ramping to avoid dangerous sparking.

It is crucial for the operation of both the kicker and separator that the beam is parallel across its transverse profile, such that a relatively small angular deflection results in a completely separate distribution in phase space. Unfortunately, during the MuSun production runs an error in the PSI database caused one of the upstream quadrupole magnets to run at only half of its nominal current. This in turn produced a divergent beam at the kicker and separator, limiting their effectiveness. Extensive tuning of the downstream magnets

¹Metal oxide semiconductor field-effect transistor

was able to compensate for this problem to some extent and acceptable extinction factors were achieved, but the separator never reached full efficiency resulting in high beam electron backgrounds.

4.3 Entrance Detectors

When muons arrive at the MuSun detector, they first pass through a series of entrance detectors as shown in figure 4.3. Together, these detectors ensure the beam is properly collimated and provide a time and position measurement for each incoming muon. They are also connected to the MORE system, as well as contributing to software based pileup protection.




Figure 4.3: The MuSun entrance detector stack. The diagram is not to scale, and the μSC is only 45 x 45 mm square. For R2015 the μSCa was upgraded to a 96 x 96 mm square scintillator with a 45 mm diameter aperture.

4.3.1 Detectors

Muon Veto (μ SCa)

The first entrance detector is the annular muon scintillator (μ SCa), designed to veto muons that are not well collimated. It consists of a 5 mm thick scintillator with a circular hole in the center, which is directly followed by a lead collimator with a matching hole. Thus, muons that are on the beam axis pass through both holes, while off-axis muons are detected by the scintillator and then stopped by the collimator. The lead collimator should minimize the impact of off-axis particles, but in some cases secondary particles from stops in the lead may still affect the experiment so it is useful to detect such events.

During the R2014 and earlier runs the μ SCa used a 70 x 70 mm square scintillator with a 35 mm diameter aperture, and the light was detected with a photomultiplier tube (PMT) coupled to one side. For the R2015 run this was upgraded to a new 96 x 96 mm square scintillator with a 45 mm diameter aperture, read out from either side by strips of 16 silicon photomultiplier (SiPM) detectors. The use of multiple readouts allowed for an increased aperture size while maintaining a uniform detection efficiency, which was a problem for the old single PMT design. The overall larger size increases coverage for off-axis muons, which can be an issue due to the difficulties with the beam focusing magnets mentioned above.

Muon Scintillator (μ SC)

After the veto counter is the muon scintillator, responsible for precisely measuring muon entrance times. The (μ SC) is 45 x 45 mm square, and is only 500 μ m thick to minimize scattering. This scintillator is still coupled to a PMT on one side, as uniformity is not a major concern without the central hole in the scintillator. The output signal is split and sent to two discriminators, one with a high threshold (μ SCHi) to identify muons and another with a low threshold sensitive to electrons. Electron triggers which pass the low threshold but not the high threshold (μ SCLo) are used to monitor the approximate beam electron rate during beam tuning.

Muon Tracker (μ PC)

The final entrance detector is the muon multiwire proportional chamber (μ PC), used to monitor the transverse positions of incoming muons. The (μ PC) consists of two 50 μ m Mylar windows separated by 35 mm, forming an airtight chamber. The chamber is pressurized to 1 bar with a gas composed of 49.9% argon, 49.9% ethane, and 0.2% Freon. Two groups of 24 anode wires, one vertical and the other horizontal, are strung across the chamber with a pitch of 2mm. 25 μ m Mylar foils on either side of each anode plane act as cathodes, and are operated at 2500V.

When a particle passes through the chamber it will leave behind an ionized track, the charge from which collects on the nearest anode wire. The argon is the primary ionizing gas, while the ethane is a quenching gas added to suppress spurious pulses. Coincidences between wires in the two differently oriented planes provide a measurement of the beam profile in two dimensions. The beam electrons are nearly minimum ionizing particles (MIPs) while the muons deposit about 5 times more energy. The detection thresholds are therefore set such that muons consistently trigger the μ PC while electrons normally do not. The electron energy deposition has a long tail so it is possible for an electron to trigger an individual plane, but it is very rare for one to trigger both planes in coincidence.

4.3.2 Muon-On-Request System

For a clean determination of the muon decay times it is important that only a single muon is present in the detector at any given time. To reduce the number of muon pileup events, the signal from the μ SC is used to trigger the kicker and divert the beam for the remainder of the event window. The kicker is also triggered periodically by a clock (μ CL) signal, providing a way to monitor beam backgrounds and the effects of the kicker in the absence of a real muon entrance. The μ CL and μ SC triggers are treated identically by the experiment, but they are also digitized separately so the events may be distinguished later in the analysis.

During R2014 the clock operated at 2.5 kHz, but was only enabled for specific datasets.

It was enabled for the entire R2015 dataset but at a reduced rate of 1 kHz, resulting in a similar number of events to R2014. A dedicated muon clock shift was also performed in R2014 in which the real muon signal was disabled and the clock rate was boosted to 30 kHz, although this alters the background structure somewhat.

4.3.3 Pileup Protection

The muon-on-request system cannot fully eliminate muon pileup, both because the kicker extinction is not perfect and because of the 800 ns delay between a μ SC signal and the diversion of the beam. Any remaining events with multiple muons must be eliminated by software pileup protection cuts. A time window from $-25 \mu\text{s}$ to $+25 \mu\text{s}$ around each muon entrance is used for pileup protection.

4.4 Target

The deuterium gas is contained in a cylindrical pressure vessel 36 cm long with a 10.1 cm radius. There are thick stainless steel flanges capping each end, but the cylinder is made of 2.5 mm thick aluminum to reduce scattering of outgoing electrons. Muons enter through a hemispherical beryllium window in the upstream flange, with a 30 mm radius and 0.5 mm thickness. Beryllium is used because its low Z minimizes scattering of the incoming muon beam, allowing a tighter stopping distribution. The vessel was stress tested and validated to withstand pressures up to 15 bar at cryogenic temperatures.

The standard operating conditions are a temperature of 31 K and a pressure of 5.1 bar. Cooling is accomplished using a condenser producing 27 K liquid neon, which drains through a series of copper tubes throughout the deuterium vessel where it evaporates and circulates back to the condenser. This arrangement avoids using a rigid connection to isolate the target from vibrations, which was a problem early in the development of the experiment. The target can be cooled from room temperature in 12 hours.

The deuterium vessel is enclosed by a second pressure vessel pumped to $\sim 3 \times 10^{-7}$ mbar vacuum, to insulate the cryogenic components from the outside air. The vacuum vessel is

made of 4 mm thick aluminum, with Mylar shielding to reflect infrared radiation. A 150 μm kapton layer provides an entrance window for the incoming muon beam.

4.4.1 Time Projection Chamber

As discussed in section 3.4, the MuSun lifetime measurement is quite sensitive to capture in materials other than deuterium. The TPC tracks incoming muons in three dimensions without introducing any foreign material in the target volume. A large electric field causes ionization electrons from particle tracks to drift to an anode plane segmented into a grid of charge collecting pads. When charge is collected the coordinates of the triggered pad along with the delay from the drift time of the charge allow a three dimensional reconstruction of the original track position. This system enables fiducial volume cuts to select only clean stops in the target gas, and is a crucial improvement over previous experiments. Because the TPC is central to the operation of the experiment it will be discussed in detail in chapter 5.

4.4.2 Spin Precession

As discussed above, the muons in the beam are spin polarized in the direction of their momentum. Negative muons lose this polarization when they form muonic atoms, as large spin-orbit couplings induce fast spin precession for random duration during the atomic de-excitation cascade. Positive muons do not form muonic atoms however, and thus retain their polarization even after coming to rest in the target. The positive muons will then precess in the earth's magnetic field (≈ 0.5 G) with a period of approximately 150 μs . As these muons decay the resulting positrons will be preferentially emitted in the direction of the muon spin, producing a time-dependent emission distribution.

Because the muon stop distribution is not perfectly centered and the electron detectors have a non-uniform coverage, the precession changes the average detection efficiency. The natural precession frequency is particularly problematic, with the 25 μs event window including roughly 1/6 of an oscillation resulting in a slope across the window. To mitigate this effect, a saddle-coil muon spin rotation (μSR) magnet is wrapped around the insulation

vacuum vessel to apply a 50 G magnetic field in the target. This reduces the precession period to $1.5 \mu\text{s}$, resulting in a visible oscillation which is much easier to separate from the exponential muon decay. The magnet is not strong enough to bend particle tracks for momentum selection, as is common for collider experiments, and is normally turned off for negative muon operation.

4.5 *Electron Detectors*

The electron detectors are the final core component of the experiment, tracking the outgoing decay electrons and providing a precise end point for the lifetime measurement. They form concentric cylinders around the target volume, as shown in figure 3.1 above. Both ends of the cylinder are open, the front to allow the muon beam to pass and the back to fit the gas circulation system. As such, the solid angle coverage of these detectors is approximately 3π .

4.5.1 *Electron Trackers (ePCs)*

The electron proportional chambers (ePCs) reconstruct the trajectories of the outgoing electrons. They consist of anode wires strung parallel to the beam axis, sandwiched between two cathode layers constructed of a series of aluminized Mylar strips wrapped in a helical arrangement. The cathode strips are oriented at a $\sim 45^\circ$ angle with respect to the anode wires, but in opposite directions for the two layers as shown in figure 4.4. An electron passing through the chamber will thus produce a signal in the unique combination of anode wire and cathode strips which intersect at that location. The ePCs are filled with the same gas mixture as the μPC .

The inner chamber (ePC1) has a length of 58 cm and a diameter of 38.4 cm, and is mounted just outside of the insulation vacuum vessel containing the target. It uses 512 anode wires and each cathode plane is composed of 192 strips. The outer chamber (ePC2) is 80 cm long and has a diameter of 64 cm, with 1024 anode wires and 320 cathode strips per plane. The ePC1 anodes are kept at 2750 V while ePC2 is slightly lower at 2600 V.

All three planes are equipped with separate electronic readout chains. Sectors of 16




Figure 4.4: Illustration of the ePC winding pattern.

wires are connected to custom built cards mounted directly to the chambers containing preamplifiers and discriminators. These are connected to a time-to-digital converter (TDC) system, which records the time of each signal. The TDCs have a time resolution of 30 ns, but the detectors themselves have a time resolution of only about 100 ns.

This system proved somewhat susceptible to noise sources such as afterpulsing and pickup, particularly for the cathode strips. The hit rates for each wire were monitored throughout production, and amplifier cards with unusually rates could usually be fixed by reseating them and ensuring they were properly grounded. Some individual "hot wires" also had extremely high hit rates, and these signals were masked out to avoid overwhelming the computer system with large amounts bad data. One of the wires in ePC2 broke during the R2014 run, resulting in very high current draw at the standard operating conditions. The sector containing this wire was ramped down to 2300 V with a 100 M Ω shunt to ground. This bad sector as well as the sectors on either side of it were therefore unable to track particles reliably, reducing the solid angle coverage of the detector.

4.5.2 *Electron Scintillator (eSC)*

The decay electron times are measured by a hodoscope with a barrel diameter of 780 mm, constructed from 16 individual panels. Each panel is 900 x 150 mm and consists of two 5 mm thick layers of BC404 scintillator. The panels are each read out by four photomultiplier

tubes, one at either end of both scintillator layers. The PMT outputs are split, and one copy of the signal is sent to discriminators and then to TDCs which record time stamps for each channel. These TDCs subdivide the external 25 MHz clock with a 32-step interpolator for a timing resolution of 1.25 ns, although non-linearity of the interpolator forces most applications to rebin this signal to 40 ns. The other copy of the signal is sent to a set of waveform digitizers (WFDs) running at a 450 MHz sampling rate. By interpolating the digitized waveforms the eSC WFDs can achieve a timing resolution of approximately 1 ns.

Using multiple readouts for each panel provides several benefits. The detectors at either end of each scintillator provide redundancy and requiring a coincidence between them suppresses noise in the individual detectors. Timing and amplitude differences between the signals from either end also provide a coarse Z measurement, although this is much lower resolution than the ePC tracking. A 4-fold coincidence between all PMTs for a given segment is usually required to minimize afterpulsing and noise, as well as to eliminate events such as neutron scattering that only trigger a single scintillator layer.

4.6 Neutron Detectors

Finally, in addition to the main detector systems the experiment also includes eight supplementary liquid scintillator neutron detectors used for systematic error testing. These are mounted outside of the scintillator segments in four groups of two, as shown in figure 4.5. In R2014 and before, the experiment used eight Bicron detectors with 13 cm diameter and 5 cm length, for approximately 1% solid angle coverage. For R2015 seven of these were upgraded to DEMON detectors with 16 cm diameter and 20 cm length, significantly increasing the detection efficiency and energy range. The neutron detector waveforms are digitized by a 12-bit flash analog to digital converter (FADC) board sampling at 170 MHz.

The Z^4 capture rate scaling mentioned in section 3.4 means that nuclear capture becomes a dominant effect in high- Z materials. In particular the silver and tungsten used in the TPC ensure that wall stops produce a large capture neutron signal. Furthermore, the neutron detectors can also detect gamma rays, and the de-excitation cascade after atomic capture



Figure 4.5: Bicorn (left) and DEMON (right) neutron detectors mounted to eSC frame.

on a high-Z material is clearly visible. Together these signals produce a quite distinctive wall stop signature, with a prompt gamma ray pulse followed by a fast exponential decay neutron signal. Measuring the size of this signal as a function of muon stop position gives an independent measurement of the TPC track reconstruction accuracy and provides a data-driven method to optimize the fiducial volume cut. The neutron detectors are also used to monitor the 2.45 MeV neutrons produced by muon-catalyzed $^3\text{He} + n$ fusions.

4.7 Gas Purification System

The stringent gas purity requirements have been discussed in section 3.4.2, with targets of 100 ppm for isotopic impurities and 1 ppb for chemical impurities. Significant work has gone into the gas purification system, mainly by the PNPI group.

The isotopic impurity target is well below that of commercially available deuterium gas, which typically contain around 2000 ppm of protium. Therefore, the gas is first purified in a cryogenic distillation procedure, utilizing the difference in saturation vapor pressures to separate the hydrogen isotopes[6]. A diagram of the distillation column is shown in figure 4.6. Liquid hydrogen is vaporized in the reboiler at the bottom, and rises through the

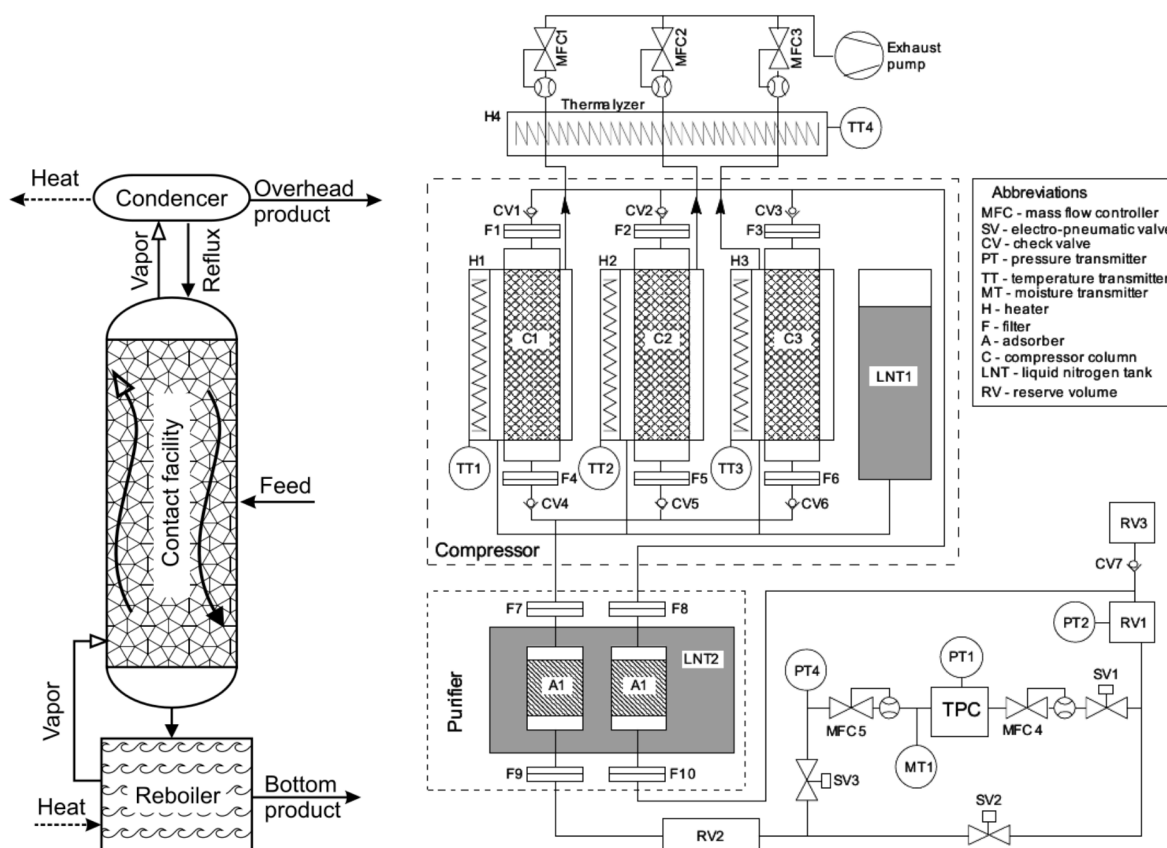


Figure 4.6: Simplified diagrams of the isotopic separation gas distillation column (left) and CHUPS gas circulation and purification system (right).

column before being liquefied at the condenser and draining back down. The column is filled with a spiral prismatic random packing designed to maximize surface area, increasing the interaction between the rising vapor and the draining reflux. The lighter protium component accumulates at the top and is discharged as a vapor, while pure deuterium collects at the reboiler and is extracted from the bottom of the column. Chromatography measurements of the resulting target gas reported a stable protium component of 16 ppm.

Figure 4.6 also shows a simplified diagram of the circulating hydrogen ultra-high purification system (CHUPS), previously developed for the MuCap experiment[31]. As the name

implies, CHUPS continuously pumps the deuterium to circulate it through the TPC and purification filters. This is accomplished using 1 liter compressor columns filled with activated carbon and cooled liquid nitrogen. The gas condenses through contact with the carbon, and then is extracted with electric heaters. There are three compressor columns connected in parallel, which are operated out of phase to maintain constant flow. A high pressure reserve volume (15 L at 100 bar) stabilizes any pressure variation.

Chemical purification is performed by passing the gas through synthetic zeolite adsorbents maintained at 77 K with liquid nitrogen. Zeolites are microporous materials commonly used as adsorbents and catalysts. Cooling the filters greatly increases the rate at which nitrogen, oxygen, and water impurities are removed. There are two purification stages, one before the gas enters the TPC and a second for gas leaving the TPC before it returns to the compressors. Gas samples were periodically extracted from the target and measured using a specially calibrated gas chromatography system capable of measuring concentrations down to one ppb[30].

4.8 Computer Systems

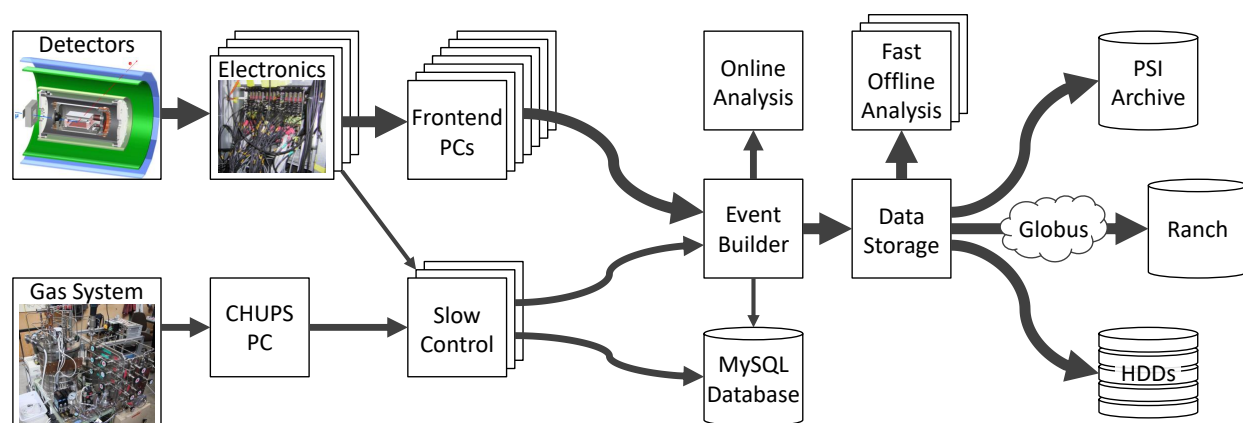


Figure 4.7: Block diagram of the MuSun data acquisition system. Data collected by the frontend PCs is consolidated by the event builder and saved to redundant storage locations.

The data acquisition system (DAQ) accumulates the data streams from the various hardware subsystems and packages them together into data files suitable for long term storage and analysis. The slow control system manages detector settings such as voltages and trigger levels, as well as environmental monitoring sensors, and records this data in a MySQL database. A small fraction of the incoming data is also processed immediately by the online analysis, allowing operators to monitor data quality in real time. Figure 4.7 shows a block diagram of the data flow through the data acquisition system.

4.8.1 Data Acquisition System

The MuSun DAQ is based on a customized version of the Maximum Integrated Data Acquisition System (MIDAS) developed at the PSI and TRIUMF laboratories. MIDAS is designed in a modular fashion, with separate programs known as frontends running on various computers. The detector system frontends run on computer boards mounted in the same electronics crates as digitizers that they manage. Data is transmitted from these computers over a private network to the central DAQ computer, where a program called the Event Builder consolidates it into a single stream. This data stream is then written out to the archives and to backup storage disks by another set of programs called loggers.

Data is collected as a series of ~ 140 ms long segments known as data blocks, with ~ 25 ms periods of downtime in between during which the data is recorded. This data acquisition cycle results in an overall DAQ live-time of 80%. Output MIDAS data files are truncated when they either reach a size limit or exceed 3000 data blocks, typically the former. Redundant copies of each data file are saved to the PSI archive, written to local disks, and sent over the internet to the supercomputer systems where the MuSun analysis is carried out. An additional slow control system manages high voltage supplies and experiment settings, performs environmental monitoring, and monitors critical safety systems. Some of this data is critical for event reconstruction and is included in individual data files, while the full set is saved to a central MySQL database.

4.8.2 *Online Analysis*

The DAQ system also includes an online analysis component which processes a small fraction of the data stream in real time. This is nearly identical to the first stage of the full offline analysis, but always skips to the newest available data block rather than proceeding sequentially. The output is fed to an Online Display program that produces several pages of detector readout and diagnostic display panels, provided an invaluable early look at the data quality. While collecting data the operators' main duties were split between monitoring the DAQ status and slow control readouts, and studying the online display screens for any anomalies. The real time feedback from these displays was also extremely helpful for diagnosing and fixing any problems that did arise.

While most potential error modes for the experiment would be immediately obvious in the online displays, some more subtle problems might only become apparent with higher statistics. For this reason, the MuSun experiment also operated a small cluster of dedicated analysis computers to process larger datasets. This "fast offline analysis" was able to process roughly half of the good data files, which were used to generate improved versions of the standard displays and for testing specialized studies such as the in-situ impurity analysis.

4.9 *Blinding Procedure*

To minimize bias introduced by the MuSun analysis, the data must be blinded. Hardware blinding is achieved by de-tuning the master clock by a secret value within 0.3% of the nominal value. This produces a corresponding shift in the measured disappearance rate of up to 1350 s^{-1} , which is sufficient to completely blind the expected $\approx 400 \text{ s}^{-1}$ deuterium capture rate. The value of the blinding was recorded and checked by an external member of the collaboration, and is unknown to any of the other members until the final analysis is completed. Different blinding values were used for each PSI run, so care must be taken when attempting to compare results across datasets.

An additional layer of software blinding is also applied when fitting the final lifetime histograms. This is produced by adding several random offsets to the fitted decay rates, with separate offsets for each PSI run, each institution, and each individual analyzer. There is also an additional offset applied to positive muon data. A script randomly generates the offset values in the range $(-500,500) \text{ s}^{-1}$, and they are stored in an encrypted format so they are unknown even to the user who generated them. This thesis will present blinded analysis results, and to avoid confusion they will typically be shown as shifts relative to a high-statistics reference fit since the exact blinded values are arbitrary.

4.10 Summary

The MuSun experiment is composed of several subsystems. It starts with a muon beam produced by the $S\mu S$ facility at PSI, which is directed through a series of upstream beamline components that shape the beam and remove unwanted backgrounds. The beam then passes through the entrance detectors before stopping in the TPC, where the majority of muons ultimately decay releasing an electron which is detected by the electron exit detectors. There are also a number of supporting systems, such as the CHUPS gas management system and the neutron detectors which can be used to monitor fusion events and wall stops in the TPC. Finally, the detector signals are collected and digitized by the DAQ, with hardware blinding achieved by a slight de-tuning of the master clock signal.

The size and complexity of the MuSun experiment is unusual, being substantially larger than a typically bench-top experiment while remaining just small enough that a single student can get hands-on experience with every individual subsystem. This is particularly true of this thesis, as while the author joined the collaboration relatively late and was not heavily involved in the hardware design phase he has worked to constrain a variety of different systematic issues to prepare for a final analysis result. Every hardware component mentioned here is relevant to some portion of the analyses presented in later chapters.

Chapter 5

TIME PROJECTION CHAMBER

The MuSun time projection chamber (TPC) is a novel and critical component of the experiment. Several of the later software and analysis chapters will be focused on particle tracking in the TPC, so it will be helpful to cover it in some detail here. This chapter will explain the TPC's operational principles and some of the physical constraints informing its design. It will then describe the specifics of the final hardware implementation and give an overview of the particle interactions in the TPC. Chapter 9 will follow up with a more detailed look at the TPC signals, and discuss ways of improving the particle tracking.

5.1 Constraints

The TPC tracks incoming muons in three dimensions without introducing any foreign material in the target volume. This system enables fiducial volume cuts to select only clean stops in the target gas, and is a crucial advancement over previous experiments. A muon passing through the deuterium gas will ionize atoms along its path, and the cloud of ionization electrons drifts to the anode plane due to an applied vertical electric field. The electrons are collected by a grid of charge-sensitive pads, with positions of the triggered pads giving the X and Z coordinates of the track while the Y coordinate may be reconstructed from the drift time. This procedure requires knowledge of the time the charge was deposited, in the case of the muon this is taken to be the muon entrance time from the μ SC. Beam muons travel at roughly 1/3 the speed of light and will come to a stop in a matter of nanoseconds while the drift times are on the order of microseconds, so this is a valid approximation.

5.1.1 Drift Field

To achieve accurate tracking the drift field must be strong enough to completely clear the electrons from the muon track and any subsequent reactions before the next muon entrance. With the standard MuSun event window being $25 \mu\text{s}$ long, we thus aim for a drift time of $15 \mu\text{s}$ for tracks at the very top of the TPC. The drift velocity v_d is obtained by solving the Langevin equation,

$$v_d = \frac{eE}{2m}\tau, \quad (5.1)$$

where E is the electric field, m is the particle mass, and τ measures the collisional timescale of the gas, proportional to the mean free path. At the MuSun target gas conditions this simplifies to $v_d \approx 5E$, with v_d in $\text{mm}/\mu\text{s}$ and E in kV/mm .

This yields the equation

$$t_{max} = \frac{h}{v_d} \approx \frac{h^2}{5V}, \quad (5.2)$$

where t_{max} is the target maximum drift time in μs , h is the total height of the active volume in mm , and V is the voltage across the active volume in kV . From this equation it is clear that it is not practical to extend the height of the TPC more than is necessary to contain the muon stopping distribution, as the required voltage would increase rapidly.

The drift field must also be large enough to quickly separate the electrons from their associated ions, otherwise recombination will reduce the observed signals. At the gas pressures used by the MuSun experiment this is a relatively small effect for the muon tracks, but more significant for fusion products with higher ionization density dE/dx .

A height of 71 mm was chosen with a cathode voltage of 80 kV , resulting in a $5 \text{ mm}/\mu\text{s}$ drift velocity and $t_{max} \approx 15 \mu\text{s}$ as desired. The relatively short height of the TPC leads to a significant number of wall stops, and a strict fiducial Y cut is always required. The horizontal dimensions of the TPC are not limited in the same way, and a larger size is used to reduce the number of stops near the edges of the TPC.

5.1.2 Frisch Grid

The positive ions also drift to the cathode, but due to their much larger mass they move three orders of magnitude slower than the electrons and remain in the TPC hundreds of times longer than the event window. As such only the induced charge from the electron is visible with the short time scale of the amplifiers, and the signal size depends on the initial Y position. This Y dependence may be eliminated by including a Frisch grid of fine wires just above the anode plane, which shields the anodes from any induced charge until the electrons pass the grid. By setting the grid voltage such that the electric field between the grid and the anode is much stronger than the field within the active volume it is possible to prevent charge loss in the grid. The condition for achieving grid transparency is derived in [19]:

$$\frac{E_g}{E} > \frac{1 + \rho}{1 - \rho} \qquad \rho = \frac{2\pi r}{d}, \qquad (5.3)$$

where E is the electric field in the main target volume as before, E_G is the electric field between the grid and the anode pad, r is the radius of the grid wires, and d is the distance between wires. This condition is achieved with a grid voltage of 3.6 kV, with r and d set to 0.025 mm and 0.4 mm respectively.

5.1.3 Gas Amplification

Typically a TPC would include a gas amplification stage beneath the grid, where the drift field is high enough that the drifting electrons themselves ionize the gas. This produces a cascade of increasing ionization referred to as gas amplification or gas gain, and when operating in this mode the TPC becomes a proportional counter. However, the unusual high pressure gas required for the MuSun experiment prevents this from being viable, as a high enough voltage cannot be reached.

Instead the TPC operates as an ionization chamber with no amplification, and must measure the small amount of charge deposited directly by the particle tracks. At the MuSun target gas density a minimum ionizing particle (MIP) deposits roughly 2 keV/mm. It takes 36.5 eV to produce an electron-ion pair, so this results in only about 50 electrons per mm!

The cryogenic preamplifiers discussed below achieve an impressive charge detection resolution of 250 electrons rms and muon tracks deposit at least 5 times more energy than MIPs, but pad dimensions larger than 10 mm are still required for good energy resolution.

The final design uses a grid of 8 pads in Z and 6 in X, with a pad length of 15.75 mm. The 4 central X rows have pad widths of 18 mm, while the edge pads have widths of 12 mm. The X edge pads were intended to be used as veto pads for muons leaving the fiducial volume, which will typically be nearly parallel to the Z axis. The non-uniform X arrangement thus slightly increases the width of the fiducial volume and .

5.2 Final Design

5.2.1 Geometry

Figure 5.1 shows the CAD design of the MuSun TPC, with a photo of the final device shown in figure 5.2 for comparison. Table 5.1 gives a summary of the most important TPC parameters. All components are designed with rounded corners to minimize the possibility of sparking. A series of 7 field wires wrap around the TPC at equally spaced Y intervals to improve field uniformity. These are connected to a voltage divider, with the resistors tucked behind and to the sides of the TPC to reduce the chances of being hit by punch-through muons. The field wires are separated from the edges of the anode plane by more than the length of a pad, so muons exiting the TPC in these directions are likely to stop before reaching the wires.

5.2.2 Materials

The anode, cathode, grid, and field wires are by far the most likely components to be hit by incoming muons, and are therefore made of high-Z materials to minimize any distortion of the lifetime measurement as discussed in section 3.4. The Anode has been coated with a layer of silver, and the cathode is constructed with a thin silver foil held by a stainless steel frame at the edges. The field wires are silver-coated tungsten, with section at the front




Figure 5.1: Top: top view of the TPC CAD design, showing the grid of anode pads. Bottom: side view, with the anode plane at the bottom and the cathode at the top. The beam enters from the left in both cases. Dimensions are given in mm.




Figure 5.2: Front view of the MuSun TPC, muons enter traveling along the Z axis.

replaced with thin wires of pure silver to reduce the probability of muons hitting the wires as they enter. The grid wires are made of tungsten and are strung in the X direction, and the X beams of the bottom frame are also tungsten to minimize differential thermal expansion and maintain constant tension on the wires. The Z beams of the bottom frame are stainless steel, as this allows a better bond between the grid wires and the frame.

The vertical support posts at the corners of the TPC must be non-conductive, as they separate the anode and cathode with a voltage difference of 80 kV. These are constructed from Macor, a glass-ceramic material designed for cryogenic applications. These posts and the stainless steel components listed above are all located at the edges of the TPC, which are much less likely to be impacted by the muon than the faces.

Parameter	Value
pad length	15.75 mm
pad width (center)	18 mm
pad width (edge)	12 mm
grid to cathode distance	71 mm
grid to anode distance	1.44 mm
grid wire radius	0.025 mm
grid wire spacing	0.4 mm
grid voltage	3.6 kV
cathode voltage	80 kV
grid to cathode field	11 kV / cm
grid to anode field	25.7 kV / cm

Table 5.1: Specifications for the final TPC design used in production.

5.2.3 Readout

Flat kapton cables carry the signals from the anode plane through the deuterium vessel flange to a set of cryogenic preamplifiers located within the insulating vacuum and cooled by a dedicated liquid nitrogen line [61]. The preamplifiers work in a charge-integrating mode to maximize the energy resolution, and this signal is then sent to a shaping amplifier. The combination of the drift time between the grid and the anode pads, the charge integration window of the preamplifiers, and the processing applied by the amplifiers produces a characteristic pulse shape with a width of approximately $1 \mu\text{s}$.

The signals are recorded by a set of custom waveform digitizers (WFDs) built by Boston University for the MuLan experiment. These WFDs have an 8-bit resolution and a 25 MHz sampling rate. In most cases a high gain readout is used which records the muon signals with

optimal resolution. However, fusion events can produce much larger signals so the shaping amplifier output is split to create a low gain signal with one tenth the amplitude, which is digitized separately.

5.3 Detector Response

The rate of energy loss, or stopping power, of particle tracks in a material is given by the Bethe-Bloch equation

$$-\left\langle \frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \left(\frac{2m_2 c^2 \beta^2 \gamma^2 T_{max}}{I^2} \right) - \beta^2 \right], \quad (5.4)$$

where Z and β are the charge and velocity of the incident particle, γ is the Lorentz factor, I is the mean excitation energy of the target gas. T_{max} is the maximum allowed energy transfer in a single collision for a particle of mass M , given by

$$T_{max} = \frac{2m_e c^2 \beta^2 \gamma^2}{1 + 2\gamma m_e / M + (m_e / M)^2}. \quad (5.5)$$

Finally, K is defined as

$$K = 4\pi N_A r_e^2 m_e c^2, \quad (5.6)$$

with N_A being Avogadro's number and r_e being the classical electron radius. Integrating equation 5.4 produces a range-energy relation that can be used to estimate the maximum range of a particle with a given starting energy.

The energy deposited by the particle is not the same as the energy observed at the anode, as some charge is lost due to recombination effects. Recombination mainly occurs when there are many electrons and ions in close proximity, and so depends on the particle's energy loss rate $\frac{dE}{dx}$ and the angle φ between the track and the electric field. The observed energy is obtained by multiplying the true energy by a saturation factor S representing the amount of charge remaining after recombination, according to the formula

$$S = 0.82 \frac{\varphi'}{\varphi' + k \frac{d^2 E}{dx^2}}. \quad (5.7)$$

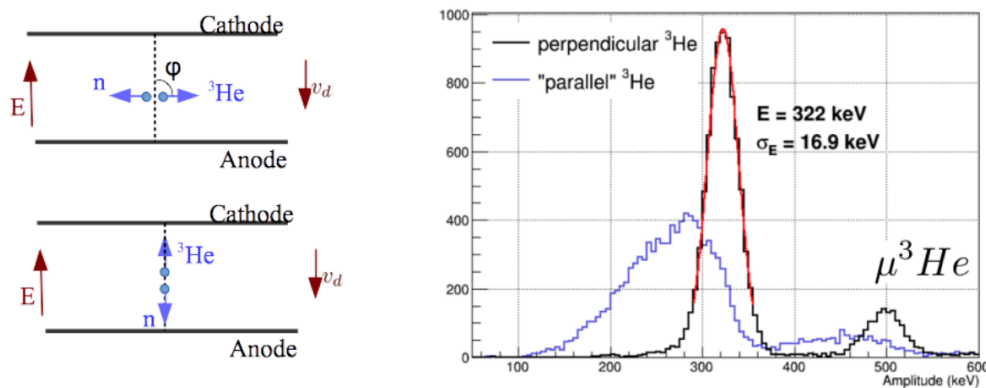


Figure 5.3: Diagram of ${}^3\text{He}$ fusions oriented perpendicular or parallel to the drift field (left), and simulated energies produced by these configurations (right).

Here φ' depends on φ while k is a constant determined by the gas density ϕ and the electric field strength, labeled \mathcal{E} here to distinguish it from the energy:

$$\varphi' = 0.158 + \sin(\varphi) \quad (5.8)$$

$$k = 53.5 \frac{4\varphi^2}{\mathcal{E}} = 0.0849. \quad (5.9)$$

The functional form of this model and the specific numerical values involved have all been carefully tuned by Duncan Prindle to achieve the best match to the muon, ${}^3\text{He}$, and p-t fusion energies observed in real data. Table 5.2 lists the energy and range obtained via SRIM¹ calculations for various particles, as well as the observed energy peaks including the effects of saturation.

The neutron scattering case is special in that it does not leave a continuous track, instead depositing energy at a single point via binary scattering with the nucleus of an atom in the gas. These are elastic collisions, so in the center-of-mass frame the impacted nucleus may recoil at any angle Θ from the incoming neutron vector. For a given nuclear mass A , the

¹The Stopping Range of Ions in Matter software, available at www.srim.org

recoil energy E_R is then

$$E_R = \frac{2A}{(1+A)^2}(1 - \cos \Theta)E_n. \quad (5.10)$$

In this case A is 2 for deuterium, so the prefactor becomes $4/9$. A maximum of $8/9$ of the neutron energy may thus be deposited, which becomes about 2 MeV including the detector response. However, this is only the case for perfect backwards scattering, and scattering at an arbitrary angle will typically deposit much less energy. In the case of very small deflection the energy deposition can be arbitrarily low.

Particle	E (MeV)	E_{obs} (MeV)	R (mm)
μ	3.3	3.0	80.8
3He	0.82	0.4	0.292
$\mu{}^3He$	0.8	0.64	0.653
t	1.01	0.8	0.935
p	3.02	2.7	13.2
n	2.45	0-2	long

Table 5.2: Energies and ranges for fusion products in the TPC. The observed energy (column 2) is reduced due to recombination of drifting charge in the gas. Ranges are obtained via SRIM calculations in D_2 at the target of 6.5% LH2 density.

5.4 Energy Signatures

5.4.1 Row Energies

Muon tracks are approximately parallel to the Z axis and span many TPC pad rows, producing a cluster of roughly coincident pulses. The details of how these clusters are identified will be discussed in chapter 8, but in most cases they are fairly unambiguous. The energy on each row in the track follows a characteristic Bragg curve with a large amount of energy

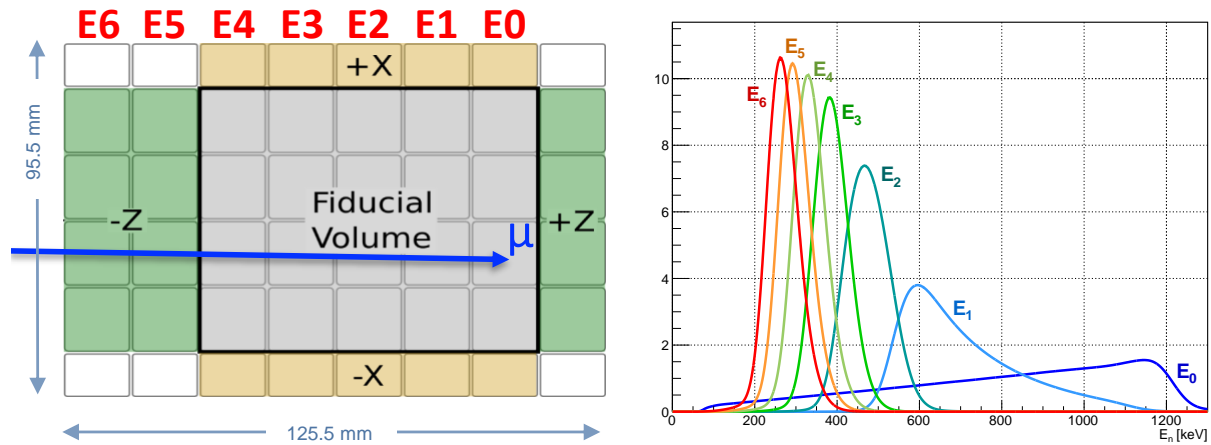


Figure 5.4: Left: Diagram of a typical muon track, with the row energy naming convention indicated. Right: Energy distributions for each pad row. E_0 and E_1 vary greatly depending on the stop depth within the final pad. The upstream energies starting with E_2 form more well defined peaks, as random variations in energy deposition on each pad become more significant than the changes in muon dE/dx with distance.

being deposited near the stop position and less on further upstream rows. The row energies are therefore labeled in reverse order, starting E_0 indicating the final non-zero row energy. Figure 5.4 shows the energy distributions seen on each of these rows up to E_6 .

E_0 and E_1 can take on a wide range of values depending on the stop depth. E_2 and higher are increasingly sharply peaked and also start to blend together, as the changes in stop position become less important than the detector energy resolution and random noise effects. E_0 and E_1 contain a great deal of information but are not particularly informative in isolation, so instead we often plot them together in 2D scatter plots. As shown in figure 5.5, these plots produce distinct bands where the position along the band is determined by the depth of the stop within the final pad row.




Figure 5.5: Plot of E_1 Vs E_0 for normal muon stops. The position along the band is determined by the stop depth $\{Z\}$ within the stop pad, as indicated by the white arrow. The insets illustrate how the stop depth shifts the alignment of the Bragg curve with the pads, thereby changing the relative sizes of E_0 (blue) and E_1 (orange).

5.5 Fusion Products

Fusion reactions deposit additional energy in the TPC, typically on the stop pad and overlapping the original muon track. Here we present two standard measures of the track energy deposition that are used to identify such events.

5.5.1 S-Energy

The 2D E_1 vs E_0 plots above are very informative and will be investigated in more detail in chapter 9, but can be unwieldy. In many cases it is helpful to have a single parameter quantifying the total energy of the muon stop, so that additional energy depositions from fusion products may be identified. One way to achieve this is to approximate the muon stop band as a diagonal line with slope $-1/2$, and define the stop energy (S-Energy or E_S) as the linear combination

$$E_S \equiv E_0 + 2E_1. \quad (5.11)$$

This parameter will be approximately constant for normal muon stops, while ^3He fusions or other particles confined to the stop pad will add directly to the S-Energy. The resulting S-Energy distribution is plotted in figure 5.6, with the sharp muon peak providing a useful calibration point and consistency check for the TPC energy.




Figure 5.6: S-Energy (left) and E_0 (right) distributions for R2015 fiducial stops. The peaks have been labeled with the particles that produce them, although the low-energy electron shoulder also has a significant contribution from muons that decay in flight. Protons from p+t fusions add to both E_0 and E_1 , producing a smeared peak.

5.5.2 E_9 Energy

Not to be confused with the row energies E_0 through E_7 , the E_9 energy refers to the sum of energies on a 3x3 square of pads centered on the estimated muon stop position. E_9 is primarily designed to identify p+t fusion events, which can escape the stop pad and deposit energy on the adjacent pads. Including all nine pads around the stop position means that in most cases the full fusion energy will be captured, even in the event of typical tracking errors. This may be seen in figure 5.6, where the muon peak is broadened in the the E_9 plot but p+t fusions are now much less distorted. The low-energy tail of the p+t peaks are caused by recombination effects for vertically-oriented proton tracks.

We further distinguish two different definitions of E_9 , one of which only includes the pulses in coincidence with the muon track while the other includes all pulses in the event window regardless of their times. The former is referred to as "Cluster E_9 ", as it is based on the pulse clustering and intended to identify events that are problematic for the TPC pulse clustering and tracking systems. The latter is simply called " E_9 " and selects fusion events in a more time-independent way, useful for studies looking at the properties of the fusions themselves rather than issues related to tracking. The details of the clustering and tracking systems will be discussed in upcoming chapters.

5.6 *Electron Tracks*

Finally, Michel electrons produced when the muon decays may also deposit energy in the TPC. As shown in figure 5.7 the energy deposited by the electron is typically very small, below 50 keV. However, there is a long tail extending to higher energies.

The electron deposits energy by scattering on atoms along the track and ionizing them, and the energy of the resulting ionization electrons follows the Landau distribution. For the most part the energy deposition is well modeled by a continuous dE/dx approximation, but the rare high-energy scattered electrons can produce a localized energy deposition that breaks this approximation. This is typically solved by using an energy cutoff to isolate the

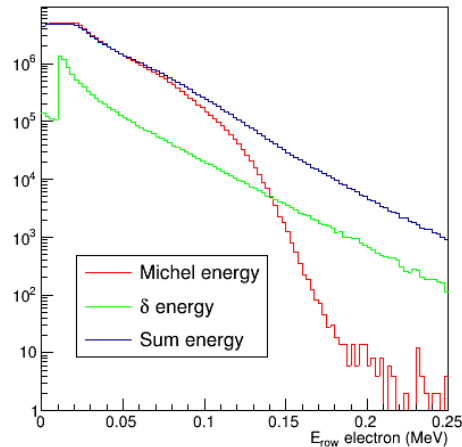


Figure 5.7: Simulated electron energy deposition on the muon stop row, divided into an approximately continuous Michel energy plus high-energy delta electrons.

highest-energy scattered electrons, which are labeled as delta electrons. The energy deposition may then be well modeled by a combination of a continuous dE/dx plus these high-energy deltas. These events are often seen in cloud chambers as short tracks emitted perpendicularly from a longer primary track. The exact cutoff energy is rather arbitrary, and is set at 10 keV for the purposes of modeling the energy deposition in the MuSun TPC. A similar model also applies to the tracks produced by muons and fusion products, but in those cases the much higher dE/dx values make these deviations from the continuous approximation less significant.

5.7 Summary

The design and operational principles of the MuSun TPC have been described in detail, including several important physics constraints that must be accounted for. The detector response and several standard measures of the event energy were then introduced. A more detailed discussion of the ways this information is used in the tracking system will be given in chapter 9, and it is also important for the systematics studies related to fusion products and electron energy deposition in the TPC which are covered in chapters 13 and 14 respectively.

Chapter 6

DATA COLLECTION

The MuSun experiment used the Swiss muon source at the Paul Scherrer Institute (PSI), located in Villigen, Switzerland. This valuable resource is shared among several other muon experiments, with beam time allocated in blocks of several weeks or months. As such, experiments are structured around brief periods of data collection, known as runs, while the remainder of the year is spent focusing on data analysis and technology development.

For the MuSun experiment individual runs generally received about 10 weeks of beam time, with additional time required before and after to set up and then shut down the experiment. The accelerator runs continuously including nights and weekends, apart from some shutdowns for maintenance, and so the experiment was also operated without breaks to fully utilize the beam allocation. During normal operations work was divided into three 8-hour shifts, although hardware problems would interrupt this schedule as everyone hurried to get the experiment back online. Less critical maintenance and hardware upgrades were planned to coincide with scheduled beam downtime and accelerator maintenance periods.

The MuSun experiment included a total of nine runs, beginning in 2008 and concluding in 2016. After a series of development and testing runs, the two main production runs in 2014 and 2015 collected a total of $1.3 \times 10^{10} \mu^-$ events, exceeding the statistics target. They also collected $2 \times 10^9 \mu^+$ events, and along with the final run in 2016 they produced a variety of specialized datasets for studying potential systematic errors. Additional simulated datasets were also generated for software development and testing purposes, and the simulation procedure will be summarized briefly here. This author joined the experiment prior to the eighth run in 2015 and participated in the data collection and analyses for both the 2015 and 2016 data sets.

6.1 Early Runs

The first three MuSun runs were focused on adapting hardware from the previous MuCap and MuLan experiments and on engineering and development of the new MuSun TPC. These led up to R2011, which was the first full-scale test run and collected $5 \times 10^9 \mu^-$ events and $0.5 \times 10^9 \mu^+$ events. This data is not used in the final analysis, but was used to develop and test the analysis software and procedures. Challenges faced during R2011 spurred further developments in 2012 and 2013 which resulted in the final design of the experiment as described in chapters 4 and 5.

The only dataset from these early runs that is directly relevant to the final analysis is from a failure of the gas purification system in R2013. While the filters were offline a high-statistics dataset was collected with approximately 20 ppb nitrogen concentration, and additional tests performed with approximately 0.5 ppm and 2 ppm impurity at a higher operating temperature. The corresponding distortion of the observed muon lifetime was large enough to be precisely measured, providing a valuable calibration point for the impurity analysis.

6.2 R2014




Figure 6.1: Accumulated statistics for the R2014 (left) and R2015 (right) data collection periods. The shaded vertical bands indicate periods of accelerator downtime.

R2014 was the first full production run using the final TPC design as described in chapter 5. Data was collected for ten weeks, yielding a total of $6 \times 10^9 \mu^-$ events as well as $1 \times 10^9 \mu^+$ events. Figure 6.1 shows the steady accumulation of statistics over the course of the run.

Dataset	Files	Events	Conditions
ds01	77628 - 79267	1.6×10^8	no μSCa
ds02	79347 - 79666	7.2×10^7	μSCa connected
ds03	78124 - 81017	7.7×10^7	variable file size
ds04	79689 - 81613	4.2×10^8	μSCa misaligned
ds05	81614 - 82018	2.3×10^8	μSCa misaligned
ds06	82038 - 82163	7.5×10^7	μCL added
ds07	82304 - 82861	1.8×10^8	μSCa realigned
ds08	82865 - 83050	5.2×10^7	unstable beam
ds09	83617 - 83859	-	low beam
ds10	83898 - 85511	7.8×10^8	collimator removed
ds11	85529 - 85726	1.3×10^8	new collimator
ds12	86197 - 86466	1.9×10^8	μPC wire swapped
ds13	86467 - 87829	9.2×10^8	μCL turned on
ds14	87892 - 88767	5.5×10^8	μPC wire back
ds15	89181 - 90693	1.1×10^8	$\mu+$ production
ds16	90704 - 92225	1.2×10^8	standard production

Table 6.1: R2014 sub-datasets based on running conditions.

Several issues affected the data quality, particularly early in R2014. For approximately the first half of the run the upstream collimator located inside the bending magnet after the separator was misaligned, resulting in high beam backgrounds. After that the collimator was removed completely for a short period, before finally being installed correctly. There were also alignment problems with the μSCa for several weeks at the beginning of the run, although this has been found to have relatively little impact on data quality. Finally, as mentioned

in section 4.2, an incorrectly calibrated upstream magnet produced a sub-optimal beam configuration throughout the entirety of the R2014 and R2015 data collection campaign. This resulted in several attempts throughout both runs to optimize the beam profile and kicker extinction by tuning the downstream magnets. To account for these changes in running conditions and beam configuration, the R2014 dataset is divided into sixteen sub-datasets as shown in table 6.1. The number of events listed indicates the number of muon-electron pairs remaining after all selection cuts are applied.

These small datasets are mainly useful for precisely adjusting calibrations, time offsets, and analysis cuts which may depend on the exact conditions. However, they lack the statistics to be useful for analyzing the muon lifetime, and most of these small changes should have little effect on the final results. Therefore, for most analyses these are combined into six larger datasets based on the most important running conditions, as shown in table 6.2.

Dataset	Files	Events	Conditions
DSM1	ds01-ds03, ds07-ds09	5.1×10^8	collimator misaligned
DSM2	ds04-ds06	8.3×10^8	collimator & μ SCa misaligned
DSM3	ds10	8.3×10^8	no collimator
DSM4	ds11-ds13	1.3×10^9	new collimator
DSM5	ds14, ds16	1.8×10^9	new collimator
DSP1	ds15	1.1×10^9	μ^+ production

Table 6.2: R2014 larger sub-datasets for high-statistics lifetime fits.

6.3 R2015

The eighth MuSun run in the Summer of 2015 was the second of the main production runs. With the experiment already performing well for the R2014 production, the purpose of this run was primarily just to continue collecting data. Data was once again collected for ten weeks, resulting in 7×10^9 μ^- events and another 1×10^9 μ^+ events as shown in figure 6.1.

6.3.1 Changes Since R2014




Figure 6.2: The new R2015 μ SCa, showing the rows of SiPM detectors along the sides.

There were relatively few hardware changes between R2014 and R2015 to maintain consistency between the two production datasets, but some auxiliary systems were upgraded. The author's first project with the MuSun group was assisting in development of the improved μ SCa detector mentioned in section 4.3. The detector itself is shown in figure 6.2, and it was connected to a control box containing a BeagleBone microcomputer programmed by the author. Seven of the eight neutron detectors were upgraded to the larger DEMON detectors discussed in section 4.6, although issues with the digitizers unfortunately limited the effectiveness of this upgrade. For some datasets an additional pulser was added to inject test signals into the TPC, which was used to study a WFD error mode described in appendix E. Finally, while the μ CL was only enabled intermittently in R2014 it was included for nearly all of R2015, albeit with the frequency reduced from 2.5 kHz to 1 kHz resulting in only a modest increase in total statistics.

Dataset	Files	Events	Conditions
ds01	94290 - 95086	3.7×10^8	
ds02	95087 - 95504	2.1×10^8	replaced a TPC WFD
ds03	95505 - 96324	2.1×10^8	μCL turned on
ds04	96334 - 97204	4.0×10^8	file size increased
ds05	97405 - 97654	1.3×10^8	μCL reduced to 1 kHz
ds06	98000 - 98041	1.7×10^7	crate 6 rewired
ds07	98057 - 98752	3.3×10^8	new beam tune
ds08	98789 - 98825	1.8×10^7	new beam tune
ds09	98830 - 98957	3.7×10^7	file size reduced
ds10	99370 - 100701	4.2×10^8	new neutron FADC, unique pulser
ds11	100702 - 101371	2.0×10^8	ePC2 voltage adjusted
ds12	101372 - 103603	6.6×10^8	unique pulser removed
ds13	103604 - 103631	1.2×10^7	reduced data blocks per file
ds14	104212 - 104624	1.8×10^8	preamp pulser removed
ds15	104641 - 105675	4.5×10^8	new preamp pulser
ds16	105678 - 107652	7.7×10^8	improved unique pulser added
ds17	107653 - 109715	8.8×10^8	separator adjusted
ds18	109716 - 111283	6.8×10^8	standard production
ds19	113072 - 114674	5.2×10^8	standard production
ds20	114767 - 115796	3.7×10^8	$\mu+$ production
ds21	115197 - 116069	1.1×10^8	new beam tune
ds22	116508 - 118657	8.3×10^8	TPC thresholds adjusted
ds23	118702 - 119984	5.3×10^8	standard production
ds24	119985 - 120119	4.5×10^7	tests with μSR magnet
ds25	120120 - 123648	1.4×10^9	standard production
ds26	123649 - 124608	3.4×10^8	standard production

Table 6.3: R2015 sub-datasets based on running conditions.

6.3.2 Datasets

Like the R2014 data, the R2015 dataset is subdivided into a number of smaller datasets based on the run conditions, listed in table 6.3. The running conditions had stabilized by this point, so there are no major changes like the collimator issues in R2014. Instead, these datasets are mainly distinguished by minor optimizations of the beam tune or adjustments to detector settings. Again, these smaller datasets are mainly used for calibration purposes where such fine details are important to keep track of.

These small datasets are once again grouped into a few large datasets for use in high-statistics lifetime fits and consistency checks, continuing the naming convention from R2014 as shown in table 6.4. As the running conditions were broadly similar across R2015, the boundaries between these datasets are rather arbitrary. DSM7 and DSM8 in particular have nearly identical operating conditions, and were chosen purely to make the four datasets roughly equal in size.

Dataset	Combination	Events	Conditions
DSM6	ds01-ds09	1.7×10^9	before first accelerator maintenance
DSM7	ds10-ds16	2.3×10^9	
DSM8	ds16-ds19	2.5×10^9	
DSM9	ds23-ds26	2.3×10^9	after $\mu+$ production
DSP2	ds20-ds22	1.3×10^9	$\mu+$ production

Table 6.4: R2015 larger sub-datasets for high-statistics lifetime fits.

6.4 R2016

The ninth and final run took place in Autumn 2016. With the main production runs complete and the full statistics collected, this run instead focused entirely on a variety of systematic studies designed to address questions that had arisen during analysis of the production data.

6.4.1 Neutron Calibration Targets

For the wall stop analysis discussed in chapter 11 it was useful to know the sensitivity of the neutron detectors to muon stops in various materials. A special detector setup was designed with a sled that sits inside ePC1 in place of the TPC, holding a sample in line with the beam. Before and after the target sample there are scintillating detectors read out by rows of SiPMs on either side, as shown in figure 6.3. The front detector is used to select events where the muon definitely hit the target, while the rear detector is for vetoing punch-through muons when using a thin sample. Finally, a plastic collimator in front of the setup blocks any muons that would impact the detectors rather than the sample itself.

Files	Events	Target
128969-129010	5.4×10^8	silver foil
129032-129047	7.2×10^8	tungsten plate
129048-129132	7.8×10^8	MACOR block
129142-129153	1.2×10^9	anode plane

Table 6.5: R2016 neutron calibration target datasets.

Several samples were used as targets, shown in table 6.5. Silver, tungsten, and Macor samples were chosen because these are the main materials used in the TPC, and should be representative of most potential wall stops. Some data was also taken with a spare copy of the anode plane as the target, to study the response from a composite material.

6.4.2 Impurity Injection

Although some data with high impurity levels was already obtained during R2013, it is not directly applicable to the production datasets since they use a different TPC. One of the major goals for R2016 was therefore to obtain new high impurity data and use this to calibrate the in-situ impurity measurement method. Gas mixtures with 13 ppb and 30 ppb

nitrogen impurity were prepared, and table 6.6 shows the datasets where these gas supplies were injected into the system. It is important to note that it took some time after the start of impurity injection for the gas in the TPC to equilibrate. The 13 ppb mixture took more than half the dataset to stabilize, while for the 30 ppb mixture the impurity level in the TPC was still rising when the gas supply ran out.

Files	Events	N2 Impurity
129261-130102	5.4×10^8	clean
130103-130652	7.2×10^8	13 ppb
130653-131259	7.8×10^8	30 ppb

Table 6.6: R2016 impurity calibration datasets. The listed impurity level is for the input gas, and does not necessarily reflect the levels in the TPC itself.

6.4.3 Beam Background Tests

Because of the miscalibrated quadrupole magnet upstream in the beamline, an optimal beam tune could not be obtained for either R2014 or R2015 and both datasets have a larger beam background than expected. To better understand the effects of this background two additional detectors were added to the entrance detector stack, as shown in figure 6.3. The beam electron scintillator (beSC) was similar to the R2015 μ SCa but with no hole in the scintillator, while the beam electron silicon (beSi) detector used a silicon chip as the detector itself. During the beam electron measurements a 1 mm lead plate was inserted between the normal entrance detectors and these new electron detectors, blocking the muons but allowing a significant portion of the electrons to pass through.




Figure 6.3: Left: The neutron calibration setup, with a Macor block as the target. Right: The entrance detector stack with additional beam electron detectors added. An optional lead plate can slot in after the μPC to block the muon beam.

6.5 Simulation

In addition to the real data collected at PSI, several additional simulated datasets were generated. This data was used to inform the development of the analysis software, to test the performance of event reconstruction against known truth information, and to verify that the analysis correctly reproduces the input capture rate. The simulation software is maintained and operated by Duncan Prindle, and will only be described briefly here.

The simulated data is generated in two steps. The first step is a Monte Carlo (MC) simulation based on the GEANT4 (GEometry ANd Tracking) toolkit developed at CERN for particle physics simulations. This program models the MuSun detector geometry and simulates the trajectories and interactions of the incoming muons, and of any additional particles produced by said interactions. The system also records the amount of energy deposited in the sensitive detector volumes. The output of this first physics simulation is

known as the Monte Carlo truth information, and a copy of it is retained throughout the analysis to allow detailed testing of individual analysis components.

The second step of the simulation procedure simulates the detector response, distorting the truth information to account for effects such as signal loss due to recombination in the TPC. The response code then applies a model of the electronics and digitizer systems, storing the output in data files with the same format as those produced in the experiment. These files may then be processed by the same analysis software as the real data. The simulation parameters and response model are carefully tuned to closely match the experimental conditions of the production datasets.

6.6 Summary

The MuSun experiment started with a series of engineering and testing runs, leading up to two successful production runs in 2014 and 2015. The R2014 and R2015 runs collected 6×10^9 and 7×10^9 μ^- events respectively for a total of 1.3×10^{10} μ^- events, exceeding the target statistics. Each production run also collected 1×10^9 μ^+ events for comparison to the known free muon decay rate. A final followup run in 2016 performed a series of dedicated systematics tests to help constrain various sources of error in the main analysis. Additional simulated data was also produced for testing purposes, and is closely matched to the conditions of the production datasets. The author assisted with the data collection for both R2015 and R2016, and has been primarily responsible for the in-depth analysis of the R2015 dataset which will be described in subsequent chapters.

Chapter 7

LOW-LEVEL ANALYSIS SOFTWARE

The MuSun analysis framework consists of three stages, each building increasingly sophisticated representations of the data. This chapter will give an overview of this framework and the process of running an analysis pass. It will then focus on the first analysis stage, named "MU", which primarily performs low-level data processing to convert raw data into detector response objects. Chapter 8 will then cover the second and third analysis stages that perform more high-level event reconstruction and histogramming.

7.1 Analysis Framework

The stages of the MuSun analysis framework are outlined in figure 7.1. Stage 1 turns the large blocks of raw data from the experiment into a series of muon-centered events consisting of objects representing particle interactions. Stage 2 then combines these individual interactions into particle track objects, and produces simplified representations of the event properties. Finally, stage 3 applies various event selection cuts and fills output histograms from the processed event data, for use in associated analysis projects. Each stage also outputs a set of calibration and diagnostic histograms which are useful for ensuring proper operation of the analysis and updating parameters for subsequent analysis passes.

The primary benefit of the staged analysis approach is increased efficiency. The low-level data processing tasks performed in stage 1 are highly unlikely to change significantly as the result of some new analysis insight. The track reconstruction performed at stage 2 is more likely to need adjustment, while the specific histograms required for an analysis often change multiple times per day as new ideas are explored. Because the stages are independent the stage 2 analysis may be rerun to apply updated tracking algorithms to the output of an

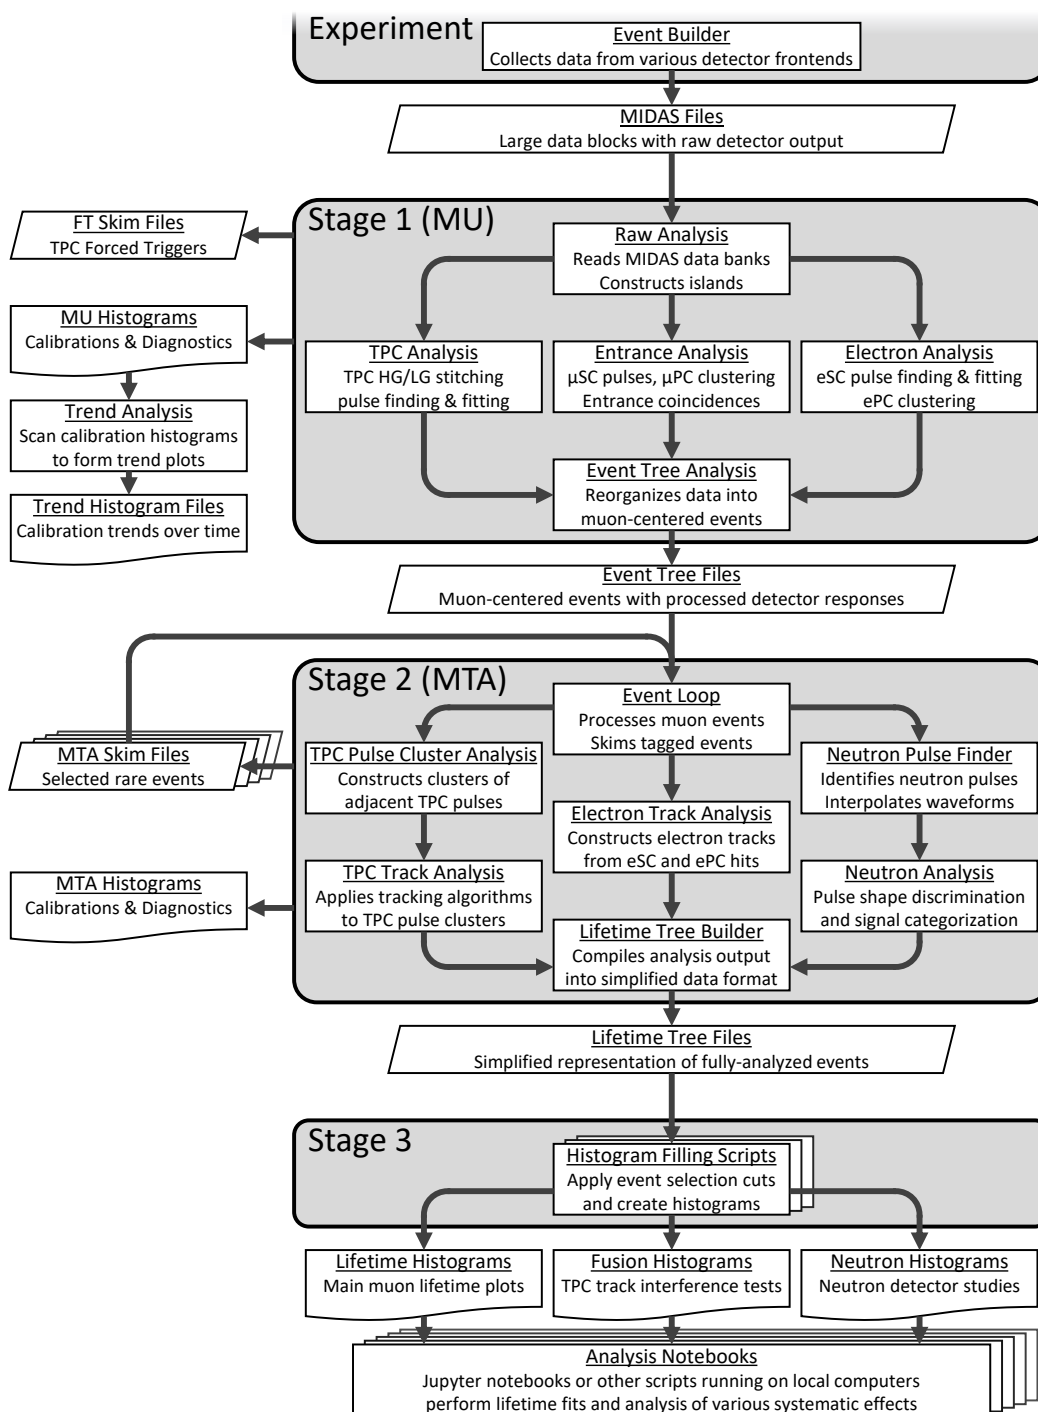


Figure 7.1: Flowchart of the MuSun analysis chain.

old stage 1 pass, and the stage 3 analysis is rerun frequently to produce new results while performing minimal additional processing.

The analysis is run on the Stampede2 supercomputing cluster provided by the Texas Advanced Computing Center (TACC), utilizing a subsystem with 1,736 48-core Intel Skylake compute nodes. Stampede2 users are granted allocations in the form of service units (SUs), representing one node-hour of compute. A full analysis pass of the R2014 and R2015 production data requires 2000 SUs for stage 1, 1000 SUs for stage 2, and roughly 30 SUs for stage 3. This again demonstrates the importance of the staged analysis framework, as for the price of a single full pass it is possible to run 3 passes starting at stage 2 or to rerun stage 3 roughly 100 times.

The tradeoff for this efficiency and flexibility is the requirement to separately store all of the intermediate files. The MuSun experiment maintains 180 TB of raw experimental data, 90 TB of which is from the two production runs. The output of stage 1 will be approximately 150% the size of the original input, while the stage 2 output will be 50%, for a total storage cost of twice the raw data. It is helpful to retain data from some old analysis passes for testing and comparison purposes, but these each require an additional 180 TB of storage space. MuSun has a 700 TB allocation on the Ranch tape archive, also provided by TACC, which is used to save these large amounts of data. See appendix C for a more detailed explanation of how MuSun utilizes computer resources.

The main analysis chain is primarily programmed in C++, using the ROOT data analysis framework developed at CERN. ROOT is designed for efficient processing of large datasets, which is critical for analyzing MuSun's roughly 90 TB of raw production data, and also provides histogramming and other data visualization and fitting tools. The final analysis of the stage 3 output histograms is more typically programmed in python, often in the form of a Jupyter notebook, using the PyROOT interface to interact with the histograms and other ROOT objects.

7.2 *MU*

MU is based on the MIDAS framework and built in a modular fashion, with dedicated analysis modules for each detector system. It takes the raw data files as input, which are composed of a series of 140 ms data blocks as described in section 4.8. These data blocks are structured as a series of so called data banks, each containing a long list of all the data packets sent from a given digitizer during that block. MU is responsible for performing low level processing to convert this raw data into C++ objects representing particle detections. The processed data is then reorganized and output as a series of individual events, to facilitate efficient analysis in the subsequent stages.

Because MIDAS is designed for data acquisition and processing in a real-time environment, it is a bit overcomplicated for simple offline analysis. In particular it makes heavy use of shared memory to allow communication between multiple active frontends, which is difficult to work with and has led to obscure problems when running many copies in parallel such as exceeding the system limit for semaphores¹ on Stampede. MU has also accumulated a significant amount of legacy code over its development, sometimes making it unclear which modules are actually relevant to the main analysis. Even setting aside these issues, the low level data handled by MU is simply difficult to work with. For all of these reasons, active development has focused primarily on the later analysis stages with additions to MU only occurring when strictly necessary.

7.3 *TPC WFD Pulses*

The WFDs send their data as 96-byte packets consisting of 4 zero bytes, a 32-bit timestamp, and 88 8-bit samples. The digitizers are configured to record the waveform when the signal passes above a threshold value, but 25 presamples are saved and included at the beginning of each waveform. This only leaves 63 samples for the waveform of the pulse itself, and pulses

¹Semaphores are a mechanism for thread or process synchronization, typically used to prevent concurrent access attempts for a given shared memory resource.

are often longer than that. If the signal is still above threshold after the last sample in the packet, this will result in a second data packet immediately following the first. Such events are easily identified by their timestamps being 88 ticks apart, and are merged together to form arbitrarily long waveforms as long as the signal remains above threshold. The resulting waveform objects are known as islands, and will need additional processing to identify and fit the pulses present in the waveform signal.

One complication with the WFD signals is that, due to a latching mechanism, the island times may be off by exactly four samples. This four clock tick ambiguity (4cta) will be discussed in more detail in the next section in the context of the eSC WFDs. For the TPCs a pulser injects a test signal into the amplifier chain every block to ensure the system is operating correctly. This sync pulse affects all pads and is injected at a known time, so it is used to align the clock signals.

7.3.1 Pedestal Calibration

The WFD waveforms generally have a pedestal value around 20 counts, to allow for undershoots which often follow larger pulses. The presamples recorded before each pulse may be used to estimate the pedestal, but this method is statistics limited. We instead use the presamples from all pulses in a data file to assign a set of precise fixed pedestals for each digitizer, improving energy resolution.

To characterize the baseline of each WFD channel a forced trigger (FT) signal causes the WFDs to record continuously for the entire $25 \mu\text{s}$ event window following a muon entrance. The drift time for ionization electrons in the TPC is at most $16 \mu\text{s}$, so the last $9 \mu\text{s}$ of each FT event is usually signal free. A Gaussian fit to the sample distribution in this region is used to extract fixed pedestal values for each TPC WFD channel. These values are stored in the MySQL database for each input file, and are retrieved by the stage 1 analysis during subsequent passes.

7.3.2 Dual Gain Processing

In the case of the TPC signals, there is an additional step that needs to be taken after the islands are created. Recall that each TPC pad is digitized by two separate WFDs, with one signal amplified by an additional factor of ten. The high gain (HG) signals are optimized for muon pulses, with dynamic ranges of approximately 1 MeV. However, some events can produce much larger pulses, producing clipped HG waveforms and requiring the low gain (LG) to see the full peak. Because most of these large pulses are caused by muon-catalyzed fusions, any bias caused by the waveform clipping may result in fusion interference effects.




Figure 7.2: Example of a clipped HG waveform, shown in red, and the corrected waveform after using the LG waveform to replace the clipped samples.

Before pulse finding can proceed, the HG and LG signals must be matched together and recombined into a single waveform in a process known as stitching. The basic idea of the stitching procedure is simple: for any waveform in which the HG signal hits the maximum value of 255 the LG signal with the same timestamp is retrieved and used to replace these maxed-out samples after scaling by the appropriate gain factor. An example of a clipped HG waveform and the final merged waveform is shown in figure 7.2.

In practice the stitching process is more complicated, because while the gain ratio between the HG and LG channels is nominally a factor of ten this is not exact. The true gain ratios must be calibrated individually for each pad, and can also vary over the course of a run. During normal operation the gains are usually constant, but they occasionally jump to slightly different values as a result of some adjustment to the experimental hardware. A different set of calibrations is therefore required for each of the small datasets listed in table 6.3. To determine the correct gain ratios, the HG and LG amplitudes may be compared for non-overloaded pulses where the peak is visible in both waveforms. Plotting the maximum amplitude of the LG vs the HG pulses shows a linear relationship which can be fitted to extract the gain calibration.

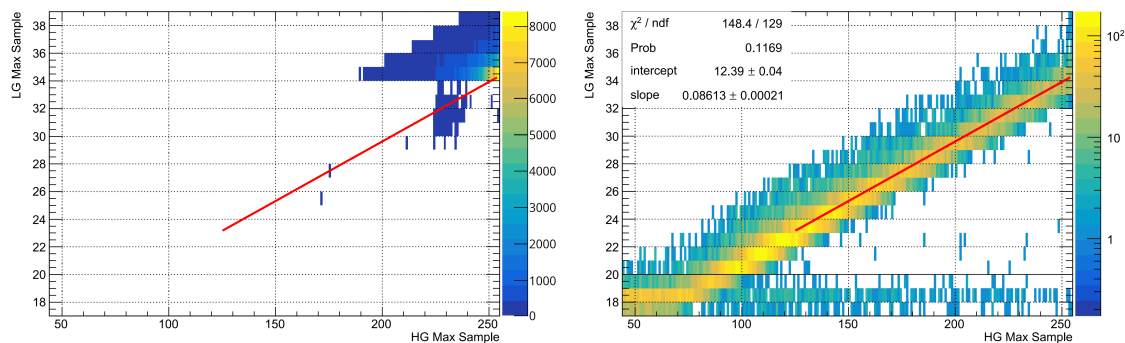


Figure 7.3: TPC relative gain calibration using normal triggers (left) or forced trigger waveforms (right). The data is for pad 44 from part of R2015 ds01, which had a very high LG threshold resulting in few usable pulses during normal operation.

This method requires pulses to be visible in both gain channels, but many small pulses fall below the LG WFD trigger threshold and are not recorded. A high LG threshold can severely limit the available fit range for some pads, rendering the method nearly unusable. A solution is to use FT waveforms, which include all pulses regardless of amplitude but only for every thousandth event. Summing an entire dataset of several hundred files together can compensate for the low FT statistics and allow reasonable fits as shown in figure 7.3,

although the calibration accuracy will vary depending on the size of the dataset. The fits to normal triggered pulses are used to create fine-binned trend plots to identify stable periods that are safe to sum together and use for the more precise FT calibrations.

Finally, if the waveform only briefly exceeds the maximum for the HG WFDs the stitching does a poor job of estimating the pulse amplitude. Because the LG signal jumps by steps of ~ 10 HG ADC counts it may be off by ± 5 compared to the true signal, which is significant for pulses with amplitudes of ~ 250 . To avoid artifacts near the peaks of these slightly clipped pulses, a better estimate of the missing samples is obtained by fitting a template function to the surrounding HG samples. If there is a LG waveform available then the algorithm will switch to standard stitching for stretches of more than 10 saturated samples, which still yields better results for large pulses.

7.3.3 Pulse Processing

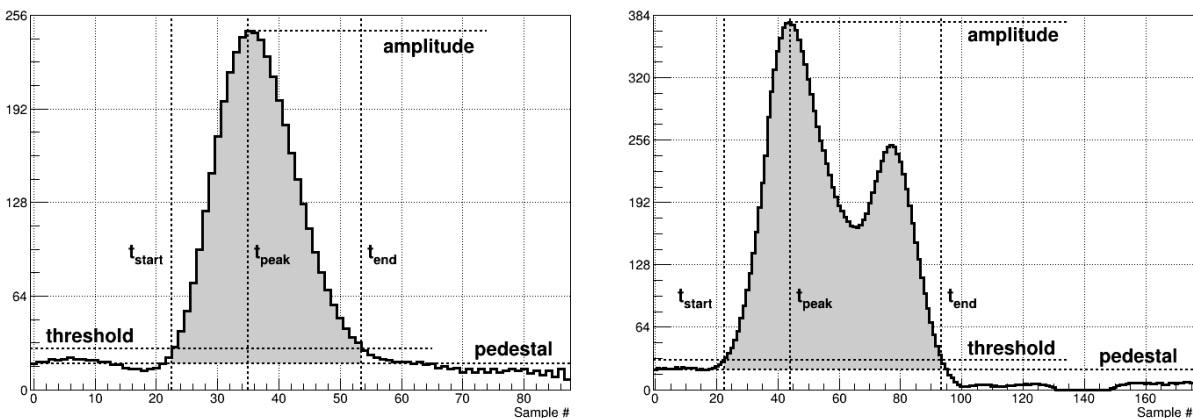


Figure 7.4: Application of the TOT pulse finder to muon (left) and p+t fusion (right) waveforms. The first and last samples to exceed a 10 count threshold set the start and stop times, respectively. The energy is determined by the integral between these samples, shaded gray. Note the undershoot in the fusion case.

With the stitched islands created, the next step is to identify individual pulses in the waveforms. A time over threshold (TOT) pulse finding algorithm was developed by Michael Murray and is explained in detail in his thesis[54]. This method looks for stretches of samples more than 10 counts above the pedestal. These sample stretches are considered legitimate pulses if the maximum sample in the stretch is above 17 counts, to avoid spurious triggers from fluctuations across the threshold. Figure 7.4 shows the application of the TOT algorithm to real waveforms.

Note the large undershoot in the fusion waveform, this can last for several microseconds after a large pulse. The undershoot is not accounted for by the fixed pedestal system, and will reduce the estimated energy of any pulse in this region while preventing small pulses from being detected entirely. This is less of a concern for the smaller muon pulses, and has little effect on tracking since fusions are typically delayed and the $25 \mu\text{s}$ event window provides ample time for the baseline to recover before the next muon entrance.

The TPC uses charge-integrating preamplifiers and shaping amplifiers, so all sufficiently short pulses produce a characteristic waveform shape with amplitude proportional to the energy of the signal. Muon tracks are nearly horizontal and the drifting charges are detected at approximately the same time, so they satisfy this condition. However, for p+t fusions with a vertical proton track the drifting charge can be smeared out over several microseconds, exceeding the integration window of the preamplifiers and producing an extended waveform as seen in figure 7.4. For such pulses the amplitude of the signal is no longer a valid measure of the total ionization energy, and the integral of the pulse must be used instead. For consistency the term energy will always refer to the pulse integral, although for normal pulses it is slightly less accurate than the amplitude due to losses where the pulse tails fall below threshold. Pulse integrals are converted to physically meaningful energy values using fixed calibration factors for each pad, which are derived as part of the stage 2 analysis.

In addition to the energy several other properties of the pulses are stored. The center time is determined using the maximum sample in the pulse, while the start and end times are determined by the first and last samples in the stretch. The amplitude is also saved for use

in specialized analyses where extended pulses are rejected. The resulting pulse objects are used for all subsequent analysis steps, while raw waveforms are only retained for extended or otherwise unusual pulses and discarded for pulses that match the normal signal shape. The forced trigger waveforms are also separated out and written to a special FT skim file for easy access.

Pulses are assigned coordinates in order to reconstruct the particle tracks in 3D. The X and Z coordinates are given directly by the coordinates of the triggered anode pad. The Y coordinate is derived from the drift time of the pulse, defined as the delay relative to the muon entrance time

$$t_{drift} \equiv t_{pulse} - t_{\mu SC}. \quad (7.1)$$

The drift time is converted into a Y value using the drift velocity of 5 mm/ μ s. It is important to note that these reconstructed Y coordinates are only valid for the muon track, where all pulses may be assumed to be roughly coincident with the entrance time. For event such as muon-catalyzed fusions where the true pulse time is unknown, the initial position is instead taken to be the muon stop position and the drift time is used to estimate the fusion time.

7.4 eSC WFD Pulses

The eSCs use the same WFD boards to record their waveforms as the TPC, so the process of island creation is similar. However, the eSC pulses are only about 20 ns long so even with a 450 MHz sampling rate only 10 samples are needed to record the pulse shape. The data is therefore recorded in shorter packets of 32 samples including 8 presamples, for a total length of 71 ns. The eSCs also only need to detect electrons which are all approximately MIPs, without the large variations in amplitude seen in the TPC signals. The 8-bit resolution of the WFDs is sufficient for capturing the electron signals, eliminating the need for the dual-gain arrangement and stitching procedures.

7.4.1 *Pulse Processing*

The pulse processing procedure is also simpler for the eSCs than it was for the TPC. It is safe to assume that there is only a single pulse per island, considering there should only be one muon decay in each event. The beam background contributes to a background rate of about 1 kHz across all 16 scintillator panels, so the probability of a background electron coinciding with a standard 71 ns island is on the 10^{-5} level. Furthermore, to eliminate after-pulsing a 100 ns artificial deadtime is applied to the reconstructed eSC pulses, which eliminates any additional pulses on the island on the rare cases where they do occur.

The largest sample in the island is used to determine the amplitude of the pulse. The pulse time is calculated as the point where the leading edge of the waveform exceeds a constant fraction of 1/4 of the pulse amplitude, using linear interpolation between the samples to obtain a more precise estimate. The pulse integral is also computed, using a fixed window from 3 samples before the peak to 5 samples after the peak of the pulse.

7.4.2 *Clock Alignment*

The eSC digitizers suffer from an effect whereby the timing can be shifted by exactly 4 clock ticks (≈ 9 ns) from its nominal value. The offset is consistent for the duration of each data block, but changes randomly between blocks. This effect is known as the four clock tick ambiguity (4cta), and must be corrected on a block-by-block basis to ensure proper clock alignment. The accuracy of this alignment is very important for forming coincidences and for estimating the eSC Z position from time of flight differences, which have a spread of about 7 ns. The eSCs do not have a sync pulse that can be used as a timing reference so the alignment procedure is somewhat complicated, and is described separately in appendix D.

7.5 *ePC Clustering*

The wire chambers use TDCs to record their signals, which produce data banks consisting of a simple list of trigger times for each parameter number. These parameter numbers are

part of a global numbering scheme for the entire experiment. The analysis at MU must use the global wiring map to fetch the relevant data banks and convert them to physically meaningful values such as wire numbers.

The author was responsible for a major rewrite of the ePC tracking software, described in detail in appendix A. In short, the analysis groups hits on nearby wires into clusters, and then finds coincidences between clusters in each of the three wire layers. Because each layer has a different winding angle these wire clusters should only overlap at a single point on the cylinder, producing a 3-fold coincidence. Often one of the expected clusters will be missing, and while 2-fold coincidences can be formed from the remaining clusters they are much less reliable. This typically occurs because of various data quality issues with the ePCs, and the new tracking system was specifically designed to robustly handle these problems.

7.6 Entrance Detector Analysis

Perhaps the most important analysis step carried out at MU is the processing of the muon entrance detector data. Detector objects are created to represent each of the three entrance detectors: the μ SC, μ SCa, and μ PC. These are then combined into muon entrance objects containing all the information about a given muon. Much of the credit for this analysis goes to Rachel Ryan, who developed the current entrance analysis framework[62].

7.6.1 μ SC

The μ SC provides precise timing of the muon entrance, and is used for both the start time of the lifetime measurement and as a reference point for all other times in an event. Because this time measurement is so critical to the experiment, the analog signal from the scintillator is discriminated and digitized by two separate TDC channels. The creation of a good μ SC object requires a coincidence between both TDC channels, if either signal is missing this is considered an error mode.

In addition to the two TDC signals, there is also a third copy of the signal that is sent to a WFD. These operate very similarly to the eSC WFDs, except the waveforms are fitted

with a pulse template function allowing multiple pulses to be identified in a single island. The TDCs are considered the official time, and the WFD times are aligned to the TDCs for each block to correct for the timing ambiguity problem. The main advantages of the WFD pulse information are a much lower amplitude threshold and minimal deadtime due to the multi-pulse fitting procedure. WFD pulses without an associated TDC trigger are not used to create full muon entrance objects, but are useful for pileup protection.

7.6.2 μSCa

The μSCa objects are very similar to the μSC objects. The time is once again taken from the TDC, while the amplitude is extracted from the WFD waveform. The pulse finding is copied from the μSC as well, but because this is only a veto detector it has relatively low count rates and it is not necessary to fit more than a single pulse. Separate WFD waveforms from the left and right SiPM strips are also available, but these are only used for a few diagnostics and the analog sum is used for most purposes.

7.6.3 μPC

Like the ePCs, the μPC X and Y wires must be grouped into clusters. The μPC is much less noisy though, so clusters are formed via a simple three step procedure as shown in figure 7.5. First, the signals are sorted in time order and preliminary clusters are formed from signals within a coincidence window. Second, the preliminary clusters are re-sorted into wire number order and split into groups of adjacent wires. Finally, these are sorted back into time order and the coincidence window is re-checked to correct a rare edge case where two separate clusters could be merged due to an unrelated third cluster.

Originally both the μPC and ePC clustering stopped after step two. The third step was added in response to investigation of the various ePC failure modes, while the ePC clustering itself was completely redesigned. The muon dE/dx is approximately 5 times that of the electron, so the μPC can use significantly higher thresholds than the ePC and is much less sensitive to noise or other data quality issues.




Figure 7.5: Diagram of the μ PC clustering algorithm. Hits are grouped by time, then by space, and then again by time.

The formation of combined XY clusters is also simplified compared to the eSC analysis, with only two planes it is simply a matter of finding coincidences between the X and Y planes. The resulting μ PC clusters provide information about the beam profile, which is useful for tuning the beam during data collection. The μ PC is also useful for pileup protection, as it is sensitive to coincident but spatially separated muons which would be indistinguishable with the μ SC alone. The separate X and Y clusters can be used to provide slightly better pileup protection, but they are much more likely to trigger on noise or beam electrons than the combined XY coincidences resulting. For the early R2014 datasets in particular beam electron triggers kill many otherwise good events, so pileup protection on μ PC singles is disabled except for some specific analyses.

7.6.4 Muon Entrances

Muon entrance objects are typically seeded with μ SC objects. They are then checked for coincidences with μ SCa and μ PC objects, and are also checked against the kicker and μ CL signals to ensure only real muons that triggered the kicker are included in the main analysis. Pileup protection flags for each entrance detector are set if there are any additional signals

detected within the $\pm 25 \mu\text{s}$ pileup window. The times of both the next and previous entrance and kicker signals are also saved, to allow later tests to extend the pileup protection window. A wider $30 \mu\text{s}$ book-ending window is used to reject events near the beginning or end of a data block. Finally, a 29 ns software deadtime is applied to eliminate any entrances immediately following another entrance.

Only those muons which satisfy a set of "HasBestEntrance" (HBE) conditions will be considered in the final analysis. Approximately 70% of muon entrances pass all HBE cuts. The full list of requirements is:

$\mu\text{SC Trigger}$	Both μSC TDC signals must be present and in coincidence.
$\mu\text{SC DT}$	The entrance cannot be within 29 ns of a previous entrance.
$\mu\text{SC BE}$	The entrance cannot be within $30 \mu\text{s}$ of either end of a block.
Kick	The entrance must have triggered the kicker.
$!\mu\text{CL}$	The entrance must be a real muon rather than a clock event.
$!\mu\text{SCa}$	There must be no μSCa signal coincident with the entrance.
μPC	There must be coincident μPC X, Y, and XY clusters.
$\mu\text{SC PP}$	Pileup protection on the μSC .
$\mu\text{SC WFD PP}$	Pileup protection on μSC WFD pulses above threshold.
$\mu\text{PC XY PP}$	Pileup protection on μPC XY clusters.

A pileup protection flag is also set for the μSCa , but is not required for a HBE event. This cut provides minimal improvement to the pileup protection efficiency, and reduces statistics by 15% in DSM2 due to the misaligned μSCa . Similarly, additional unused pileup protection flags are set for individual μPC X and Y clusters as well as μSC WFD pulses below the amplitude threshold. These are sensitive to beam electrons, and significantly reduce statistics for DSM1-3 due to the issues with the upstream beam collimator. For DSM4 and later the losses from all of these cuts are low, but the standard pileup protection already achieves an inefficiency lower than 1×10^{-3} which should be sufficient. See section 10.3.2 for more information about how the pileup protection efficiency is determined.

7.7 *Outputs*

Each analysis module can define its own set of output histograms, typically containing diagnostics and calibration data. Some specific modules also generate histograms to perform data quality checks and calculate some preliminary analysis results that are used to verify proper operation of the experiment. These histograms are saved to a separate histogram file, and may be used as input to the online display program mentioned in section 4.8.2 or a similar but more complicated display program specialized for offline analysis.

The MU output histograms are also used as the input to another script designed to generate a variety of trend plots tracking changes between data files. These trend plots are helpful for identifying stable periods that may be summed and used for high-statistics calibration fits. They are also used to identify and eliminate files that have an unstable muon beam or various detector error conditions. The calibration and quality cut procedures were mainly developed by Rachel Ryan for the R2014 dataset, and are described in detail in her thesis[62]. The author repeated this analysis for R2015, with only minor adjustments.

The TPC forced trigger waveforms are also of interest for some specialized analyses, so they are separated into a FT skim file for easy access. These files are structured as a list of forced trigger events, each containing the waveforms for all TPC islands as well as the trigger time and a flag indicating whether the TPC pulser was activated. They also include the associated muon entrance object, to allow some rudimentary event selection cuts.

7.7.1 *Event Tree*

The large data blocks used by the MIDAS files are good for efficiently collecting data during the experiment, but not conducive to efficient analysis. Only a few diagnostic histograms make use of the time within the block or look at long range correlations, and the block ID is also only relevant for a few calibration procedures. In general the analysis is only concerned with signals directly related to a given muon entrance, so it is beneficial to reorganize the output data after these calibrations and diagnostics are completed by MU.

The new format is based on the concept of individual "events" centered on each muon entrance. All detector objects within $35 \mu\text{s}$ of the entrance time are added to the event. The $25 \mu\text{s}$ kicker window after the muon entrance contains most of the information associated with the interactions and decay of the muon, while the $25 \mu\text{s}$ before the entrance are useful for characterizing the backgrounds. The additional 10μ on either end of the event ensures no relevant information is inadvertently cut from the event, such as drifting charges in the TPC that may have a significant delay between their creation and detection times. These padding regions also include information about adjacent muon entrances, which is useful for some studies. This system does introduce some redundancy, with objects often duplicated across two or even three events with overlapping time windows, but this is more than made up for by the computational efficiency allowed by having all relevant data readily accessible.

The events are saved in a TTree data structure, part of the ROOT data analysis framework. This data structure, generally referred to simply as a tree, is essentially an advanced version of a database table or spreadsheet. It is related to the standard computer science notion of a tree in that the simple list of column headers in a normal table is replaced with a branching tree structure defining the arrangement of the data. The tree consists of a collection of "branches" that can represent any C++ objects, which each contain a combination of either additional branches or "leaves" representing individual data values. In the case of the MuSun Event Tree the structure represents the format of an individual event, with branches for each of the detector objects created by MU.

Individual events are recorded as entries in the Event Tree. If the branch structure is equivalent to the the column headers in a normal data table then these entries are equivalent to the table rows containing actual data values. Internally, the data associated with each branch or leaf is stored in a separate buffer known as a basket. When these buffers are full they are automatically compressed before being written to the data file. This system allows efficient storage of large amounts of data, while still allowing relatively quick random access to specific event or branch data by only decompressing the relevant baskets.

The Event Tree is primarily intended to store the processed detector response objects

produced by MU, and the raw input data such as waveforms or individual TDC triggers are generally discarded. However, some low level data is retained if it is deemed potentially interesting for further analysis, most commonly for the TPC where forced trigger waveforms are saved as well as extended or double-pulse waveforms which might be associated with muon-catalyzed fusion. The neutron detector data is also retained in full, the analysis for these detectors was developed later than the rest of the software framework and primarily takes place in stage 2. Despite discarding as much unnecessary low level data as possible, the Event Tree files are about 60% larger than the original raw data files.

7.8 Summary

The first stage of the MuSun analysis software, MU, is responsible for unpacking the MIDAS data files produced by the experiment and converting them into useful objects representing the responses of the various detectors. For waveform data produced by the WFDs this involves reconstructing the waveform from the raw data packets and then identifying pulses in the signal. In the case of the TPC WFDs there is added complication from the dual-gain system used to increase the dynamic range, requiring high and low gain signals to be stitched together to form the final reconstructed waveform.

TDC data consists of a sequence of hit times and associated parameter ID numbers, which must be decoded and assigned to the corresponding detector. For most detectors the time of each hit is recorded, for the wire chambers they are grouped into clusters of coincident hits. Coincidences between hit clusters in different detector planes are then used to triangulate the position of the particle detection.

Finally, the analysis of the muon entrance detectors is fully completed at MU, including identification of multiple detector signals associated with a given muon and the creation of various entrance quality and pileup protection flags. The resulting muon entrance objects are used as the center points for creating events containing all of the processed detector objects within a $\pm 35 \mu\text{s}$ window, which are saved to an output event tree file. Further processing of this output data is performed by the second stage analysis described in chapter 8.

Chapter 8

HIGH-LEVEL ANALYSIS SOFTWARE

The stage 2 analysis program processes the Event Tree files output by MU, and is therefore named Muon Tree Analysis (MTA). The primary goal of MTA is to use this data to reconstruct particle trajectories representing the muon stop in the TPC and the outgoing decay electron. The neutron detector data is also analyzed at this point. After this processing the results are saved as a simplified format designed to enable easy creation of lifetime histograms with various analysis cuts.

MTA is built in a modular fashion somewhat similar to MU, but to simplify development it is written as a self-contained C++ class hierarchy rather than being based on MIDAS. Several miscellaneous analysis modules also perform calibrations or study specific systematic effects. Often these specialized analyses only apply to specific unusual events, so to streamline development such events may be flagged and separated into smaller skim files with the same format as the original event tree.

Stage 3 will also be discussed briefly. Unlike the first two analysis stages which are essentially synonymous with the MU and MTA programs, stage 3 is a general term encompassing several smaller scripts. For the most part these were originally dedicated histogramming modules included in MTA, which have since been spun off and adapted to analyze the output of stage 2. The purpose of these scripts is still primarily to generate histograms, and they perform very little additional processing.

8.1 TPC Tracking

The most important processing step performed at MTA is the tracking of muons stopping in the TPC. MU has already processed the TPC waveforms and identified individual pulses, so

now the task is to group these pulses into track objects and to estimate the muon stop position. This seemingly straightforward task is complicated by the presence of muon-catalyzed fusion reactions, and several tracking algorithms have been developed in an attempt to minimize interference due to additional energy deposition from fusion products. Most of the tracking algorithms involve specific energy cuts to distinguish these particles, so precise energy calibrations are important and will also be discussed here.

8.1.1 Clustering

The first step of tracking is to group related pulses together into clusters. To qualify for inclusion in a cluster each pulse must be close to one of the other pulses in the cluster, both spatially and temporally. The spatial cut is relatively simple, merely requiring the two pulses to be on adjacent pads in both X and Z, including the diagonally adjacent pads. For runs prior to R2014 a gap of one pad row in Z is allowed, since the earlier iterations of the TPC could potentially miss the smaller pulses created when the muon energy is split across two pads as it crosses an X boundary.

The time difference Δt between two pulses is calculated from the end of the first pulse to the start of the next pulse. Two different clustering windows are used, a tight 500 ns window and a less restrictive 1 μs cut. Muon pulses are typically about 1 μs in length, so these correspond to effective clustering windows of about 1.5 and 2 μs as measured using the peak times. However, recall that the muon pulses are actually produced at nearly the same time and differences in detection time result from time for charges to drift from the original Y position down to the anode plane. These effective time windows are therefore only valid for straight tracks, and the distance between peaks may be larger for a steeply angled track. basing the coincidence cut on the pulse tails minimizes the effect of the track angle, since the positions where the track leaves one pad and where it enters the next pad should be essentially identical. This means that over 99.9% of tracks would pass the cut even with no time gap allowed, and the non-zero coincidence windows exist mainly to prevent the case where the signal briefly drops below threshold.

Clusters must additionally include a pad on the first Z row and have a Z extent of at least two pads to be valid muon track candidates. Isolated clusters in the middle of the TPC are most likely the result of neutron scattering or fusion events. Clusters confined to the first Z row do include some muon stops, but could also be secondary particles from muon scattering upstream and too little information is available to make a positive identification.

8.1.2 Tracking Algorithms




Figure 8.1: Example of an event display for a p+t fusion event. The muon track is sloping downward causing downstream pulses to arrive earlier. The fusion adds to the energy on the final pad to produce the large peak, with the proton then traveling upstream and creating the earlier pulse on the next to last row. The right panel shows the pads included in the cluster along with their amplitudes (A), energies (E), and peak times (t).

If the TPC clusters merely contained muon pulses then the most downstream pad in the cluster could trivially be identified as the stop position. The most problematic of these are the protons emitted by p+t fusions, which have a range of nearly a full pad length and deposit nearly 3 times the maximum energy deposited at the muon on a single pad. Figure

8.1 shows an example of an event display showing a muon track with a p+t fusion included in the cluster, see chapter 9 for a discussion of the various ways fusion products affect the data from the TPC. If the fusion pulses are mistaken as part of the muon track it can easily cause a mis-reconstruction of the stop position.

Several muon tracking algorithms, or trackers, have been developed that attempt to compensate for the presence of fusion products in the TPC. Much of the initial tracker development was done by Michael Murray, and extensive studies of several tracker variations may be found in his thesis[54]. Significant upgrades have been made to the more sophisticated algorithms since then, ultimately resulting in four commonly used trackers:

Basic Simply takes the most downstream energy deposition to be the stop, which performs well for clean muon tracks but is easily thrown off by fusions.

Road Stop Threshold (road) Uses an energy cut to identify the first pad including p+t fusion energy as the stop, and extrapolates the X and Y position from the upstream pads rather than using the stop pad itself.

Upstream Developed by Duncan Prindle, this tracker improves on the road tracker by estimating the stop Z using a fit to the expected range-energy relation for the upstream muon pulses.

Proton Direction (P-Dir) Takes the opposite approach of examining the stop cluster more closely and identifying upstream and downstream-going protons from their subtly different energies, caused by differences in how the fusion overlaps with the original muon track. The author has since expanded further on this idea by adding additional energy cuts to handle more event types such as ^3He fusions or electron scattering, for a full discussion see section 9.4.

In summary the basic and road trackers both produce biased Z estimates but in opposite directions, as shown in figure 8.2. The road tracker produces much better X and Y predictions




Figure 8.2: Diagram of fusion interference from protons emitted either upstream (top) or downstream (bottom). In both cases the basic and road trackers report stops on the 6th and 5th Z row, respectively. The upstream and P-Dir trackers should correctly report the true muon stop position.

than the basic tracker though, and its relative simplicity and intuitive behavior make it a useful comparison point. The upstream tracker fit is relatively unbiased but is less precise than the other trackers as it completely omits the stop pad from the analysis, making it primarily useful for specific analysis applications aiming to be agnostic about the details of the stop cluster. Finally, the P-Dir tracker leverages the high energy resolution of the TPC to distinguish different event topologies with finely tuned energy cuts. Although the P-Dir tracker is rather complicated it produces the most accurate reconstruction, and is the primary tracking algorithm used in this thesis.

8.1.3 Absolute Gain Calibration

The TPC gain calibration from MU only scaled the high gain signal to match the low gain for each pad, and did nothing about the overall size of the pulses. A separate absolute gain calibration is required to convert the pulse amplitudes or integrals into actual energy values. Most importantly these energies must be consistent between pads, as a particular pad with a bad energy calibration could interact very strangely with the trackers.

The preferred calibration method is to look at isolated ${}^3\text{He}$ fusions, which have a consistent energy and are normally localized to a single pad due to their short range. These are selected by looking for a single additional pulse on the muon stop pad but delayed from the main muon stop cluster. Plotting the amplitude of these pulses produces a clean peak, which is calibrated to sit at 400 keV as shown in figure 8.3. The energies could also be calibrated using this peak, but that was found to produce inconsistent results so the energy calibration will be handled separately.




Figure 8.3: Amplitudes (left) and energies (right) produced by delayed muon catalyzed fusion products. The ${}^3\text{He}$ peak is ideal for defining calibrations.

Selecting delayed p+t fusions in a similar way also appears to produce a useful calibration peak, but only when multi-pad clusters are allowed. Restricting the selection to single pulses on the stop pad makes the peak much less distinct, which is not ideal for generating pad-by-

pad calibrations. The p+t fusion peak is instead used as a consistency check to ensure the rest of the calibration procedure is working correctly.

Because the relation between the pulse amplitude and energy values depends on the pulse shape, it is best to perform the energy calibration using muon pulses since they will be used for tracking. With the amplitude calibration already determined by the ^3He pulses, the energy calibration may be obtained by plotting the muon energies vs amplitude and fitting the result. This produces a linear relationship, but to achieve a good fit an offset value is required in addition to a scale factor as shown in figure 8.4. The energy offset is negative, and is most likely the result of the TOT pulse finder failing to account for the tails of pulses after they fall below threshold.




Figure 8.4: Energy calibration using muon pulses on the stop pad. The central solid line is the fit, with the box indicating the estimated width of the band. The dashed lines show the extrapolation of the fit results to high and low amplitudes.

Finally, the first Z row requires special consideration since the trackers all require two or more pads to form a track. Without any valid stops on the first row, none of the procedures above will work properly. The energy calibration can be corrected by simply selecting stops

in the second row and then looking at the pulses on the first row, which limits the usable energy range compared to the stop pad itself but is still usable. The amplitude calibration has no choice but to use the upstream muon pulses as well, but this performs poorly due to a wide spread in muon amplitudes. The majority of this spread comes from variation in the muon stop depth within the pad, so a sharp calibration peak may be obtained by using a narrow amplitude cut to constrain the stop depth as shown in figure 8.5. An amplitude of 400 keV is chosen to match the ${}^3\text{He}$ amplitude and minimize any additional error introduced by the calibrations on the stop pad.




Figure 8.5: Left: Plot of pulse amplitudes on the stop pad (A_0) and the pad before the stop pad (A_1). Right: A_1 distribution for a narrow slice around $A_0 = 400$ keV, with the full A_1 distribution shown for reference. The low-amplitude peak is caused by muon decay in flight (dif), and may be reduced with an electron time cut.

8.2 Electron Tracking

The second main task of MTA is to reconstruct electron trajectories using timing from the eSC and spatial information from the ePCs. This section will describe the processing for each of these detectors, as well as their combination to form full tracks. There will also be a

discussion of the relative merits of several alternative electron definitions used by the MuSun experiment.

8.2.1 *eSC Coincidences*

With the eSC pulse finding and time correction already completed at MU, the processing at MTA mainly consists of simple coincidence checking. Tubes connected to the same scintillator are first matched to form 2-fold coincidences, using a 30 ns coincidence window to avoid any bias from the time of flight delay which may be up to ± 6 ns for extreme Z. Inner and outer 2-folds are then matched together using a tighter 20 ns window. The final eSC objects contain pulse times, amplitudes, and integrals for each of the four tubes on the eSC panel. They are also categorized based on the number of coincident pulses found:

4-Folds Most electrons should deposit energy in both scintillator layers and produce pulses in all four detectors. These signals are required for most analysis purposes.

3-Folds These are most commonly produced when one pulse in a 4-fold is too small and fails to trigger the digitizer, or is otherwise missed for some reason. They are not used for the main analysis, but are useful for studying detector efficiency.

2-Folds These represent events where energy was deposited in only one scintillator layer. This is possible for electrons but usually requires high-angle scattering or a highly angled incoming track to avoid hitting both layers. Instead 2-folds are often caused by neutrons, which naturally cause single localized scattering events rather than forming continuous tracks like charged particles.

Singles Individual isolated pulses are produced mainly by electronic noise and afterpulsing, and are unlikely to be physically meaningful.

The signals from the individual eSC PMTs exhibit significant afterpulsing effects, shown in figure 8.6. The sharp peaks seen in some channels are the result of ion emission from




Figure 8.6: Autocorrelations for eSC pulses from individual PMTs (left) and for coincidences (right). Ideally these plots should show a constant background plus some residual contribution from muon decays, as shown by the dashed green line on the right. The prominent peaks and oscillations seen in the individual PMTs are the result of afterpulsing following a real electron detection, and are significantly reduced by requiring coincidences. The final 4-fold coincidences have nearly eliminated afterpulsing and are in good agreement with the expected background. Data from eSC segment 8 is shown as it has the worst afterpulsing, the other segments typically show less pronounced effects.

the electrodes during the last few amplification stages of the PMTs. The smoothly varying signal is likely an indication of a slight loss of vacuum in these relatively old PMTs, resulting in ion production throughout the tube instead of only at the electrodes.

All tubes have a prominent afterpulsing peak approximately 75 ns after the initial pulse, which is most likely an artifact of the 71 ns eSC island length. This signal is removed by applying a 100 ns artificial deadtime after each pulse, which has the added benefit of ensuring only a single pulse may be present in the coincidence windows described above. Aside from this initial peak the afterpulsing is uncorrelated between tubes, so these signals are unlikely to form proper coincidences. The afterpulsing artifacts are significantly reduced for 2-fold coincidences and nearly eliminated for full 4-fold coincidences.

8.2.2 *eSC Position*

Although the primary purpose of the eSC is to precisely measure the electron times, they may also be used to obtain coarse position information. Because the eSC is composed of 16 individual segments, a detection in one segment immediately localizes the angular coordinate ϕ to the $\pi/8$ angle covered by that panel. An approximate Z position may also be extracted from the time of flight delay between the detectors at either end, according to the formula

$$Z = \frac{1}{2}V_{eff}(t_u - t_d) \quad (8.1)$$

where t_u and t_d are the upstream and downstream pulse times and $V_{eff} = 142$ mm/ns is the effective velocity of light inside the scintillator.

The time difference between the eSC tubes may be determined with roughly 1 ns accuracy, and the results may be improved somewhat by averaging the Z estimates from both scintillator layers. The result is an average Z resolution of approximately 50 mm for electrons near the center of the panel, degrading somewhat towards either end as the pulse in the far detector becomes more degraded due to attenuation in the scintillator. The attenuation length was estimated at 126 cm, in reasonable agreement with the literature value of 140 cm for BC404 scintillator, so pulses traveling the full length of the scintillator should be attenuated by roughly half. In principle the attenuation provides an independent measure of the Z position, but the eSC waveforms are too noisy to measure the pulse size with much precision making this method far less reliable than the time of flight measurement.

8.2.3 *ePC Tracking*

The ePC clusters already have Z and ϕ values determined by the analysis at MU, which along with the known radii of both chambers is enough to fully specify the interaction location in three dimensions. Tracking should therefore be a simple matter of selecting coincident clusters from ePC1 and ePC2 and performing a linear extrapolation from the hit coordinates. The challenge lies in determining which pair of clusters to use, with noise creating additional clusters while inefficiencies potentially cause missing information in the real clusters. Once

again, the electron tracking has been completely redesigned with these issues in mind and the new system is described in detail in appendix A, but an overview will be provided here.

The eSCs are always used for timing so the ePC tracking can only affect the lifetime by introducing a time-dependent efficiency factor. The general philosophy of the new electron tracking system is therefore to be as inclusive as possible to avoid inadvertently rejecting good electrons. Any combination of electron detector hits that pass a set of relatively loose cuts is considered a potential track, which may result in a very large number of tracks in events with many extra noise clusters.

It is assumed that each eSC hit is associated with a single real electron, so including all of these track permutations would result in double counting. The list of candidate tracks associated with a given eSC hit are passed through a series of quality checks designed to select a single "best track" for that electron. Ideally an ePC track should agree with eSC ϕ and Z , pass near the muon stop position (for decays) or the beam axis (for backgrounds), and be composed of high quality triple coincidences in both ePCs. To check for bias in the track selection, a few alternative "best tracks" are selected prioritizing a different one of these objectives. These should generally give similar results, but can differ when strict geometric cuts are applied.

8.2.4 *Electron Definitions*

Any given analysis may specify its own set of criteria for determining which electrons qualify for inclusion, known as the electron definition. For the eSCs there is the choice of allowing 2- or 3-folds in the analysis, while the ePCs permit a wide range of cluster quality and geometric cuts. However, the most fundamental choice is whether to use the ePC track information at all, with many analyses opting for scintillator-only electrons based solely on the eSC data.

Scintillator-only electrons have the advantages that the eSCs are very reliable and their analysis process is relatively simple, so they should not have any strange detector effects. The disadvantage is that, although rough position information is available, there is no way to determine the track angle. This makes it impossible to reliably distinguish decay electrons

from beam particles or cosmic rays, and any analysis based on the scintillators alone must accept these backgrounds.

In contrast, the ePCs may be used to select electron tracks that actually originate near the muon stop, greatly reducing backgrounds. The downside is that the ePCs suffer from a variety of data quality issues including noise, electronics pickup and feedback, sparks, and more. Historically there was a major discrepancy between the full tracks and the scintillator-only analysis, which the author eventually traced to a software bug in an analysis cut intended to eliminate low quality ePC tracks. Fixing the bug was not enough to completely resolve the problem, and this investigation prompted the major rewrite of the tracking software mentioned previously in an attempt to handle these errors more carefully. Although the new tracking is a major improvement it is still somewhat sensitive to these data quality issues.

The electron definition thus represents a choice between two systematic effects. If scintillator-only electrons are used then the effects of the backgrounds must be corrected for, while for full tracks a correction is needed for the ePC noise and efficiency effects. Both methods should agree after all relevant corrections are applied, and comparison of different electron definitions provides an important cross check of the results. Ultimately however, the beam backgrounds are much better understood and quantified than the details of the ePC behavior, so while the full tracks are very useful for studying other systematic effects the scintillator-only electrons will continue to be preferred for the final analysis.

8.3 Neutron Detectors

The liquid scintillator detectors used to study neutrons are also sensitive to gamma rays or charged particles, and in fact detect many more gamma rays than actual neutrons. These particles are distinguishable because they interact differently in the detectors, with the gamma rays and charged particles primarily affecting electrons while the neutrons deposit energy via binary collisions with atomic nuclei. As shown in figure 8.7 the resulting pulses have subtly different shapes, with the neutron pulses tending to have a significantly elevated tail relative to those produced by electromagnetic interactions. The process of identifying neutron pulses

using their tail is known as pulse shape discrimination (PSD), and is crucial for removing the large gamma background and extracting usable neutron signals from these detectors.




Figure 8.7: Comparison of neutron and gamma pulse shapes. Vertical dashed lines indicate the integration windows used to identify the peak and the tail of the pulses.

8.3.1 Pulse Finding

The neutron pulse finding algorithm faces different constraints from those of the other detectors. Effects such as afterpulsing can increase the apparent size of the pulse tail and interfere with the PSD cuts, so while an artificial deadtime would normally be sufficient to remove afterpulsing it renders the waveform unusable in this case. To identify waveforms with multiple peaks the pulse finder starts from the maximum sample the pulse finder iterates through each sample in turn, keeping track of the smallest sample seen. If any sample before the peak exceeds this minimum by more than a designated threshold it is considered a new peak, and the waveform is flagged. Because the pulse tails can fluctuate somewhat, samples after the peak must also exceed the value of a pulse template by the same threshold to avoid false positives. These pulse templates are used in the interpolation procedure described below.

In addition to the double pulse check, events with multiple islands in the same detector are also rejected to avoid any interference with subsequent pulses. Again this would normally be accomplished with an artificial deadtime, but considering the large majority of pulses are not actually neutrons they will typically be removed by later analysis cuts. This results in an odd situation where the remaining neutron pulses may be influenced by deadtimes from particles that have later been eliminated from the analysis, which was deemed unnecessarily complicated compared to rejecting these events entirely. A minimum amplitude threshold is also applied for each detector prior to the multiple island check, to avoid rejecting signal pulses due to islands containing only noise.

8.3.2 Interpolation

The neutron digitizers operate at a nominal sampling rate of 170 MHz, corresponding to a bin width of 5.88 ns. For the majority of the R2015 production run the sampling rate was actually half this amount, producing a bin width of 11.76 ns. This is problematic since the neutron pulses have a typical width of about 15 ns, resulting in only a few samples measuring the peak particularly at the lower sampling rate. To obtain a waveform accurate enough to be used for pulse shape discrimination the raw samples must be interpolated to reconstruct a smooth pulse shape.

A simple cubic spline was used as a first attempt at pulse interpolation, but tended to result in unwanted oscillations particularly near the pulse onset. A solution for this problem has been developed for engineering applications which explicitly sets the derivatives at each point to prevent overshoots and oscillations[42]. The constrained spline performs well for the background and tail of the pulse but cuts off the peak, so it is blended together with the standard spline near the peak to produce a smooth interpolation.

This modified spline results in acceptable performance at the nominal sampling rate, but at half the rate it becomes unreliable. The issue is that the shape and amplitude of the interpolated peak changes significantly depending on how the samples are aligned with the true time of the pulse. To solve this problem a set of pulse templates was introduced to

ensure a consistent shape for all waveforms and also to allow fitting of overflow pulses. To avoid biasing the estimation of the size of the pulse tail, the template is only used for the peak of the pulse and then transitions back into the spline for remainder of the waveform.




Figure 8.8: Left: Neutron template generated from many combined waveforms. Right: Interpolation of a single waveform using various methods.

Templates are obtained by combining many different pulse waveforms, offsetting the samples based on the estimated pulse time. This may then be iteratively refined by inputting the template into the interpolation function and the improved pulse time estimates to create a new merged waveform. One of the resulting templates as well as a comparison of the different interpolation methods is shown in figure 8.8.

8.3.3 Particle Identification

Electrons and other charged particles coming from within the eSC cylinder are easily identified by looking for coincidences between the neutron detectors and the electron detectors. With the neutron detectors mounted directly behind the eSC panels, it is difficult for any electrons to enter the neutron detectors without first passing through the eSCs. Coincident tracks also clearly align with the neutron detectors, as shown in figure 8.9.




Figure 8.9: ePC1 electron clusters coincident with neutron detector pulses.

To distinguish the neutron and gamma pulses we quantify the size of the tail by integrating starting 26 ns after the start of the pulse. The tail fraction is defined as the ratio of this tail integral to the integral of the entire pulse from -40 ns to +200 ns around the maximum. Two dimensional histograms of tail fraction against total energy are known as PSD plots, and gamma and neutron pulses produce two distinct bands as shown in figure 8.10. The integration ranges as well as the details of the pulse interpolation procedure have all been optimized in an attempt to maximize the separation between these two bands.

A PSD cut is a curve defined in this space tuned to separate these bands. Preliminary PSD cuts were set by fitting vertical slices of the PSD plots with a pair of peaks and setting a cut value between them. More sophisticated cuts have since been developed by selecting events with different populations of neutrons and gamma rays, which are described in appendix F. In either case, the general idea is to generate cuts at various separations from the gamma band, which are then used to assign a quality rating to each neutron pulse. The numbers of neutrons and gammas in a sample will vary dramatically depending on the other event cuts being used, so different neutron quality levels may be appropriate depending on the analysis.




Figure 8.10: Example neutron PSD plot, with several PSD cut variations shown in red. The large lower band is from gamma rays, while the much smaller upper band is produced by neutrons.

8.4 Outputs

Once again, each analysis module may generate its own set of output histograms containing diagnostic information used to check the operation of the software. An optional set of dedicated histogramming modules can also create more physically meaningful histograms using the output of the main analysis modules. The histograms are saved in a separate histogram file, and are typically merged to create high statistics histograms for each dataset or even an entire run. There are no standard scripts for further analysis of these histogram files as there were for MU, but the MTA histograms are grouped into folders by analysis module and are relatively easy to look through manually.

8.4.1 *Skimming*

In the course of developing a new analysis module, it is normally necessary to run many test passes. In the case of specialized analyses which are only concerned with rare events, repeatedly processing data files containing mostly unrelated events represents a significant waste of time and compute resources. To streamline the development of such modules, a system was implemented whereby any analysis module may flag events as potentially interesting. These flagged events are then copied and saved to a separate output file, which can be used as input to subsequent analysis passes instead of the original file.

To ensure the analysis of the skimmed files works correctly, it is important that the skimmed events are reloaded from the input file and copied without modification. The original skimming implementation neglected this step, resulting calibrations and other corrections being applied twice when processing the skim files. In addition to the flagged events, 0.1% of all events are saved to each skim file as "minimum bias" events, ensuring that some normal events are available for comparison even when processing skimmed files. Without these events it would often be necessary to also process some standard files separately, potentially resulting in version mismatches or other discrepancies between the different files. In some cases it is also important for the analysis to differentiate between various types of skimmed events, so the set of skim flags is saved and is readable by the analysis modules.

Several skim files are created, containing events flagged for various reasons:

Neutron Selects events containing pulses in the neutron detectors. Allowing gammas selects 13% of events, which is easily reduced to 2.5% by requiring a PSD cut although this requires re-skimming if the neutron pulse finding is adjusted.

Clock Selects events with a μCL entrance, for use in background studies. Includes about 2.6% of events for R2015, but depends on the clock rate.

Fusion Selects events using a loose E_0 threshold to identify p-t fusion reactions, for use in fusion interference studies and tracker development. Includes about 2.2% of events.

Multi-Cluster Selects events with isolated TPC pads or extra clusters in addition to the main muon track. These are useful for the impurity analysis, as well as for selecting delayed fusion reactions. Includes about 0.53% of events.

8.4.2 *Lifetime Tree*

The results of MTA are compiled into a simplified representation known as the Lifetime Tree, discarding most low-level details but retaining the majority of the information normally needed for making event selection cuts and generating lifetime histograms or other results:

Entrance All the information from the original MU entrance objects is retained. Some additional flags are also added to keep track of error conditions for the event, such as the WFD errors discussed in appendix E.

TPC Tracks A separate TPC track object is saved for each of the tracking algorithms described above. These include reconstructed stop positions and various flags for properties of the stop such as any fusion products or additional tracks were identified. They X, Y, amplitude, and energy values for each Z row in the TPC, which might vary between trackers due to different clustering windows.

Extra Clusters Besides the muon tracks, the tree includes a list of the pad number, energy, and time of any extra pulses and clusters found in the TPC.

Electrons Detailed tracking information is saved for every electron marked as some variant of "best track" and including an eSC 4-fold coincidence. These include time, energy, and Z values from the eSC, and time, ϕ , and Z values for each ePC as well as a variety of flags and track quality information. Some minimal information is also saved for tracks without 4-fold coincidences.

Neutrons For neutron detector pulses the time, energy, and tail fraction are saved as well as a PSD rating and some flags for selecting specific types of particle.

To avoid unnecessarily cluttering the lifetime tree, events must pass a few checks before being saved. Events are rejected if they have any error conditions, except for TPC WFD errors which are included to study their effects. Clock events must pass all best entrance conditions, while real muons must pass the μ SC and kicker checks but the other standard cuts are skipped to allow testing of various entrance cuts. TPC Pulser events are also rejected. Finally, events that have no TPC track and no electrons detected are not particularly interesting, so most of these are rejected with only 0.1% retained for testing purposes.

8.5 Stage Three

Originally MTA included an analysis module dedicated to the generation of output lifetime histograms for use in the final analysis. When the lifetime tree was introduced this module was separated from MTA and adapted to use the new tree file, becoming the primary stage 3 analysis script simply named Lifetime Tree Analysis (LTA). Without the memory constraints imposed by operating as part of MTA alongside many other modules, this script was free to expand and now includes many variations of the standard lifetime histograms with different cuts applied. The script has also been forked several times, with collaboration members free to add any histograms they deem interesting to their own version. Because the stage three analysis is so cheap to run there is no real need to coordinate in order to avoid redundancy or develop a standard set of shared histograms, and making it easy for individuals to quickly implement and test new analysis ideas is much more efficient from a user time perspective.

Several other similar scripts are more specialized for use specific analyses. A script for generating neutron lifetimes was used for the analysis in chapter 11, and one dedicated to analysis of fusion and TPC tracking issues was used for the analyses in chapters 9, 13, and 14. These scripts have a more limited scope than the main LTA script, and so are easier to modify and faster to run. Typically it is possible to produce new results using the full production statistics within only a few hours.

8.6 *Summary*

The second stage of the MuSun analysis software, MTA, is responsible for combining individual detector objects produced by MU into higher level objects representing complete particle trajectories. Muons in the TPC are tracked using specialized algorithms designed to minimize interference from muon-catalyzed fusion products. Electron tracks are also reconstructed by combining information from the eSC and ePCs, although in practice a minimal electron track definition is often used which omits the ePC information. The neutron detector pulses are also processed at this point.

Several outputs are produced in addition to the standard diagnostic histogram files. A variety of skimmed files contain only events relevant to specific analyses, in the same format as the original event trees. Most importantly, a simplified lifetime tree is created containing most of the important event information. Several stage three analysis scripts further process the lifetime tree data to produce various histograms needed by the final analysis, and may be rapidly iterated to test new ideas with minimal delay.

Chapter 9

MUON TRACKING

TPC track interference refers to instances in which a secondary particle emitted after a muon stop deposits additional energy in the TPC, resulting in a mis-reconstruction of the stop position when this energy is added to the muon track. These particles may be either muon-catalyzed fusion products or the decay electrons themselves, both of which have a time-dependent probability of causing tracking errors. As such, if these errors shift such events across the fiducial volume boundary in a biased way it has the potential to significantly distort the observed muon lifetime.

This chapter will describe the various ways in which migration may occur, as well as the development of specialized tracking algorithms that are robust against these effects. The performance of these trackers will also be evaluated using simulated Monte Carlo data. Chapter 13 will follow up with data-driven techniques for the measurement and correction of any remaining lifetime shift due to interference.

9.1 Formalism

Before delving into the specific behaviors of individual trackers, it will be helpful to develop some formalism to describe fusion interference events and their effects on the muon lifetime. This will be primarily focused on p+t fusions since they have by far the largest effect, but the same formalism also describes n+³He fusions with only minor modifications. Electron interference will be discussed briefly at the end of this section, but because it is both less important and more complicated than fusion interference it has not been analyzed to the same level of detail.

9.1.1 Fusions Event Decay Times

For the purposes of this chapter a simplified kinetics model will be used which ignores the difference in capture rate between the hyperfine states as well as the population of $\mu^3\text{He}$. This is still a reasonable approximation to the full kinetics described in section 3.3, and is sufficient to capture the important behavior for fusion interference. With these simplifications, the probability to observe a muon decay at time t takes a pure exponential form

$$e(t) = \lambda_0 n_\mu(t) = \lambda_0 e^{-\lambda_- t}, \quad (9.1)$$

where $n_\mu(t)$ represents the probability of a muon surviving until time t , λ_0 is the free muon decay rate and λ_- is the total decay rate including the muon capture rate Λ_d .

The decay time distribution may be decomposed into separate distributions for events with and without a fusion, denoted by the superscripts f and nf respectively:

$$e(t) = e^f(t) + e^{nf}(t). \quad (9.2)$$

Although the total is a pure exponential decay, the two components are both non-exponential.

To derive the time distribution for fusion events, consider the probability for a fusion at time t_f followed by a decay at time t

$$e^f(t, t_f) = \tilde{f}(t_f) e(t - t_f). \quad (9.3)$$

The function $\tilde{f}(t)$ is the probability of observing a fusion at time t , and is the product of an underlying fusion time distribution $f(t)$ and the probability $n_\mu(t)$ for the muon to survive until that time

$$\tilde{f}(t) \equiv f(t) n_\mu(t). \quad (9.4)$$

In general throughout this chapter a tilde (\sim) over a variable will be used to label observable quantities, to distinguish them from the idealized theoretical values which are written without a tilde. The probability of the decay occurring before the fusion is zero, since the muon is necessary to catalyze the fusion reaction. For decays occurring after a fusion the probability

depends on the fusion time and the delay between the fusion and electron times:

$$\begin{aligned} e^f(t, t_f) &= f(t_f)n_\mu(t_f) \times \lambda_0 n_\mu(t - t_f) \\ &= e(t)f(t_f) \end{aligned} \quad (9.5)$$

The probability of a muon decay with any prior fusion is then obtained by integrating over the fusion time:

$$\begin{aligned} e^f(t) &= \int_0^t e^f(t, t_f) dt_f \\ &= e(t) \int_0^t f(t_f) dt_f \end{aligned} \quad (9.6)$$

The corresponding non-fusion decay time distribution is simply the complement:

$$e^{nf}(t) = e(t) - e^f(t) = e(t) \left(1 - \int_0^t f(t_f) dt_f \right). \quad (9.7)$$

It is clear that the integral will approach zero for early decays independent of the exact fusion time distribution $f(t)$, as there is simply too little time for fusions to occur.

The fusion time distribution itself may be derived from the target gas conditions as described in section 3.3. Fusions predominantly occur from the quartet hyperfine state, as shown in figure 9.1. The large majority of fusions occur within the first few microseconds after the muon stop, before the quartet state is fully depleted.

9.1.2 Acceptance Function

The discussion so far has assumed perfect detection efficiency for both fusion and non-fusion events, which is rather unrealistic. All muon tracks will have a time-independent acceptance efficiency α which includes the effects of the fiducial volume cut as well as any random tracking errors unrelated to fusions. Fusion events will also have an additional time-dependent efficiency $\alpha_f(t_f)$ to account for the effects of fusion interference. With these efficiencies included, the observed decay time distributions become

$$\tilde{e}^f(t) = \alpha e(t) \int_0^t f(t_f) \alpha_f(t_f) dt_f \quad (9.8)$$

$$\tilde{e}^{nf}(t) = \alpha e(t) \left(1 - \int_0^t f(t_f) dt_f \right). \quad (9.9)$$




Figure 9.1: Population of the quartet (N_d) and doublet (N_q) hyperfine states as well as the $\mu^3\text{He}$ population (N_{He}) as a function of time. Fusions predominantly occur from the quartet hyperfine state.

The time-dependent efficiency term only appears in the fusion time distribution because it represents events being incorrectly accepted or rejected by the tracking cuts, not mis-identification of fusion events as non-fusions.

It is helpful to introduce the differential fusion acceptance function η , defined as

$$\eta(t) \equiv \alpha_f(t) - 1. \quad (9.10)$$

This function characterizes the change in acceptance efficiency for fusion events relative to non-fusion events, with a positive value indicating an increase in fusion events while a negative value indicates losses. The observed electron time distribution then simplifies to

$$\tilde{e}(t) = \alpha e(t) \left(1 + \int_0^t f(t_f) \eta(t_f) dt_f \right). \quad (9.11)$$

To a good approximation the shape of $\eta(t)$ is independent of the muon stop position or track angle, so it is useful to split it into a magnitude and a normalized shape

$$\eta(t) \equiv \eta_0 \bar{\eta}(t). \quad (9.12)$$

With this formulation $\bar{\eta}(t)$ essentially reflects the probability of a given fusion event being included in the muon stop cluster. Fusions occurring before the end of the clustering time window are always included in the muon cluster, so $\bar{\eta}(t)$ takes the value 1. At very late times the fusions will be entirely separate from the stop cluster, and $\bar{\eta}(t)$ takes the value 0. Between these two extremes the precise shape depends on both the clustering time window and the type of fusion, as shown in figure 9.2.




Figure 9.2: Monte Carlo clustering probabilities for p+t and n+³He fusions as a function of fusion time, with 1 μ s or 500 ns clustering windows. These estimates lack the requirement that the fusions cause migrations, which will shorten the tail of the true $\bar{\eta}$ functions.

The longer tail on the p+t $\bar{\eta}$ function is caused by the long range of the proton, which makes it much easier for even quite delayed fusions to connect to the stop cluster. This only applies to downward-going protons that can catch up to the drifting charge cloud, which is why $\bar{\eta}$ quickly drops to about 0.5 before transitioning into the extended tail section. In contrast the $\bar{\eta}$ function for ³He lacks this long range component and quickly drops to near zero after the end of the clustering window. The P-Dir tracker uses the shorter 500 ns time window for clustering, while the other standard trackers all use the 1 μ s window.

Finally, the claim that this clustering probability is independent of the muon track is not completely accurate. With a downward-sloping muon track the fusion may be clustered with the first upstream pad rather than the stop pad, slightly extending the effective clustering window. In principle this results in a slight Y dependence as the track angle is correlated with the stop position, but the tracks are not steep enough to have much effect. For stops near the edges of the TPC it is also possible for protons to leave the active volume, reducing the clustering probability.

The probabilities plotted in figure 9.2 are derived by counting the number of fusions of each type that are included in the muon cluster. Not all of these fusions cause migration, particularly in the ^3He case, but $\bar{\eta}$ should include only migrating events. The long tail of the p+t $\bar{\eta}$ function results from protons emitted directly downward, but such events will be very unlikely to migrate since the fusion products do not move horizontally. When restricted to those fusions that actually cross a pad boundary, $\bar{\eta}$ will have a somewhat shortened tail which depends on the proximity of the stop to the edge of a pad. Determination of the true shape requires the Monte Carlo data to be re-analyzed with fusion products removed, but it appears to be a relatively minor correction as only the tail of the distribution is affected. From figure 9.1 it is clear that the length of the clustering window is sufficient to capture the majority of the quickly decaying fusion time distribution.

The decay electron time distribution now becomes

$$\tilde{e}(t) = \alpha e(t) \left(1 + \eta_0 \int_0^t f(t_f) \bar{\eta}(t_f) dt_f \right). \quad (9.13)$$

The integral may be taken to have a fixed form depending only on the muon kinetics and the clustering parameters, so the degree of distortion in the muon lifetime is entirely determined by η_0 . For small amounts of distortion the change in the measured disappearance rate will be approximately proportional to the size of the distortion:

$$\Delta\lambda \propto \eta_0. \quad (9.14)$$

9.1.3 Migration

The parameter η_0 measures the extent to which fusion events are gained or lost by the tracking cuts. This is most commonly caused when the true stop position is on one side of the fiducial boundary but fusion interference causes it to be mis-reconstructed on the other side. This process is known as migration and may be pictured as stops moving across the boundary, although it is important to remember that the muons are not physically moving in the detector.

Let N denote the total number of muon stops that would be reconstructed in a given target volume, in the absence of any fusion interference. Let F refer to the number of those muon stops that have a subsequent fusion. These quantities are similar but not identical to the true numbers of stops in the volume, as fusion-independent tracking errors still apply. However, since such errors by definition apply equally to both types of events, N and F are simply proportional to each other. In terms of the formalism developed above, this is equivalent to including α but setting η to zero.

The number of events that are actually observed is labeled \tilde{N} , and the corresponding number of fusions is \tilde{F} . For non-fusion (NF) events these definitions are equivalent, so

$$NF \equiv N - F = \tilde{N} - \tilde{F}. \quad (9.15)$$

The observed number of fusion events differs from the idealized value by the number of migrations into the volume, denoted M :

$$M \equiv \tilde{F} - F = \tilde{N} - N. \quad (9.16)$$

Because these parameters all refer to the total number of events, to relate them to η it is necessary to integrate over the decay electron time distribution. Performing this integral, one finds

$$\int_0^\infty \tilde{e}(t) dt = \alpha \frac{\lambda_0}{\lambda_-} \left(1 + \eta_0 \int_0^\infty \tilde{f}(t) \bar{\eta}(t) dt \right). \quad (9.17)$$

Note that while this appears similar to equation 9.13 above, it is now the observable fusion time distribution $\tilde{f}(t)$ inside the integral instead of the theoretical value $f(t)$. The factor

λ_0/λ_- is just the probability of a muon decaying and producing an electron rather than being captured, and is not relevant when counting the number of muon stops observed in the TPC. Obtaining the total number of events is then simply a matter of replacing the single-muon detection efficiency α with the total number of events N :

$$\tilde{N} = N \left(1 + \eta_0 \int_0^\infty \tilde{f}(t)\bar{\eta}(t)dt \right). \quad (9.18)$$

The number of migrations now becomes

$$M = N\eta_0 \int_0^\infty \tilde{f}(t)\bar{\eta}(t)dt. \quad (9.19)$$

Once again, these functions are approximately independent of the details of the muon stop distribution, so this integral is simply a constant. The migration fraction is therefore proportional to η_0 , and thus also to the rate shift assuming relatively low amounts of migration:

$$\frac{M}{N} \propto \eta_0 \propto \Delta\lambda. \quad (9.20)$$

The exact size of the correction and methods for quantifying the migration will be discussed further in chapter 13. However, a rough estimate is that a 1% change in the number of p+t fusions is expected to produce about a 5 s^{-1} shift in the measured muon disappearance rate. The fusions are quite distinctive due to their high energy and it would be easy to veto them entirely, but to ensure they are reconstructed identically to the non-fusion events to the required level is quite challenging.

9.1.4 *Electron Interference*

Finally, consider the case where the tracking interference is caused by the decay electron itself rather than by a muon-catalyzed fusion event. The low energies of the electron pulses mean they have a much less drastic effect on the tracking system, and require a less precise correction to reduce the influence on the muon lifetime to an acceptable level. There are also complications involved in modeling the electron interference so a detailed formalism has not been developed in this case, but some general observations can be made.

To use notation similar to that for fusion interference, let $\eta_e(t)$ represent the differential acceptance function for decay electrons. Rather than integrating over a separate fusion time, in this case the migration probability depends directly on the muon decay time:

$$\tilde{\epsilon}(t) = \alpha e(t) (1 + \eta_e(t)). \quad (9.21)$$

As in fusion interference, any non-zero value of η_e indicates migration.

In addition to time η_e also depends strongly on the decay electron emission direction, since the electron itself is producing the migrations. Electron tracks produce signals in the full TPC volume, so η_e will have a long tail for downward-going electrons that can catch up with the drifting charge cloud in the TPC. For the low energy electron signals the migration effects are also sensitive to the overlap between the electron and muon pulses, so η_e is not constant within the clustering window like the fusion $\bar{\eta}$.

Electron interference affects the muon lifetime in two different ways. First, on average η_e is higher for decays at early times and will therefore increase the observed disappearance rate of $\tilde{\epsilon}(t)$. Second, the ability of downward-going electrons to add to the stop cluster even at late decay times creates a vertical asymmetry effect. The asymmetry is very pronounced when selecting specific eSC segments or electron directions, and can also appear when studying all electrons due to variations in detector efficiency. All of these problems may be minimized by reducing the magnitude of η_e .

9.2 TPC Energy Deposition

Although the formalism is somewhat abstract, the conclusions are clear. For muon-catalyzed fusions there is a correction to the muon lifetime which scales with the net quantity of fusion migrations into the fiducial volume. The story for electron interference is more complicated, but again the number of electron migrations should ideally be minimized.

If precise spatial information were available then the correct muon stop location could be identified easily, but this is not the case in practice. Recall from chapter 5 that the TPC was forced to use a coarse spatial resolution to maximize energy resolution, while Y

tracking additionally requires knowledge of the true pulse time that is typically unavailable for fusions. Studying the energy deposition of a track in detail is therefore the best way to extract information about the event and to estimate the stop position.

Muon tracks are approximately parallel with the Z axis, so the primary concern for tracking is estimating the Z stop depth. The total energy deposited on each Z row was previously defined as E_n , numbered in reverse order with E_0 referring to the final row of the cluster. To be consistent with the formalism developed in this chapter these observable values should properly be denoted \tilde{E}_n , with E_n instead numbered relative to the true muon stop pad. This convention will be used for clarity when comparing observed and theoretical energy depositions, in cases where this is helpful to understand the observed migration behavior. Otherwise this thesis will mainly refer to the observable values \tilde{E}_n , and the tilde will be omitted to simplify the notation.

9.2.1 E_0 vs E_1

As discussed in chapter 5 E_0 and E_1 are both strongly correlated with the muon stop depth within the pad, producing a distinct band in a 2D scatter plot. In addition to the primary muon stop band there is a great deal of additional structure in these plots, as shown in figure 9.3. At these energies the relevant processes are mostly $n+{}^3\text{He}$ fusions and electron interference, $p+t$ fusions deposit much higher energies and are off the scale of this plot.

In events with $n+{}^3\text{He}$ fusions the resulting ${}^3\text{He}$ or $\mu^3\text{He}$ ion deposits about 400 or 600 keV of additional energy, respectively. This typically occurs on the stop pad, producing duplicate copies of the original stop band that have been shifted in E_0 . If the muon stops close to the pad boundary then it is also possible for the helium to add to E_1 , extending these bands past the top of the main band albeit with reduced amplitude.

Double-fusion events produce two additional bands with energy shifts of 800 and 1000 keV. In the first fusion the muon must be recycled rather than sticking to the helium, so there is no 1200 keV band corresponding to two $\mu^3\text{He}$ ions. Note that it is highly unlikely for both helium ions to migrate even for stops very close to a pad boundary, so the high- E_1 tails are

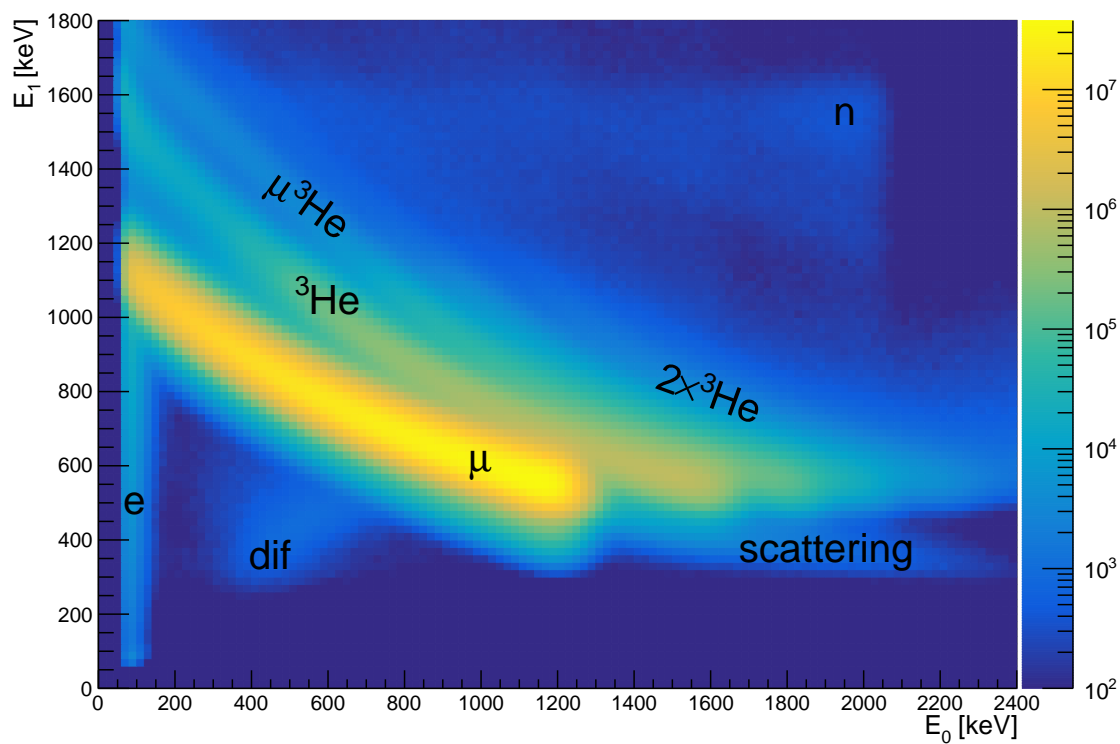


Figure 9.3: Plot of E_1 vs E_0 where each visible band has been labeled with the particle or process that produces it. Particles such as electrons or fusion products create additional bands by adding to the muon energy. An electron time cut has been applied to minimize decay in flight (dif) events, which have low energies due to not stopping completely. Scattering events can also cause a muon to stop early, but in this case its remaining energy is deposited on E_0 .

less prominent relative to the single-fusion bands. In principle this pattern continues with a series of progressively fainter bands corresponding to increasing numbers of fusions, but events with three or more fusions are very rare with the MuSun target conditions. Finally, the structure in the top right of the plot is caused when the fusion neutron also scatters in the TPC, and will be described in more detail below.

The vertical band at the bottom left (labeled e) is due to electron interference, which in rare cases may deposit up to about 100 keV in a TPC pad. When this occurs on the pad following the muon stop it can add a row to the cluster, shifting the labels of the row energies. Using the tilde notation to distinguish observed from theoretical values, we have:

$$\tilde{E}_0 = E^e \qquad \tilde{E}_1 = E_0^\mu, \qquad (9.22)$$

where superscripts e and μ indicate the energies deposited by the electron and muon respectively. Because the observed E_0 only includes the electron energy it does not depend on the stop depth. This band extends all the way to $E_1 \approx 1200$ keV, although it cannot be distinguished from the primary stop band. The decay electron often adds energy to E_0 and E_1 as well, but since these energies are low and smoothly distributed they result in smearing the main band rather than creating a distinct new band as with $n+{}^3\text{He}$ fusions.

There are two additional features to discuss, related to truncated muon tracks. Decay in flight (dif) events have low energies since the downstream portion of the track is missing. These are largely eliminated with an electron time cut but a small fraction remains, while at the edges of the TPC there is also a similar signal from tracks leaving the active volume. The other structure at higher E_0 is caused by large-angle scattering events, which may similarly cause muons to stop early but result in the remaining energy being deposited on the stop row rather than being lost.

These details demonstrate the power of the TPC and are important for analyses presented in subsequent chapters, but the primary concern is p+t fusion events. These deposit approximately 4 MeV of additional energy in the TPC, so in figure 9.4 the energy range has been expanded to show them.




Figure 9.4: Plot of E_1 vs E_0 with expanded range to show p+t fusion events, where each band has been labeled with the particle that produces it. Subscripts denote energy deposition on the muon stop pad (0), one pad upstream (u), or one pad downstream (d).

Because of the long range of the proton, p+t fusions produce a more complicated structure than n+³He fusions. Roughly half of p+t fusions are entirely contained on the stop pad, producing another shifted duplicate of the original stop band labeled pt₀. In the other half of events the proton crosses a Z row boundary, creating the large diagonal band in the center of the plot. This is actually two different bands created by protons going either upstream or downstream, the former having slightly higher energy as a result of increased overlap with the muon track. This distinction is used by the P-Dir tracker to identify the proton direction and reconstruct the stop position accordingly, and is discussed in more detail in section 9.4.1.

Double-fusion events once again add duplicate copies of the fusion bands. The combination of a p+t and a n+³He fusion actually produces four separate diagonal bands depending on the proton direction and whether the muon remains bound to the helium, but for downstream-going protons the double-fusion bands overlap with the upstream-going single-fusion band. This structure may be seen more clearly in figure 9.8 below, where Monte Carlo truth information has been used to isolate the two proton directions. A pair of p+t fusions produces another diagonal band at much higher energies, in this case the substructure is mostly lost due to both protons being emitted at independent random angles. It is also possible to have both electron interference and fusion interference in a single event, causing the vertical electron migration band to extend all the way to the highest energy scales.

Finally, with the expanded energy range it becomes apparent that there are in fact three separate clusters of neutron scattering events, depending on which pad the neutron scattered on. Neutrons scattering on or upstream from the stop pad add to the corresponding E_n value, while downstream scattering extends the cluster producing observed energies

$$\tilde{E}_0 = E^n \qquad \tilde{E}_1 = E_0^\mu + E^{He}. \qquad (9.23)$$

As in electron migrations the observed E_0 is produced solely by the neutron scattering and is independent of the muon stop depth. Neutron scattering differs from the other processes discussed in that it can deposit anywhere from 0 to 2 MeV as mentioned in chapter 5. Instead of distinct bands the resulting structures are therefore smeared across a wide energy range,

although they have a sharp cutoff at the maximum neutron energy deposition. In double-fusion events the neutron scattering may also add to the p+t fusion energy, producing the "pt+n" cluster at the top left of figure 9.8.

9.2.2 Upstream Energies

Unlike E_0 and E_1 which can take on a wide range of values depending on the stop depth, E_2 and higher are less sensitive as the dE/dx curve quickly flattens out away from the Bragg peak. Starting with E_3 the detector energy resolution and random noise begin to dominate, with only small shifts in the peak position from one pad to the next. The upstream tracker can extract an approximate stop position by fitting all of these upstream peaks to a model of the Bragg curve, but individually they are of little interest.

E_2 is informative because it occupies an intermediate zone between the stop pads E_0 and E_1 and the upstream pads E_3 and up. Like E_0 and E_1 it still has some sensitivity to the muon stop depth, but like the upstream pads it is typically not sensitive to fusion products. The peak of the E_2 distribution is also fairly well separated from both E_1 and E_3 , making it possible to detect instances where migration events have caused a shift in the row numbering.

9.3 Standard Trackers

Because migrations are essentially the result of tracking errors, several specialized tracking algorithms have been designed with fusion interference effects in mind. Section 8.1 has given a brief overview of the tracking system, but it is worth covering these trackers in more detail. The P-Dir tracker will be discussed separately in the following section. The other three trackers are all based on TPC clusters derived using a $1 \mu s$ time window.

9.3.1 Basic Tracker

The basic tracker takes the naive approach of simply picking the last Z row of the cluster as the stop position. X and Y are then determined by the pad and center time of the largest

pulse on the final Z row. In fact this hardly qualifies as a tracker, as these coordinates are output directly by the clustering algorithm with no additional processing step.

For clean muon stops the basic tracker would perform well, but any energy deposited after the muon stop results in a downstream mis-reconstruction. Even when the Z coordinate is correct the X and Y determination will select p+t fusions over the muon, as they deposit much more energy. The basic tracker represents a worst-case scenario where no attempt is made to correct for interference effects.

9.3.2 Road Tracker

The road tracker attempts to divide the input cluster into a "road" cluster containing only clean muon pulses and a "stop" cluster containing any potential fusion products. The division between the two clusters is found by setting an energy threshold of 1300 keV, slightly above the highest E_1 energy possible for a clean muon stop. The first Z row to exceed this energy threshold must include additional energy from fusion and is identified as the stop position, or the last Z row is taken if the threshold is not reached.

The upstream rows up to but not including the estimated stop row are defined as the road cluster. The stop row is omitted even for the tracks that do not exceed the road energy threshold, as there may still be interference from ^3He fusions on the stop pad. The X position is then taken to be the X position of the last row before the stop pad, while Y is estimated by calculating the slope of the road cluster and extrapolating forward onto the stop row.

For p+t fusions the road tracker has purely upstream Z migrations, opposite to the behavior of the basic tracker. Electron interference and ^3He fusions can still migrate downstream as they do not deposit enough energy to reliably trigger the energy threshold. The X and Y reconstruction is less accurate than the basic tracker for non-fusion events, but is far more consistent which is the real goal here. The large amount of upstream fusion migration is not ideal, but the road tracker has the benefit of being simple and transparent and until recently was the standard tracking algorithm used in MuSun.

9.3.3 Upstream Tracker

The upstream tracker extends the basic idea of the road tracker by fitting muon pulses to the expected Bragg curve, producing a new estimated Z position. By completely omitting the stop row from even the Z reconstruction the upstream tracker achieves a very low sensitivity to most interference effects, whether it be p+t or ^3He fusions or electron interference. A check for outliers is also included to detect the rare cases where a fusion neutron scatters on an upstream pad. However, the downside of this tracker is that it is much less precise than the other trackers, so it is only used for some specialized analyses.

9.4 P-Dir Tracker

The original P-Dir tracker created by Michael Murray[54] used a diagonal cut in $E_0 - E_1$ space to separate the upstream and downstream-going proton bands. This general idea has since been expanded on by the author, with additional cuts to identify migrations caused by most of the other event types identified above. The proton direction cut itself has also been changed to a more precise curve, since the accuracy of this cut is crucial to the performance of the tracker. The result is a rather complicated algorithm with many parameters, but when properly tuned it performs much better than the simpler trackers.

9.4.1 Proton Direction Cut

Figure 9.5 illustrates how upstream and downstream-going protons align with the muon stop. The proton deposits an energy of approximately 3 MeV over a range of 13.2 mm, more than 80% the length of a pad. The triton deposits roughly 1 MeV of additional energy, but over a range of under 1 mm. An energy cut of $E_0 + E_1 > 4$ MeV is therefore used to identify fusion events, which will easily pass this threshold with the addition of the muon energy. In the case of an upstream-going proton the fusion typically overlaps the last two pads of the muon track, while a downstream-going proton only overlaps with the final pad of the muon track and creates an additional row containing only the proton energy.




Figure 9.5: Diagram of the Bragg curves for events with upstream-going (left) and downstream-going (right) protons. The arrows indicate the track vectors for the muon (blue), the proton (orange), and the triton (red). The total energy of the p+t fusion does not depend on direction, but in the upstream-going case an additional pad of the muon track overlaps with the fusion.

It is important to note that even with an upstream-going proton, the triton may still cause a downstream migration if the stop position is within 1 mm of the pad boundary. In fact, because the proton and triton are emitted back-to-back at least one of them must travel downstream to some degree. This means that any fusion event with E_0 below 1 MeV must necessarily be a downstream migration, as even a stop near the front of a pad with an upstream-going proton will still see the full energy of the triton deposited on the stop pad. Similarly, events with E_0 greater than 3 MeV must have the stop on the final row, as this cannot be produced by the migrating proton alone.

To define the cut between the two proton directions, the E_0 vs E_1 plot is first rotated by 50° to approximately align the bands with the vertical axis as in figure 9.6. The bands are not entirely straight, so to maximize accuracy a curved cut is fit between them. The separation of the bands may be improved somewhat by subtracting E_2 from the horizontal




Figure 9.6: Left: Proton direction cut defined in $E_0 - E_1$ space that has been rotated clockwise by 50° so that the p+t fusion bands are approximately vertical. Right: Alternate cut where E_2 has been subtracted from the X axis, moderately increasing the band separation. The left band corresponds to downstream-going protons, while the right band is produced by upstream-going protons. Dashed lines indicate the estimated transition between the bands, while the solid lines indicate the final cuts including minimum and maximum E_0 limits.

axis, since for downstream migrations E_2 is actually the E_1 of the muon track and thus higher energy. This is only possible for tracks that are at least three pads long, so for very short tracks the first cut must be used instead.

Because the proton direction cut is so sensitive, even a slight variation in energy calibration between datasets may significantly reduce its accuracy. Manually specifying these cuts in a consistent way across all datasets would be both tedious and error-prone, so instead an automated procedure has been developed to generate the cuts. Each horizontal slice of the histograms in figure 9.6 is individually fit with a pair of Gaussian peaks to approximate the

two bands. This is not entirely accurate, but appears to agree reasonably well with the data as shown in figure 9.7.




Figure 9.7: Fits to selected slices of the proton migration bands.

The cut value for each slice is determined by the point at which the two Gaussians intersect, as this represents the minimum classification error. The cut values estimated for each slice have some jitter, so a high-order polynomial is used to fit these points and produce a smooth curve. The points at the bottom of the plot which are above the maximum E_0 value are fit to a separate polynomial, since the shape of the downstream proton band changes significantly after the muon and triton begin being included on the stop pad. This procedure results in a kink at the end of the cut which ideally should also be replaced with a smooth curve, but this is a minor concern.

9.4.2 Double Fusions

Events with $E_0 + E_1 + E_2 > 7.5$ MeV are identified as having at least two separate p+t fusions. This threshold is slightly lower than twice the expected fusion energy, since unlike with the single-fusion threshold there are no major event types of interest at lower energies which must be excluded. E_2 is now added to the total to account for the possibility of

back-to-back protons resulting in both upstream and downstream migration simultaneously. If all three rows are above the 1.3 MeV road threshold then it may immediately be identified as such an event, and the true stop position must be on the middle row.

Assuming only two rows are above the road threshold, one of the two should be the stop row. However, it is not possible to define a double-fusion proton direction cut since the energies of these events are rather smeared out. Instead, whichever row has the highest energy is chosen as the stop position, since the two randomly oriented fusions are unlikely to both migrate.

Consider the worst-case scenario of a muon stopping very close to a row boundary. The probability of each proton migrating will then be just under 50%, with a 25% chance of both doing so. The highest energy row will therefore be correct over 75% of the time, as in the case where the protons travel in opposite directions the symmetry will be broken by the fact that the tritons primarily deposit energy on the stop pad. For a more typical stop position somewhere in the middle of the row the accuracy will be significantly higher, as many protons will be emitted at too high of an angle to reach the boundary even if they do travel in the right general direction.

9.4.3 Helium Fusions

The migrating portions of the ^3He and $\mu^3\text{He}$ stop bands are cleanly separated from the primary muon stop band, as shown in figure 9.8 below. They are selected with hand-drawn polygonal cuts, but there is no clear way to identify the migration direction of these events. Because of the short range of the helium ions, these events necessarily have a muon stop position very near the pad boundary. Events that stop at the back of a pad with a downstream-going helium have a nearly identical energy signature to events which stop at the start of a pad and have an upstream-going helium. This was not a concern with p+t fusions, where the presence of the triton made these two configurations quite distinct.

To minimize the number of migrations, the E_0 cut for these bands is set to half of the expected energy deposition of the helium ion. A ^3He deposits about 400 keV and a $\mu^3\text{He}$

deposits about 600 keV, so that becomes a 200 keV and a 300 keV threshold respectively. Although these cuts cannot definitively identify any given event, on average approximately equal numbers of events will be shifted in either direction over any given boundary which should be enough to avoid any net migration.

One final consideration is that, while this cut does a reasonable job of reproducing the true muon stop position, it may instead be better to deliberately induce migrations to better mimic the behavior of non-fusion events. The short range of the helium means migrations are only possible for stops at the very beginning of the pad, and in the absence of fusion many of these events would have failed to cross the pulse-finding threshold and been reconstructed upstream. Slightly adjusting the ^3He tag as well as the p+t E_0 thresholds may be something to consider for the upcoming final analysis pass.

9.4.4 *Electron Interference*

Although it is the final issue to be discussed here, electron interference is actually the first effect that must be corrected for by the tracker. This is to account for events affected by both electron and fusion interference, resulting in a normal fusion event followed by one or more rows containing only small electron pulses. If the electron interference is not corrected first then the carefully calibrated fusion cuts described above may be applied to entirely the wrong set of pads. The correction is achieved by dropping the small electron pulses from the end of the cluster and redefining E_0 and E_1 relative to the new truncated cluster.

All electron migration events will have E_0 below 200 keV. The electron interference correction occurs in three steps. First, in rare cases the electron may deposit energy on more than one pad row, resulting in a string of low-energy pads following the muon stop. To correct this effect rows are repeatedly dropped as long as both E_0 and E_1 are below 200 keV.

Second, if E_0 is below 200 keV and E_1 is above 200 keV but below the muon stop band, the event may be an electron migration but could also be a decay in flight or punchthrough event. The migrations should be corrected while the decay in flight events should not, and

punchthrough muons should perhaps be shifted further downstream instead although this is not implemented. An E_2 cut may distinguish these cases, since for electron interference we observe $\tilde{E}_2 = E_1^\mu$, while for these other low-energy events we instead lose the last pad and the observed energy is at most $\tilde{E}_2 \lesssim E_3^\mu$. A minimum E_2 threshold of 600 keV is effective at isolating the low-energy electron interference events.

Finally, the case with both electron and fusion interference is characterized by an E_0 below 200 keV but either E_1 or E_2 above 1300 keV. This may also overlap with some purely fusion events, but as discussed above any fusions with a low E_0 are likely also downstream migrations. In fact, above the ${}^3\text{He}$ band the E_0 threshold may be increased to 300 keV and above the p+t band it may be further increased to 1 MeV, to match the ${}^3\text{He}$ and p+t E_0 thresholds respectively.

After these correction steps E_0 and E_1 should include either the muon pulse or high-energy fusion pulses. The fusion cuts may then be safely applied to these corrected energies, potentially resulting in an additional fusion correction. The only exception is the electron pulses which overlap the main stop band, which cannot be cleanly identified and result in some uncorrected electron interference for non-fusion events as discussed in chapter 14.

9.4.5 Summary

Figure 9.8 shows the approximate regions identified as downstream migrations by the P-Dir tracker cuts. Note that this diagram does not capture the E_2 dependence of the p+t or electron cuts, so some events included in the indicated regions will not actually be flagged as migrations and vice versa. The cuts are clearly not perfect, and the lower edge of the p+t fusion cut in particular could use some refinement to avoid biting into the upper edge of the neutron scattering events. There is also currently no attempt to cut on the downstream neutron scattering events, which might be something to add for a final analysis pass. Neutron scattering events are quite rare, so precisely tuning the cuts in this region makes little difference to the results and was considered low priority.

To test that these regions do indeed represent migrating events, Monte Carlo data may




Figure 9.8: Top: E_1 vs E_0 for R2015 data. The P-Dir energy cuts are shown in red with crosshatched regions considered downstream migrations, although note that several of these cuts also depend on E_2 which is not represented here. Bottom: similar plots for Monte Carlo data, split into events with true muon stop positions on the next to last (left) or last (right) row of the cluster.

be used to see the energy distributions for migrating events. Figure 9.8 also shows the cases where the muon stop is on the last or next to last row of the cluster. Muon stops after the last row of the cluster are only possible if E_0 is below threshold, and stops two or more pads upstream are rare since they require long-ranged electron interference or neutron scattering. The P-Dir tracker accurately identifies most regions that contain primarily downstream migrations, but there are also significant areas of overlap where no definitive cut can be made.

In general the P-Dir cuts do a very good job of picking out the regions of predominantly downstream migration. Some small refinements may still be implemented, but the overlap between these two plots makes any major improvements impossible without some alternate way to distinguish these events. It is important to note that, since these plots are based on Monte Carlo truth, they do not include non-fusion tracking errors and threshold effects which should ideally be reproduced to minimize any difference with the non-fusion events. See appendix G for some more discussion of this issue.

One final consideration is that while the P-Dir tracker and the road tracker use the same road energy threshold, they identify the road cluster differently. The road tracker starts at the beginning of the track and finds the earliest row above threshold, labeling that as the stop position and every row before it as the road cluster. The P-Dir tracker goes in the opposite direction, working backward from it's estimated stop position until it finds a row that is below threshold. In most cases these select the same road cluster, but if upstream neutron scattering causes a single row to exceed the threshold then the road tracker will truncate the track while the P-Dir tracker will not be affected. The track angle is extrapolated from the first and last pads of the road cluster, so neutron scattering will only affect the stop Y estimate if it occurs on these rows.

9.5 Tracker Performance

The behavior and relative performance of each tracker is most easily assessed using Monte Carlo data, where they may be compared directly against the true stop position. The

simulation should model fusion migrations fairly accurately since it has been carefully tuned to reproduce the observed energies of fusion products. The electron energy deposition in the TPC is not as good of a match to real data, so non-fusion migrations may be compared between different trackers the actual size of the effect will be somewhat off.

The amount of migration also depends strongly on the stopping distribution, as even p+t migrations are relatively short-ranged and require muon stop positions near one of the fiducial boundaries. The Monte Carlo has been adjusted to approximately match the stopping distribution of the data, but this is sensitive to the exact beam tune and changes somewhat between datasets. Thus, while the simulation is very helpful for developing a qualitative understanding of the tracker behavior and tuning them to be more accurate it cannot give a precise quantitative migration measurement. Instead data-driven approaches are used to estimate a final lifetime correction, which will be discussed in chapter 13.

9.5.1 Longitudinal Migration

Longitudinal migration along the Z axis is the largest concern, as the muons predominantly travel in this direction and a wide range of stopping depths is possible. In contrast the X and Y coordinates of the muon stop are usually strongly correlated with the positions of the upstream pulses, only deviating significantly for large-angle scattering events which anyway do not affect the measured lifetime. In fact, because the road and P-Dir trackers determine the Y coordinate by projecting the slope of the track forward to the stop pad one of the largest sources of fusion-related Y migration is Z migration causing the track to be projected by the wrong distance.

The natural first choice to estimate the amount of Z migration is to plot the difference between the true and reconstructed Z positions, as in figure 9.9. However, these plots can be counterintuitive because, apart from the upstream tracker, the trackers only determine a stop pad rather than a precise depth. The reconstructed stop position is therefore always taken to be the center of the stop pad, and only errors larger than half a pad length actually constitute migrations. Also recall that the true goal is to minimize the difference between




Figure 9.9: Difference between true and reconstructed Z for events with (left) and without (right) p+t fusions.

fusion and non-fusion events, although since fusions are so much more likely to migrate this essentially amounts to minimizing the fusion migrations.

For p+t fusion events the basic and road trackers are nearly mirror images of each other, as expected since they are designed to produce uni-directional migration. The P-Dir tracker is far more symmetrical and the migrations in either direction are reduced by nearly an order of magnitude compared to the other trackers. For events without a p+t fusion the P-Dir and basic trackers have similar small amounts of upstream migration due to events where E_0 falls below the pulse finding threshold, while the road tracker has some extra migrations caused by events where E_1 managed to exceed the road threshold. The road and basic trackers have nearly identical downstream migration tails, while the electron cuts included in the P-Dir tracker allow it to significantly outperform both although it still cannot identify short-ranged electron migrations which overlap too much with the primary stop band. See appendix G for more information about tracker errors for events without p+t fusions.

Another way to visualize p+t fusion migrations is shown in figure 9.10. Here events are classified as upstream or downstream migrations based on whether the reconstructed stop Z is less than or greater than the true Z stop row, where stops in the first 1% of a pad are

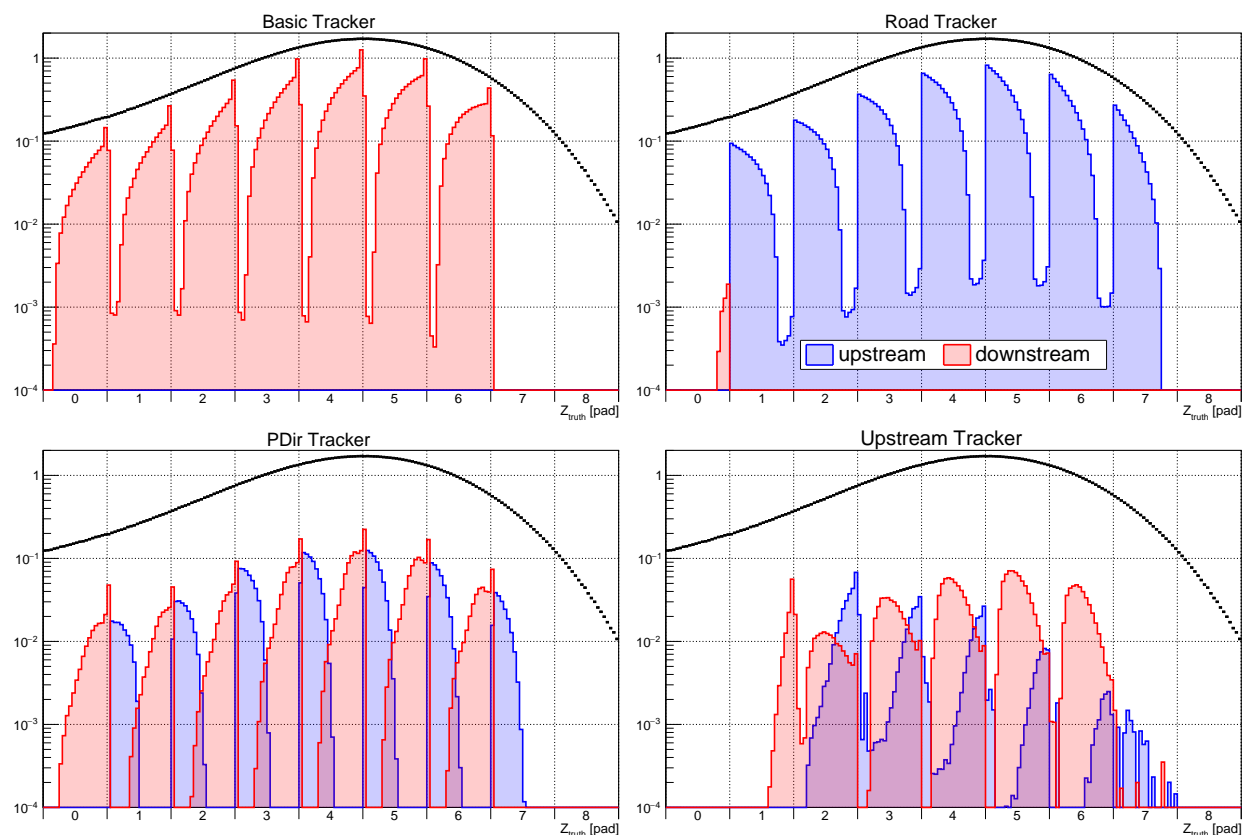


Figure 9.10: Number of migrations caused by $p+t$ fusion events, as a function of true stop Z position. The expected non-fusion migrations have been subtracted to isolate the effects of fusion interference specifically. The stop distribution for all events is shown by the black curve, while the number of events which migrate either upstream or downstream are drawn in blue and red, respectively. The histograms are normalized based on the number of events in the fiducial volume, but do not represent migration probabilities which would be obtained by dividing by the black curves.

counted as part of the previous row to approximate the E_0 pulse-finding threshold. The number of migrations in either direction is then plotted as a function of the true stop Z position, to show how the migration probability changes with proximity to the pad boundaries. This is performed for events with and without p+t fusions, and the non-fusion migration probabilities have been subtracted from figure 9.10 to isolate the effects of fusion interference. Note that $z=7$ is the last physical pad while $z=8$ represents muons stopping past the end of the TPC, although the upstream tracker may report stops in row 8 or even 9 as it is not constrained to the size of the TPC.

Tracker	Downstream	Upstream	Total	Signed
Basic	21.38	0	21.38	21.38
Road	0	18.76	18.76	18.76
Upstream	1.77	0.52	2.28	1.25
P-Dir	2.63	2.19	4.83	0.44

Table 9.1: Probabilities of tracker Z errors caused by p+t fusions, as percentages. The signed column refers to subtracting the upstream migrations from the downstream migrations. Note that the majority of these errors remain inside the fiducial volume and do not actually constitute migrations across the boundary.

Table 9.1 shows the estimated probabilities for tracker errors caused by the presence of p+t fusions. This includes tracking errors that merely shuffle events within the fiducial volume, so the number of real migrations across the fiducial boundary will generally be much lower. Once again the basic and road trackers exhibit purely downstream or upstream migration respectively, with similar Z dependence. Stops near the edge of a pad are very likely to migrate, but the probability drops quickly with distance from the pad boundary and is suppressed by over two orders of magnitude for stops at the other end of the pad. The P-Dir tracker instead has migrations in both directions across each pad boundary in a

nearly symmetrical manner, resulting in significant cancellation and a very low migration probability overall. Finally, the upstream tracker shows even fewer fusion migrations than the P-Dir tracker, but it is less symmetrical so the upstream and downstream migrations do not cancel as well.

9.5.2 Transverse Migration




Figure 9.11: Difference between true and reconstructed Y for events with (left) and without (right) p+t fusions. The road and P-Dir trackers produce nearly identical results in Y since they use the same extrapolation procedure from the upstream Y slope.

Migration in the X and Y directions is less critical than in Z, but can still be a significant contribution to the net migration. Y migration in particular is important to study, as it is crucial to avoid stops in the solid anode and cathode planes. Fortunately the reconstructed Y stop position may be compared to the true value much more precisely than X or Z, as shown in figure 9.11, and the Y fiducial cut may be adjusted to an optimal value. In contrast the X fiducial cut can only be set at pad boundaries, and in most recent analyses the X cut has actually been removed entirely rendering X migrations largely irrelevant.

The basic tracker has a distorted Y distribution for p+t fusion events, since it identifies




Figure 9.12: Direct comparison of P-Dir tracker Y errors for events with and without p+t fusions. The red line indicates the number of excess fusion migrations beyond the given Y value.

the largest pulse on the final row as the stop position which typically selects the fusion itself. For events without p+t fusions the basic tracker actually performs better than the others, with a sharper central peak and smaller long-range tails. However, it still has medium-range shoulders produced by a combination of electron and ^3He fusion interference. These are approximately 10 mm wide as a result of the roughly $2 \mu\text{s}$ effective clustering window that was shown in figure 9.2, since the drift velocity is $5 \text{ mm}/\mu\text{s}$.

The P-Dir and road trackers produce nearly identical Y distributions since they use the same method of calculating the slope of the road cluster and then projecting forward onto the stop pad. This projection loses some precision for non-fusion events, but produces similar looking peaks for all event types with little of the distortion seen by the basic tracker. Figure 9.12 shows a more direct comparison on the fusion and non-fusion events reconstructed by the P-Dir tracker, and while fusions do broaden the peak somewhat it is a fairly small effect. The tracker error is estimated by subtracting the non-fusion distribution from the fusion

distribution, and then taking the integral of the tails to find the total number of fusions which migrated past a given Y boundary. This error estimate quickly falls below the 10^{-3} level for any stops more than a few millimeters from the fiducial boundary, and even for stops directly on the boundary there only appears to be a difference of about 2%

9.6 Conclusion

A formalism has been developed to describe tracker interference and its effects on the muon lifetime, which indicates that the observed rate shift is proportional to the number of migrations into the fiducial volume. To minimize migrations specialized tracking algorithms are used which attempt to consistently reconstruct the muon stop position without being affected by fusion products or electron tracks in the TPC. The energy that these extra particles deposit in the TPC was studied in detail, and an improved tracking algorithm was developed that identifies and corrects for various types of migrating events using their distinctive energy signatures. Finally, the performance of this new algorithm was compared to the other standard trackers using Monte Carlo simulation data, where it produced significantly more accurate results.

Chapter 10

SYSTEMATICS OVERVIEW

Recall that the muon capture rate on deuterium is expected to be approximately 400 s^{-1} , and the MuSun experiment is aiming for a precision of 1.5%. With an even split of statistical and systematic uncertainties, added in quadrature, the total systematic error budget from all sources is just over 4 s^{-1} . Because the lifetime method involves subtracting the positive and negative muon disappearance rates (both approximately $455,000\text{ s}^{-1}$) this corresponds to an uncertainty below 1 in 10^5 for the original rate measurement before subtraction.

To achieve this ambitious precision target, it is essential to form a detailed understanding of each potential source of systematic error. This chapter will give an overview of all major systematic effects that have been identified, with emphasis on the work performed by other collaboration members. More detailed discussions of specific effects that have been extensively studied by the author will follow in the subsequent chapters.

10.1 Stops on Other Nuclei

Ensuring that the incoming muons actually stop in the target deuterium gas is perhaps the most important task for the MuSun experiment. As mentioned in section 3.4, the capture rate on heavier nuclei scales roughly as Z^4 . Even a small fraction of captures on other nuclei can therefore produce a large change in the observed muon lifetime, which has been a major limiting factor for the precision of previous experiments. The hardware and design decisions aiming to minimize such captures have already been described in detail, so this section will focus mainly on efforts to quantify their effects.

10.1.1 Wall Stops

The most obvious source of captures on other nuclei is muons stopping directly in non-deuterium materials such as the beam pipe or deuterium vessel, collectively known as wall stops. The inclusion of the TPC virtually eliminates such events, but introduces the possibility of wall stops in the TPC structure itself. A fiducial cut is applied that selects muon stops that are well separated from any solid material, but tracker errors may allow a small fraction of wall stops to be mis-reconstructed inside the fiducial volume. An analysis using the neutron detectors to estimate the fraction of wall stops as a function of position is presented in chapter 11.

10.1.2 Chemical Purity




Figure 10.1: Gas chromatography measurements for R2014 (left) and R2015 (right). The system was not yet fully calibrated during R2014, resulting in relatively large uncertainties.

As explained in section 3.4.2 the experiment requires ultra-high gas purity, and the N_2 concentration in particular must be kept at a 1 ppb level. This is achieved by the CHUPS gas circulation and filtration system, but the residual impurity after filtration must be monitored and corrected for. During data collection the impurity levels were monitored with periodic gas chromatography measurements performed by the PNPI group. To achieve sub-ppb level accuracy, samples taken from the target volume are allowed to accumulate on an

adsorbent bed, increasing sensitivity by three orders of magnitude compared to conventional chromatography techniques[30]. The system required careful calibration to reach the design sensitivity, which was not fully achieved until R2015. These measurements were also a rather complicated manual process, resulting in relatively few measurement points through the production runs as shown in figure 10.1.




Figure 10.2: Left: Impurity capture recoil energy spectrum observed in R2013. Right: energy distribution of delayed pulses on the stop pad, showing the minimum between the electron scattering and ${}^3\text{He}$ peaks. A decay electron veto enhances the capture recoil signal.

An alternative approach has been developed to monitor the impurity level in-situ by detecting TPC energy deposition due to recoils from muon capture on nitrogen impurities:



These recoils produce an energy peak around 150 keV, which lines up nicely with the minimum between the electron and ${}^3\text{He}$ peaks produced by stops in deuterium as shown in figure

10.2. In clean data any signal will be extremely small, but it can be enhanced by comparing events with and without electrons because captures prevent the muon from decaying normally. The identification is challenging, but by carefully measuring the observed shapes the capture recoil contribution may be identified. This method offers increased precision, particularly for R2014, and also allows for continuous tracking of the impurity between the chromatography sample points.




Figure 10.3: Capture recoil impurity estimates for the R2016 impurity injection test data, calibrated to the gas chromatography measurements.

The capture recoil method is calibrated to match the gas chromatography measurements using the R2016 impurity injection test data, as shown in figure 10.3. There is also an important additional calibration step to estimate the contribution from neutron recoils after deuterium capture, which will be selected by the same cuts as the impurity captures. Unlike

the pulses from capture recoils or fusions which are confined to the stop pad, neutron scattering may occur at any location in the TPC. The neutron pulses from the rest of the TPC may be used to subtract the neutron background component on the stop pad itself, reducing the signal in the production data by roughly a third. The final impurity estimates for R2014 and R2015 are 1.30 ± 0.14 and 1.03 ± 0.13 ppb, respectively.

The R2013 dataset reached extremely high impurity levels above 1000 ppb, producing a large shift in the measured decay rate. Analysis of this dataset indicates that nitrogen impurities produce a rate shift of about $4 \text{ s}^{-1}/\text{ppb}$. Applying this scaling factor to the impurity estimates above predicts shifts of $5.2 \pm 0.6 \text{ s}^{-1}$ for R2014 and $4.1 \pm 0.5 \text{ s}^{-1}$ for R2015.

10.1.3 Isotopic Purity

Protium (^1H) impurities in the target gas can combine with a muonic deuterium to form a $pd\mu$ molecule. These molecules undergo a different set of fusion processes than those for $dd\mu$, as shown in figure 10.4. The muon may stick to the resulting ^3He , slightly adding to the $\mu^3\text{He}$ background, but the main concern is the case where the muon is unbound. Unlike the fusion reactions in pure deuterium where any emitted muons are simply recycled, in this case the resulting "Alvarez" muon is ejected with an energy of 5.3 MeV[7]. At the target gas density this corresponds to a range of 26 cm, more than twice the total length of the TPC, so these muons leave the target volume and most likely become wall stops.

These fusion reactions have been studied extensively and the reaction rates in table 10.1 are well known. MuSun's sensitivity to the isotopic purity has been estimated by including a 100 ppm concentration of protium in the kinetics model, which predicted a 5 s^{-1} shift in the observed decay rate. The gas chromatography measurements taken throughout the production runs reported lower concentrations in the range of 16-20 ppm, corresponding to a less than 1 s^{-1} correction.




Figure 10.4: Fusion processes in an isotopic admixture of protium and deuterium gas. The "Alvarez" muon emitted by these reactions has a range of ~ 26 cm.

Process	Symbol	Rate (10^6 s^{-1})
Molecular Formation	$\lambda_{p\mu d}$	5.6
Total Fusion	λ_f	0.3
Fusion, γ Emitted	$\lambda_{f\gamma}$	$0.85 \lambda_f$
Fusion, μ Emitted	$\lambda_{f\mu}$	$0.15 \lambda_f$

Table 10.1: Rates for $p\mu d$ molecular formation and fusion processes, reproduced from [17].

10.2 TPC Track Interference

TPC track interference has been mentioned several times already, as it has been one of the most difficult systematic effects to account for. Because a fiducial volume cut is required to eliminate wall stops, tracking errors may result in events with true stop positions inside the cut window being reconstructed outside, or vice versa. See chapter 9 for a detailed discussion of the TPC tracking and the strategies to minimize mis-reconstructions. These errors are time dependent and can have a large effect on the observed muon disappearance rate.

Fusion interference refers to the case where fusion products are added to the track, and biases the electron distribution towards later times due to the requirement that the fusion has time to occur before the muon decays. It may be quantified either by detecting fusion

products in the TPC or by observing its characteristic distortion of the lifetime histograms. Chapter 13 will discuss corrections based on both of these methods, as well as a simplified model of the tracker errors.

Electron interference occurs when the decay electron deposits enough energy in the TPC to change the reconstructed stop position. In rare cases the electron energy may be detected in isolation, but typically it must overlap with a pre-existing muon pulse and increase its energy above either the TPC pulse-finding threshold or some tracking cut. Chapter 14 will discuss modeling approaches for these events, as well as a method for estimating their effect by comparing decay electrons emitted in different directions.

10.3 Beam Backgrounds

A third broad category of systematic effect are those produced by additional particles in the beam. The lifetime method requires plotting the decay electron time relative to the muon entrance time, but if there are multiple muons or electrons detected it becomes unclear which should be paired together. A constant background term would not be an issue, but several beam background effects are time dependent and could interfere with the lifetime fits.

10.3.1 Muon On Request

Recall that the MORE system includes a kicker that deflects the beam after a muon entrance is detected. Because of the signal travel time and muon time of flight delays, the beam remains at the full rate r for the first 800 ns before dropping to the kicked rate. While the average beam rate within the kicker window is fully determined by the MORE system, the spacing between events is randomly determined by the delay before the next muon is detected by the μ SC. The muon entrances obey Poisson statistics so this results in an exponential distribution in the delay, which in turn produces an exponential change in the beam background outside the kicker window depending on the probability of falling within the kicker window of the next event. The probability of the kicker being inactive at a given

time therefore becomes

$$P_b(t) = \begin{cases} e^{rt} & t < 0 \\ 1 & 0 \leq t < t_k, \\ 0 & t > t_k \end{cases} \quad (10.2)$$

Where t_k is the time of the kicker step and the muon entrance occurs at $t = 0$.

One might expect the beam background to scale with P_b , but while this is true in general it is not the case for muons "seen" by the μ SC. Because the μ SC triggers the kicker and defines the event window, it would be a contradiction for it to detect any background muons between two events since then that muon should have been identified as the next entrance. A general equation for the beam background is therefore given by

$$B(t) = \begin{cases} B_k + (sB_0 - B_k)e^{rt} & t < 0 \\ B_0 & 0 \leq t < t_k, \\ B_k & t > t_k \end{cases} \quad (10.3)$$

where B_k and B_0 are the background detection rates with the kicker on or off, and s is 0 for muons seen by the μ SC and 1 for unseen muons.

Backgrounds seen in other detectors typically have a similar form to the the unseen muons, since particles such as beam electrons are produced upstream in the beam pipe and have no effect on the μ SC. Also, in most datasets the beam electron background actually increases when the beam is deflected, resulting in an upward kicker step as shown in figure 10.5. There are a number of other complications that have been omitted from this simple model, for a more detailed discussion of the kicker behavior and its effects on backgrounds seen in the various MuSun detector systems see appendix B.

The kicker primarily causes problems for the lifetime measurement only when the fit range includes the step itself. Thus, while in some cases it is beneficial to explicitly model the step, the standard fitting procedure simply uses a start time of 1 μ s to exclude the step and focus on the region where the background is approximately flat. After the kicker step




Figure 10.5: Beam background seen in the R2015 μ CL data.

there may be a small exponential tail coming from the decay of unseen background muons, but this is estimated to cause less than a 0.25 s^{-1} change in the fitted decay rate.

10.3.2 Muon Pileup

Muon pileup refers to events where multiple muons enter the detector in the same event window. Incorrectly comparing decay electrons produced by the second muon to the time of the first muon entrance will result in invalid measurements, potentially interfering with the lifetime estimate. If the beam background were uncorrelated with the muon entrance then the extra decay electrons would simply add an additional constant background term, but due to the MORE system the large majority of extra muons will enter within the first 800 ns. Pileup events where both muons stop in the deuterium gas have relatively little effect on the lifetime, but if the second muon entrance is a wall stop this can introduce a fast decay component even when using the TPC to select clean stops in the fiducial volume.

As mentioned in section 7.6, several software pileup protection cuts are applied to eliminate events with multiple muon entrances. An independent estimate of the pileup protection




Figure 10.6: Muon tracks observed in the TPC for R2015 μ CL events, with no pileup protection (left) or full pileup protection (right). Red lines show fits to a model of the beam background. The cuts reduce the number of pileup events by more than 5 orders of magnitude.

efficiency is determined by checking for muon tracks in the TPC during μ CL events, as shown in figure 10.6. With no pileup protection applied there are a large number of tracks observed within the event window, but the pileup protection cuts reduce the number of tracks by more than 5 orders of magnitude for R2015. The observed time distribution of muon tracks depends on the MORE background and on the drift time in the TPC, and is fit to a model described in appendix B. The lack of μ CL data for some R2014 datasets makes the pileup protection efficiency estimate more difficult, Rachel Ryan performed an alternate analysis for those datasets looking for additional muons detected after the normal maximum drift time in the TPC[62].

By performing these fits with a variety of pileup protection cut definitions and comparing to the corresponding lifetime fits, we can produce a trend of muon decay rate shift as a function of the pileup muon probability. There is a trend towards faster decays as expected, but the high pileup protection efficiency results in a small affect on the fitted rates. A

zero-extrapolation estimates a shift of less than 1.5 s^{-1} for the full pileup protection cuts.

10.3.3 Background Slope

An additional background-related issue that has been identified more recently is that there appears to be a slight time-dependent component to the nominally flat background after the kicker step. This was originally observed by studies selecting events with specific numbers of electrons within the event window, which find differences of several hundred s^{-1} between single-electron and double-electron events. The exact source of this effect is not completely understood, but is likely caused by either instabilities in the kicker voltage or slow decays of secondary excited particles produced by wall stops. Several techniques for studying the effect are explained in chapter 12, along with some discussion of the possible sources and the

10.4 Electron Detectors

Finally, hardware problems with the electron detectors could distort the observed time distribution. For an exponential decay signal small shifts in time are equivalent to a small change in signal size, so offsets such as the WFD clock ambiguity or the time smearing mentioned previously have minimal effect on the lifetime. However, it is critical that the electron detectors have a constant detection efficiency as a function of time. Effects such as afterpulsing or deadtimes after an electron detection could affect the probability of seeing additional electrons at later times, thereby skewing the fitted decay rate.

10.4.1 eSC Afterpulsing

As mentioned in section 8.2.1, the individual PMTs in the eSC exhibit significant afterpulsing. This problem is largely eliminated by requiring 4-fold coincidences of all tubes on a given eSC segment, but any remaining afterpulsing would produce additional electrons at late times and skew the fits to longer lifetimes. Afterpulsing always occurs in the same eSC segment as a previous real pulse, so the size of this effect may be estimated by comparing

events with two electrons in the same segment to events with two hits in different segments, as shown in figure 10.7.

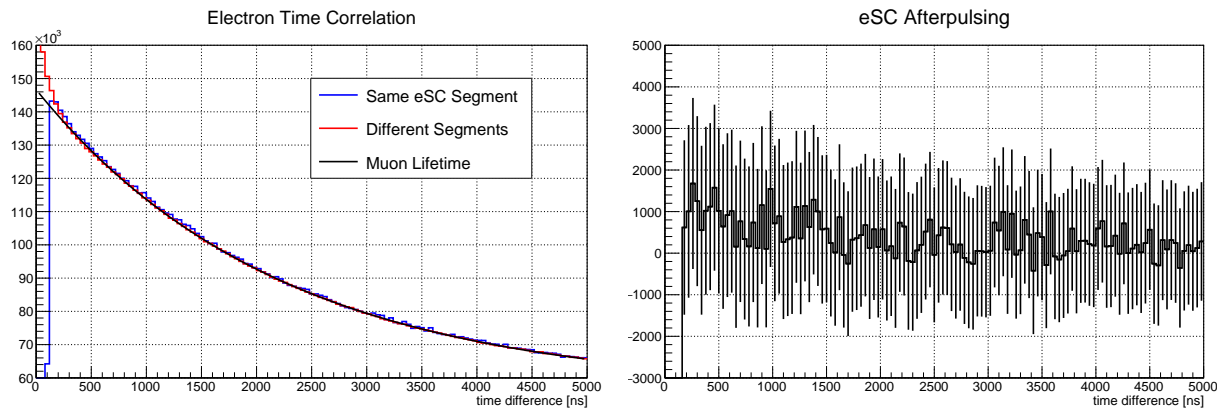


Figure 10.7: Left: electron time correlations for hits in the same eSC segment or different segments, normalized to times after $5 \mu\text{s}$. The expected muon lifetime curve is shown for comparison, and while deadtimes and coincidences produce different behavior at early times all three curves agree well after about 300 ns. Right: The difference between the two histograms on the left, possibly showing a very slight afterpulsing effect.

There may be a very small amount of afterpulsing, but it is not statistically significant. Performing lifetime fits while selecting events with two electrons in the same or different segments results in rate shifts of 158 ± 414 and $297 \pm 385 \text{ s}^{-1}$ for R2014 and R2015, respectively. Only 4.4% of R2014 and 3.2% of R2015 electrons occur in pairs, and of these only 1/16 share the same eSC segment, so scaling appropriately the expected rate shift for normal events becomes 0.44 ± 1.14 for R2014 and $0.59 \pm 0.77 \text{ s}^{-1}$ for R2015. Combining both production runs yields an estimated rate shift of $0.54 \pm 0.64 \text{ s}^{-1}$ due to afterpulsing.

This test also demonstrates that the background slope mentioned above is not caused by this sort of afterpulsing effect. If that were the case then the fits with electrons in different segments should be unaffected, while the electrons sharing a segment would show a large.

Instead the different eSC segment configurations make relatively little difference, indicating the majority of the effect really is caused by the background.

10.4.2 *ePC Noise*

When using full electron tracks, ePC noise can also distort the muon lifetime. This is slightly different than afterpulsing in the eSC, since the eSC is always used as the primary source of timing information. However, the ePC tracking can cause a distortion in the muon lifetime if it has a time-dependent efficiency, causing real electrons to be rejected when full tracks are required. The ePCs suffer from a variety of data quality problems, so achieving a constant detection efficiency is very difficult. Appendix A discusses a major rewrite of the electron tracking software aiming to increase the reliability of the system.

After these improvements the full tracks show much better consistency with the simpler eSC-only electron definition, but there is still a shift of about 8 s^{-1} caused by using the ePCs. Because this shift is the result of ePC inefficiency it can more than double when applying strict tracking cuts, and the exact shift is complicated to model. The beam background effects mentioned above have been much better studied and should be easier to correct for than the ePC detector issues, so despite the tracking improvements the eSC-only electron definition will continue to be the standard for the final analysis.

However, while the full tracks are not suitable for high precision lifetime fits, for most other purposes a $\sim 10 \text{ s}^{-1}$ rate shift is a relatively minor concern. The directional information provided by the full tracks is useful for characterizing several other systematic effects, allowing one to isolate muon decay or beam electrons and to study the TPC electron interference as a function of electron emission angle. These detectors are thus very important for informing the analysis cuts and corrections applied to the final lifetime fits, despite not being included in the fits directly.

10.5 Conclusion

The major sources of systematic error in the MuSun experiment have all been studied extensively. Table 10.2 shows the expected rate shift and uncertainties associated with the effects focused on in this chapter. The following chapters will discuss the remaining systematics that were only briefly summarized above, including the wall stops, background slope, and fusion and electron interference in the TPC.

Effect		Shift (s^{-1})		Uncertainty (s^{-1})	
		R2014	R2015	R2014	R2015
Beam	Pileup	1.5		1.5	
	MORE	≈ 0		0.25	
Purity	Isotopic	0.8		0.2	
	Chemical	5.2	4.1	0.6	0.5
Electrons	Afterpulses	-0.54		0.64	
	Tracking*	10			

Table 10.2: Summary of the systematic effects discussed in this chapter. Note that the electron tracking uncertainty does not apply to the main eSC-only analysis.

Chapter 11

WALL STOPS

Ensuring that only muons that actually stop in deuterium are considered when measuring the muon lifetime is the most fundamental aspect of the MuSun experiment. The TPC is included in the experimental design for the express purpose of ensuring clean stops in the target gas, and is one of the major improvements over previous experiments. Many of the most challenging systematic effects to analyze have resulted as side effects of the TPC tracking, but those difficulties are a small price to pay for largely eliminating wall stops. However, even with the TPC it is possible for large-angle scattering and tracker errors to result in wall stops which are mis-reconstructed inside the fiducial volume, and it is important to quantify this effect.

The probability of mis-reconstructing the stop position drops rapidly with the distance from the true stop position. Requiring a reconstructed stop inside the TPC should thus prevent any muons from reaching the walls of the deuterium vessel, with the main concern becoming wall stops in the TPC structure itself. For this reason the TPC is constructed using primarily high-Z materials, so that any wall stops which do occur will disappear rapidly before the $1 \mu\text{s}$ fit start time due to the approximately Z^4 scaling of the capture rate. Because the capture rate in these materials is much higher than the free muon decay rate, it also means that the majority of wall stop muons will capture rather than decaying into electrons. Captures on high-Z materials typically eject up to 5 neutrons[50], making the neutron detectors an attractive choice for measuring the number of wall stops.

11.1 Neutron Detector Response

The neutron detectors are tanks filled with liquid scintillator, with the scintillation light detected by photomultiplier tubes. These are sensitive to gamma rays and charged particles as well as neutrons, but these particles interact in different ways. In low- Z materials gammas interact via the Compton effect, charge particles produce direct ionization, and neutrons deposit energy via binary collisions with atomic nuclei. These processes produce subtly different pulse shapes, which may be identified with pulse shape discrimination cuts as discussed in section 8.3.

Neutron scattering also produces less scintillation light compared to an electron of the same energy, as shown in figure 11.1. Energies measured by the neutron detectors are therefore reported as "electron equivalent (ee)" energies, which have a non-linear relationship with the true neutron energy. The MuSun experiment has used two different types of neutron detector. The Bicron detectors used until R2014 were sensitive to an energy range from about 0.2 to 2.5 MeVee, corresponding to neutron energies between 1 and 5 MeV. The larger DEMON detectors introduced in R2015 are sensitive to a higher range from about 0.5 to 5 MeVee, or neutron energies of 2 to 9 MeV.

The neutron also deposits only a portion of its energy, recall from chapter 5 that the recoil energy from a binary collision is

$$E_R = \frac{2A}{(1+A)^2} (1 - \cos \Theta) E_n.$$

In this case the neutron will typically be colliding with a hydrogen nucleus in the scintillator, so $A=1$ and the prefactor becomes $1/2$. Even a mono-energetic neutron source will thus be observed to have a smooth energy distribution ranging from 0 up to a cutoff at the actual neutron energy.

11.2 Neutron Sources

Although the neutron detectors are quite sensitive to wall stops, there are several other neutron sources that must be considered. Particularly near the center of the TPC the wall




Figure 11.1: Comparison of scintillator light output produced by neutrons and electrons. The blue dot indicates the maximum energy deposition of a mono-energetic 2.45 MeV fusion neutron.

stop contribution will be extremely low, with the observed signal dominated by reactions in the deuterium gas itself as well as beam backgrounds. To identify a small wall stop contribution among these other signals will require detailed understanding of all neutron sources, of which we identify five in total.

11.2.1 Muon Capture in Deuterium

First, there are of course the neutrons produced by nuclear muon capture in deuterium. The exact time dependence of the deuterium capture signal depends on the muon kinetics, shown in figure 11.2. The kinetics have been discussed previously, so only a brief reminder will be given here. Initially approximately 2/3 of the atoms will be in the quartet state, with a capture rate of only $\approx 10 \text{ s}^{-1}$ compared to the $\approx 400 \text{ s}^{-1}$ doublet capture rate. This causes the detection rate of capture neutrons to rise at early times as the quartet state is rapidly

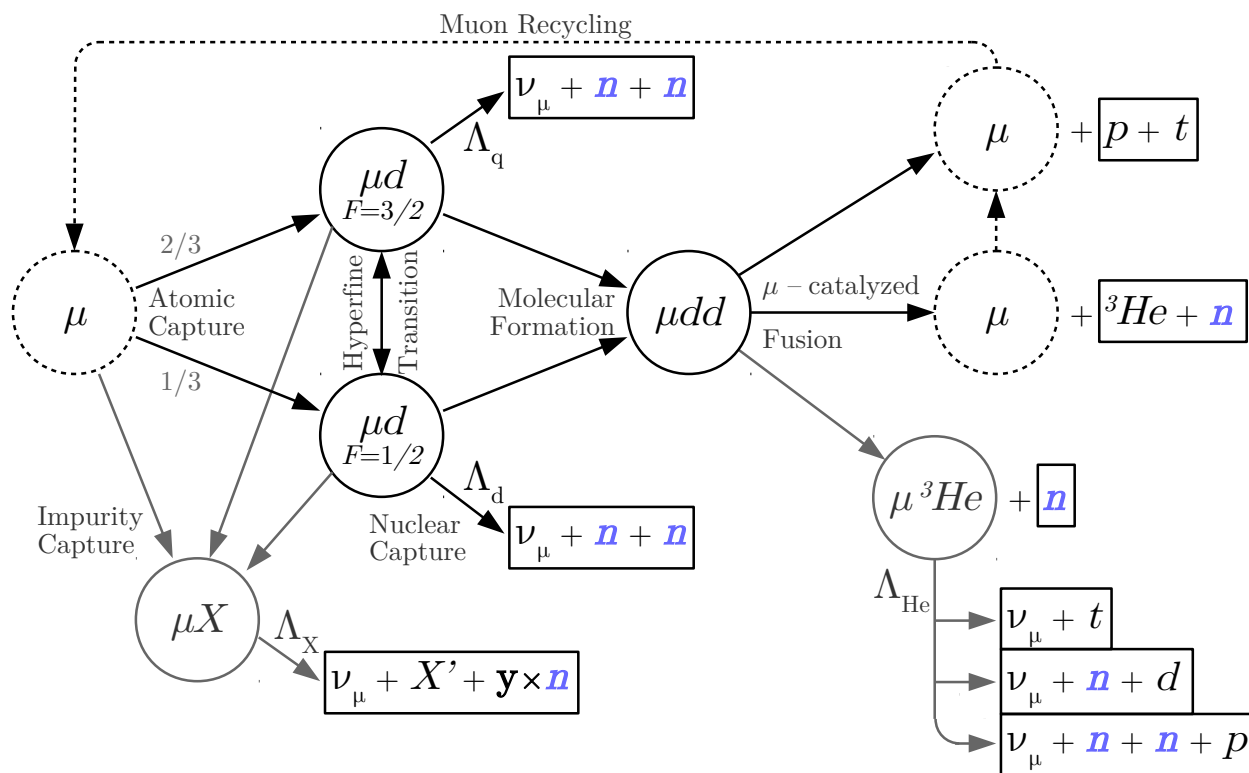


Figure 11.2: Muon kinetics in D_2 gas, with reactions that emit neutrons highlighted in blue.

depopulated, before falling again as the muons decay. The resulting double-exponential time distribution may be seen in figure 11.3. Incidentally this raises an interesting possibility of constraining the quartet capture rate using the shape of the time distribution, although this is quite difficult to do precisely[56].

Muon capture and muon decay are mutually exclusive, so capture events may be selected by rejecting events with an observed decay electron although the 70% solid angle coverage of the electron detectors limits the effectiveness of this cut. In principle the two coincident neutrons produced by these events would also be a distinct signature, but because they are typically emitted in approximately the same direction they are difficult to distinguish in practice. The three-body final state of the capture reaction also yields a wide range of neutron energies with few distinctive features, particularly after the additional smearing




Figure 11.3: Expected time distributions for capture and muon-catalyzed fusion in deuterium. The shape of both curves is governed by the transition from the quartet hyperfine state at early times to the doublet state at late times.

from the neutron detector response. Figure 11.4 shows the observed capture neutron energy spectrum, the apparent peak at low energies and the high-energy cutoff are both the result of the limited dynamic range of the detector and are not particularly meaningful.

11.2.2 Muon-Catalyzed Fusions

Fusion events have already been discussed extensively, although in this case it is the $n+{}^3\text{He}$ fusions that are important rather than the $p+t$ fusions that typically cause problems in the context of muon tracking. The fusion neutron time distribution has the same functional form as the capture neutron distribution, as they both depend directly on the hyperfine state populations. However, because the quartet rate has a much higher molecular formation rate than the doublet state the fusion time distribution falls sharply at early times rather than rising, as shown in figure 11.3

Fusion neutrons have several distinctive features. First, as discussed in section 9.2 the energy deposited in the TPC by the ${}^3\text{He}$ or $\mu^3\text{He}$ can be selected using E_0 vs E_1 cuts. There is some overlap between the ${}^3\text{He}$ band and the normal muon stop band, but if the overlap




Figure 11.4: Comparison of observed fusion and capture neutron energies, identified using TPC energy cuts. To best show the structure of the fusion energy distribution data from R2014 NU3 is shown, as it has the lowest energy threshold.

region is omitted it is possible to obtain clean samples of both fusion and non-fusion events. The $n+{}^3\text{He}$ fusion reaction also leads to a two-body final state with characteristic 2.45 MeV neutrons, resulting in a sharp cutoff in the energy deposition as shown in figure 11.4. Finally, because the muon remains in the TPC after the fusion there is typically a decay electron produced after the fusion neutron. Requiring an electron detection in a short time window following the neutron is an effective way to identify fusions, or a looser check for an electron anywhere in the event window may be used to avoid introducing a time dependence to the background.

11.2.3 Wall Stops

Wall stop neutrons are very similar to those from muon capture on deuterium, with no associated electron and energies spanning the full dynamic range of the MuSun detectors. Captures on different materials are distinguished mainly by their capture rates, with stops in

silver or tungsten having lifetimes of approximately 87 and 78 ns respectively[64]. Relaxing the TPC stop requirement also permits stops in the aluminum vessel walls or the steel front flange, with lifetimes on the order of 900 ns and 200 ns. A simple exponential decay appears to be a good match to the neutron time distribution for capture on a single element, although note that steel contains nickel and chromium in addition to iron.

It is also possible for muons to initially stop in the deuterium gas and then diffuse to the walls. The amount of diffusion may be estimated by analyzing the impact parameter of the decay electron tracks as a function of time, and is expected to be a small effect. This analysis will focus solely on the direct wall stop case.

11.2.4 *Photo-Neutrons*

Decay electrons may also undergo secondary reactions resulting in the production of additional neutrons. This most commonly occurs in a two-step process where the electron first emits high-energy bremsstrahlung gamma rays, which are then absorbed and emit neutrons via photo-nuclear reactions. The resulting photo-neutrons will be nearly time coincident with the decay electron, but the electron will not necessarily be detected near the triggered neutron detector. If the electron escapes through the ends of the eSC or is otherwise missed by the detectors then these events can appear similar to capture neutrons. The probability of producing a photo-neutron depends on the amount of material in the path of the electron track, resulting in a larger signal in the bottom detectors as they happen to be directly aligned with the TPC support structure.

The photo-neutron signal is quite small and difficult to measure directly, but it is an important correction when studying the rare capture neutrons. Because they are coincident with the decay electrons the signal simply decays exponentially, so an admixture of photo-neutrons in the capture spectra will result in an underestimation of the size of the fast component. This may be calibrated by comparing to the data for positive muons, which produce photo-neutrons from decay electrons but do not produce nuclear muon capture or fusion events.

11.2.5 Beam Background

Finally, the neutron detectors have high background levels relative to the size of the signals. Some of these are due to constant background sources such as cosmic rays, but the majority are capture neutrons produced by beam background muons. As such, they once again have a wide range of energies that cannot be distinguished from other events. The time distribution of the beam background is shown in figure 11.5, and follows the standard form discussed in appendix B.




Figure 11.5: Neutron background, isolated by using a TPC energy cut to select non-fusion events but requiring a decay electron detection to minimize captures. Some residual capture component remains and is included in the fit.

The most likely source of these backgrounds is muons stopping in the lead collimator included in the entrance detector stack. The fast decay after the kick is consistent with the muon lifetime in lead, and the fact that the signal drops after the kick also suggests the

source is near the entrance detectors. It may be interesting to revisit this analysis with a μ SCa pileup protection cut applied, as that may significantly reduce the number of beam muons stopping in the lead collimator.

11.3 Gamma Sources

The neutron detectors detect significantly more gamma rays than they do real neutrons, and typically these are considered backgrounds and rejected. However, in this case the gamma rays also provide distinctive wall stop signals and should be included in the analysis. We identify four classes of gamma signal:

11.3.1 Bremsstrahlung

A large amount of bremsstrahlung radiation is produced by the decay electrons scattering as they exit the detector. As with the photo-neutrons this results in a strong directional dependence in the bremsstrahlung signal, with more than a factor of two difference between the upper and lower detectors. Bremsstrahlung is the main form of gamma signal from stops in deuterium.

11.3.2 Nuclear Capture

After muon capture on a high- Z material, the nucleus is left in an excited state. In most cases it will quickly return to the ground state by emitting gamma rays, resulting in a gamma signal following the same exponential decay curve as the corresponding wall stop neutrons and decay electron signals. In principle these gamma rays should form distinct transition lines, but the complicated structure of high- Z nuclei produces a large number of closely spaced transitions that cannot be distinguished by the MuSun detectors.

11.3.3 Muonic X-Rays

Recall that during the atomic capture process the muon is initially at a very high atomic energy level, and rapidly cascades down to the ground state. In the process it emits characteristic X-rays with well-defined transition energies, and while the signal from deuterium is below the dynamic range of the neutron detectors the wall stops in high-Z materials are clearly visible. The cascade is faster than the time resolution of the neutron detectors, and appears as a single pulse coincident with the muon entrance time. The shape is well approximated by a rectangular pulse, with the width most likely being an artifact of the timing uncertainty of the digitizers.




Figure 11.6: Atomic capture energy spectra from R2016 calibration data.

The atomic transition lines are spaced widely enough that distinct peaks may be observed in the energy spectra despite the poor energy resolution of the neutron detectors. Figure 11.6

shows energy spectra for atomic capture on TPC materials obtained during the R2016 run. These distinct peaks offer an additional way to identify the stop material, although this is only usable near the walls where there is a large signal. In principle stops in deuterium should produce a similar signal, but the ground state of muonic deuterium only has an energy of 2.65 keV so these photons will be well below the minimum energy threshold for the neutron detectors.

Stops in Macor are difficult to identify using capture neutrons, but as shown in figure 11.6 they do produce a small prompt atomic capture signal. Comparing the sizes of the atomic and nuclear capture signals may therefore be a way to identify stops in Macor or other intermediate- Z materials. In practice there should be very few stops in Macor, and the analysis will primarily focus on stops in silver and tungsten.

11.3.4 Beam Background

Finally, there is a large beam background consisting mainly of bremsstrahlung radiation from particles in the beam. Unlike the neutron background the gamma ray background increases when the beam is deflected, as a side effect of the increased electron background. Calibrating with the clock data, we extract the background shape seen in figure 11.7. The gamma background contains a strong accelerator RF oscillation component, which must be included in the background model to maintain good timing resolution and to avoid bias when estimating the size of the atomic capture peak.

11.4 Analysis Procedure

Table 11.1 summarizes the key properties of the signals discussed in the previous sections. In the context of the wall stop analysis the relevant signals are the fast exponential decay signals produced by nuclear capture as well as the prompt atomic capture gamma ray peak. The goal of the wall stop analysis is to study how the size of these signals varies as a function of reconstructed stop position.

We are interested in stops that are mis-reconstructed inside the TPC but actually stopped




Figure 11.7: Gamma background signal for R2015 data.

outside, so the TPC energy cuts are not usable in this analysis. Instead a minimum energy threshold is used to eliminate the majority of the fusion neutrons, at the cost of a significant reduction to the capture signal. An electron veto helps to suppress fusions, and eliminates the majority of the bremsstrahlung and photo-neutron signals. However, it is not possible to devise a cut to distinguish between wall stops and deuterium captures or to eliminate the beam background.

The analysis relies primarily on the distinctive time distributions of the wall stop signals to identify them. Figure 11.8 shows the time distributions for various signals emitted from wall stop materials. The atomic capture signals are clearly distinguishable from the much slower muon lifetime, and the atomic capture peak is only about 10 ns long and unlike any signal expected for clean stops in deuterium. By fitting the observed time distributions with

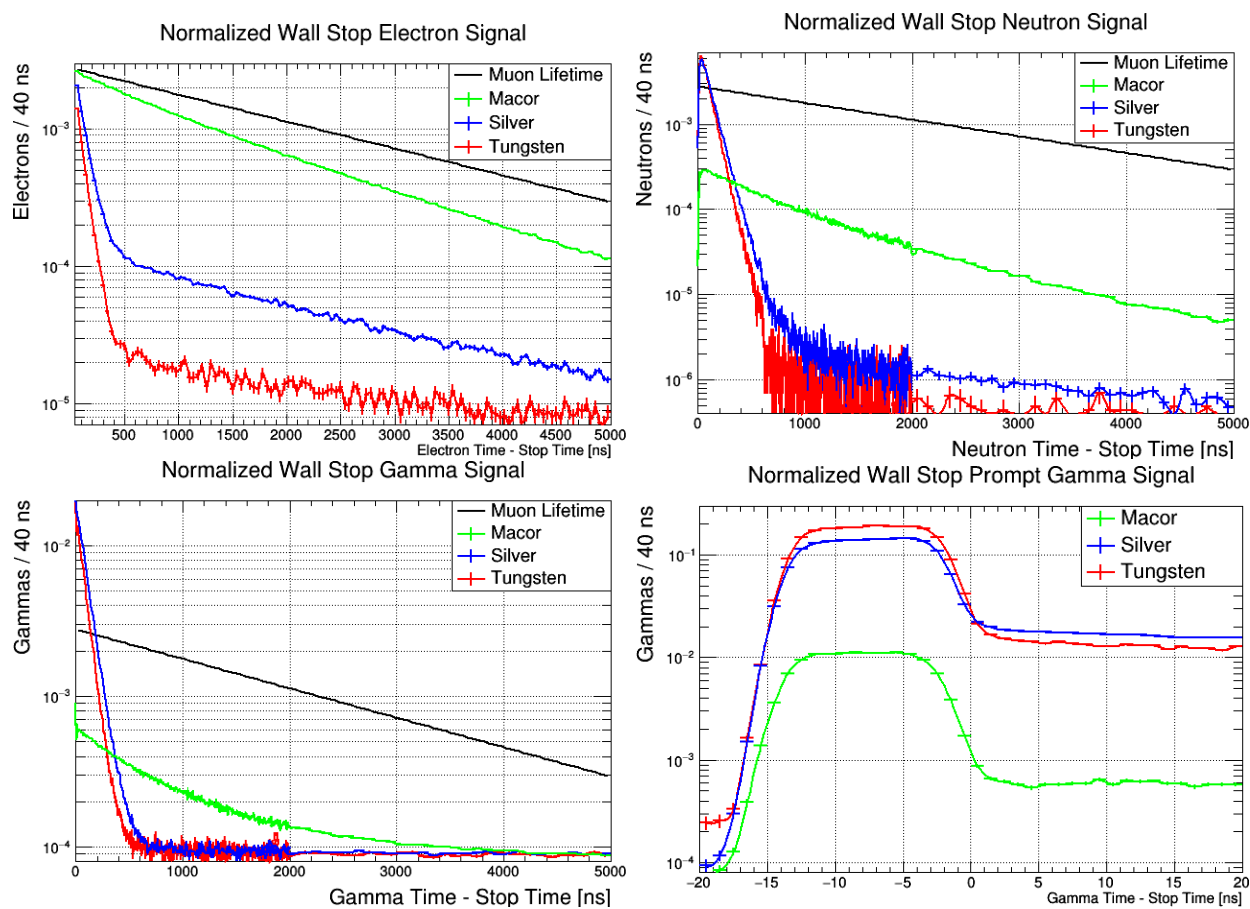


Figure 11.8: Panels 1, 2, and 3 show observed electron, neutron, and gamma signals for various target materials measured during the R2016 systematics testing, with an ideal muon decay curve shown for comparison. The slow decays seen in the high-Z samples are backgrounds from muon stops in the entrance scintillator. Panel 4 shows a close-up of the gamma signal to highlight the prompt atomic capture peak. All signals are normalized by the number of detected events.

Particle	Signal	Source	Energy	Electron	Properties
Neutron	D Capture	D stops	<2.54 MeV	none	2 neutrons
	³ He Fusion	D stops		delayed	TPC energy
	High-Z Capture	wall stops		none	0-5 neutrons
	Photo-Neutrons	scattering		prompt	
	Background	beam			
Gamma	Bremsstrahlung	scattering	peaks	prompt	$t_N = 0$
	Nuclear Capture	wall stops		none	
	Atomic Capture	wall stops		none	
	Background	beam			

Table 11.1: Summary of signals in the neutron detectors.

a model for the relevant signals we can extract an estimate for the number of wall stops. Repeating the fit with various muon tracking cuts then results in a measurement of the wall stop contribution as a function of reconstructed stop position.

11.4.1 Signal Modeling

The lifetime fits typically used by the MuSun analysis avoid a great deal of complication by starting at 1 μ s and using a coarse time binning. The fits to the wall stop signals are different, requiring high time resolution and directly fitting the area around the muon entrance time that is normally avoided. Normally neglected effects such as the detector timing resolution and time of flight delays therefore become important, and must be accounted for in the fit. For this reason an entirely separate fitting procedure was developed for the neutron analysis, based on the sophisticated RooFit framework included in ROOT. Each of the signals listed above are represented with a probability density function (PDF), which are then combined to form a detailed model of the system.

The detector timing resolution is accounted for by convolving the predicted time distribution for each signal with a Gaussian representing the error in the time reconstruction. This turns the various exponential decay curves into a functions known as exponentially modified Gaussians. The background is already sufficiently smooth so the convolution is omitted in this case. The importance of including this smoothing is clear in figure 11.9, where the prompt atomic capture peak directly overlaps the onset of the capture signals. The goal will be to identify much smaller wall stop signals in the interior of the TPC, which could easily be influenced by inaccurately modeling the shapes near the entrance time.




Figure 11.9: Neutron (left) and gamma (right) signals for stops within 5 mm of the anode plane. Individual fit components are shown, with signal referring to muon stops in deuterium, wall captures referring to nuclear captures, and the prompt peak showing muonic X-rays from atomic capture in the walls. Note that the time axis covers a range of less than 200 ns, decay of the signal component from deuterium is not apparent.

For the gamma signals time of flight delays can shift the observed times within about a 1.5 ns range, which is small enough to have no significant effect even on the sharp atomic capture peak. For the neutron detectors the time of flight makes a larger difference due to the relatively slow speed of the neutrons, with delays of up to 5 ns possible for a 2.45 MeV neutron. Furthermore, since the neutron speed depends on its true energy which is not accurately measured by the detectors, there is always an unavoidable broadening of the time distribution even if the stop position and observed energy are controlled. Fortunately

the neutron signal fits only need to identify the nuclear capture signal and not the prompt peak, so this extra timing uncertainty is not a major concern.

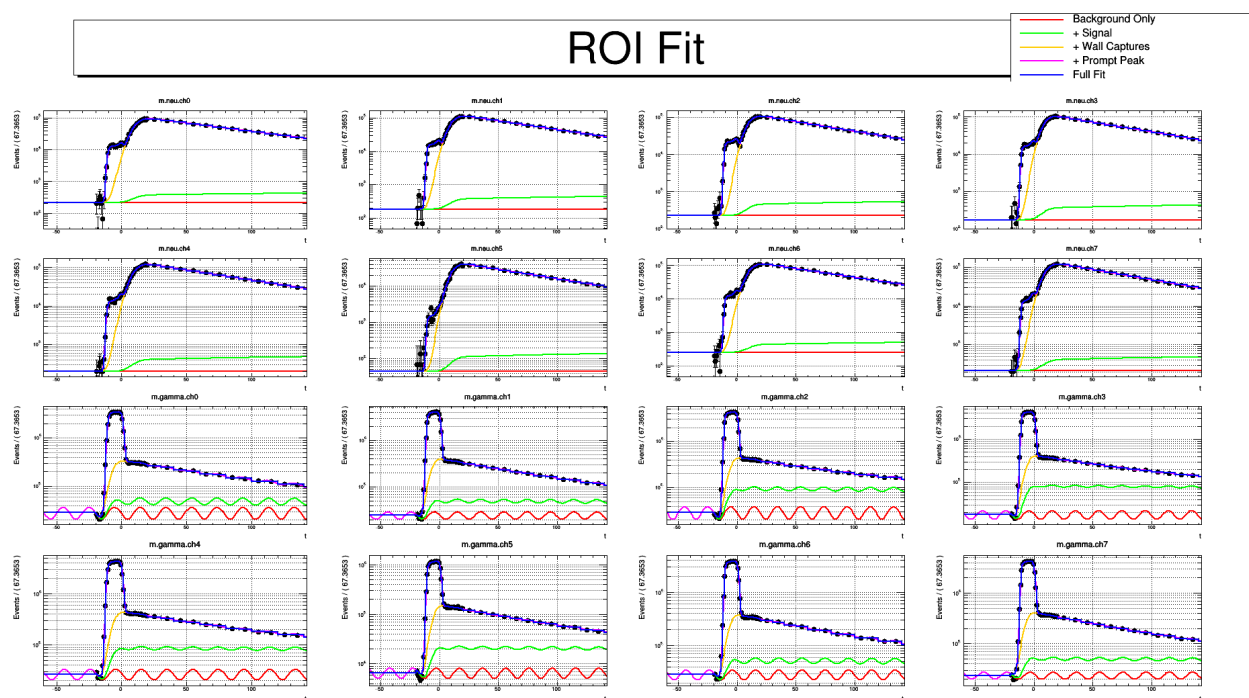


Figure 11.10: Combined fit to all detector channels for R2015 stops anywhere in the TPC. Here the individual fit components are not shown separately but instead added to the combined signal one by one, building up to the full fit function.

Because the behavior of each of the eight neutron detectors is somewhat different, simply adding the signals from all detectors would also result in an unwanted smearing effect. Instead a separate model is created for neutron and gamma signals from each detector, for a total of sixteen channels. Using the RooFit system these histograms may all be fit simultaneously, as shown in figure 11.10.

11.4.2 Calibrations

The full model clearly has far too many parameters to fit reliably, and most of these must be fixed before proceeding with the analysis. Ideally this would leave only two parameters representing the number of stops in deuterium and in the walls, although in practice additional parameters are needed to produce good fits. Global properties such as the number of estimated stops or the beam and kinetics parameters are therefore shared between all detectors, with individual channels only differing by their specific detector response parameters. The relative signal amplitudes in each detector may also be fixed by estimating the solid angle coverage based on the observed muon stop distribution, although individual amplitudes are needed to account for the directional dependence of the bremsstrahlung radiation.




Figure 11.11: Observed electron time distribution for stops at $Y < 5$ mm.

The remaining free parameters are fixed by a series of high-statistics calibration fits performed before the main wall stop analysis. These include fits to stops in the center of the TPC to isolate the deuterium signals, and to slices at high and low Y to calibrate the wall stop signals. The selection of extreme Y positions also produces a wall stop contribution large enough to be clearly visible in the electron detectors, as shown in figure 11.11. This

is used to relate the observed neutron signals to the number of wall stops, by fitting the electron time distribution with a double-exponential of the form

$$f(t) = N_d \lambda_\mu e^{-\lambda_\mu t} + N_w \lambda_\mu e^{-(\lambda_\mu + \lambda_w)t} + BG, \quad (11.1)$$

Where N_d is the number of stops in deuterium, N_w is the number of stops in the wall, and λ_w is the capture rate in the wall material. To account for the differences between wall stops at the top or bottom of the TPC which are mainly silver and tungsten respectively, the wall stop fit parameters are set using either the high or low Y calibrations as appropriate.

After these calibrations the majority of the parameters in the model are kept fixed for all subsequent fits. There are 11 free parameters remaining, representing the number of stops in deuterium and the walls, the amplitude of the background, and 8 parameters controlling the bremsstrahlung amplitude in each detector. This is sufficient to produce fairly reliable fits to lower statistics samples, particularly since the relatively slowly varying background and deuterium signal components are not strongly correlated to the fast wall stop signals.

11.4.3 Lifetime Effect

Before discussing the results, it is important to establish the safe levels for wall stops. The relation between wall stop fraction and the fitted lifetime is difficult to directly observe, as tracker interference is by far the dominant effect in any stop position scans. Instead, a simple simulation is used to generate time distributions for various wall stop fractions and then measure their effect on the standard MuSun fits. Figure 11.12 shows the estimated fraction of wall stops in various materials required to produce a $1s^{-1}$ rate shift.

The standard MuSun fits starting at $1 \mu s$ are nearly unaffected by high-Z materials unless they outnumber the number of stops in the deuterium itself. This clearly demonstrates the benefit of using these materials in the TPC, as stops in the aluminum vessel must instead be controlled to the 10^{-5} level regardless of the fit start time. The main concern is thus not actually the high-Z inner surfaces of the TPC, but instead the iron or Macor components that make up parts of the TPC frame. These are not directly measured by the wall stop




Figure 11.12: Wall stop fraction required to produce a $1s^{-1}$ rate shift, as a function of fit start time. The standard MuSun fits starting at $1\ \mu s$ are nearly unaffected by stops in the high-Z TPC wall materials, in contrast to stops in the aluminum vessel that must be controlled to the 10^{-5} level.

analysis, but should be significantly lower than the observed wall stop fraction as the muons must exit the TPC diagonally to reach these components rather than simply hitting one of the faces. A wall stop fraction of 10^{-3} is therefore a good conservative target, as it should ensure even these intermediate-Z materials are well below the maximum safe level.

With the approximately 10^{-4} wall stop fraction observed in the fiducial volume a much earlier start time of around 100 ns would be needed before any significant effect would be expected. The TPC does include some steel components mainly in the bottom frame which are a larger concern, but a muon would need to leave the TPC in both X and Y the number of stops should be significantly lower than the observed stops in high-Z materials and certainly far below the 10^{-3} level estimated to be safe for iron. Similarly stops in the Macor posts at the corners of the TPC will be significantly less common than X or Z wall stops individually, and are also not a concern.

11.5 Results

To assess how different fiducial volume cuts would affect the wall stop signal we will perform scans over the reconstructed muon stop position in the TPC. Note that the final rate shift depends on the total number of wall stops integrated over the full fiducial volume, so while the differential scans can show large effects near the edges they are still typically a small number of events in an absolute sense. There has also been interest in using cuts on the first TPC pad row to focus the muon stop distribution, as described in section 13.2, and fit results with and without these cuts will be shown.

11.5.1 Stop Y




Figure 11.13: Wall stop fraction as a function of reconstructed stop Y. A variable width binning scheme is used to capture the rapid changes near the walls. Black points indicate the result with first row cuts applied, while the red points include all events.

The Y position is the most critical to study, as the solid anode and cathode planes at the top and bottom of the TPC mean muons escaping vertically are almost certain to stop

in the walls. The Y axis is also the shortest dimension of the TPC, so there are significantly more losses in Y than in X or Z. A scan of the fraction of wall stop events as a function of reconstructed Y stop position is shown in figure 11.13. The low-Y calibrations are used for $Y < 40$ mm while the fits above that threshold use the high-Y calibrations, resulting in a slight discontinuity in the middle of the scan. The shape of Y scan is very reminiscent of the tracker performance testing using Monte Carlo truth shown in section 9.5. In both cases there is a rapid drop of more than two orders of magnitude over a distance of about 5 mm, followed by a more extended tail. At wall stop fractions around the 10^{-4} level the neutron fit begins to have difficulty picking out the small wall stop component from the other background signals, resulting in somewhat inconsistent results near the center of the TPC.

11.5.2 Stop X and Z




Figure 11.14: Wall stop fraction versus reconstructed stop X (left) and Z (right). Black points indicate the result with first row cuts applied, while the red points include all events.

Muons exiting the TPC in the X or Z directions must travel an additional 18 or 24 mm respectively to reach the walls. These walls are also a series of individual wires rather than a solid plane, so the majority of muons that do reach them will simply pass through. As a

result, while the final Z row and the X edge rows do include some punch-through muons, the wall stop fractions remain below the 10^{-3} level even for these differential measurements. The downward trend as a function of Z is the result of less accurate Y reconstruction for shorter tracks, as shown in figure 11.15. The first Z bin in particular contains two-pad tracks, which have only a single pulse in the road cluster and so cannot use the track slope to extrapolate to the correct Y stop position. These short tracks are not included in the standard fiducial cut, and starting with Z=2 the wall stop fraction is again at the 10^{-3} level.




Figure 11.15: Wall stop fraction as a function of both Y and Z. The trend seen in the Z scan is largely the result of less accurate Y reconstruction for shorter tracks.

11.6 Summary

This chapter described the various signals seen by the MuSun neutron detectors, as well as a sophisticated fitting procedure for identifying the distinctive wall stop signals. The wall stop fraction is estimated to be approximately 10^{-4} in the fiducial volume, far below the level necessary to produce a noticeable effect on the muon lifetime fits. The standard 15 to 56 mm fiducial Y cut is sufficient to minimize the wall stop component, and could potentially be relaxed by about 5 mm with little issue. The X and downstream Z cuts may actually be removed entirely, and this has become the standard configuration used in most analyses.

Chapter 12

TIME-DEPENDENT BACKGROUND

In addition to the well-understood MORE background described in the previous chapter, there appears to be a slight downward trend in the background within the fit window from 1 to 25 μ s. This chapter will describe this effect and methods to measure it, as well as several possible sources of the time dependence. Finally, the sensitivity of the lifetime measurement to such an effect will be discussed and a correction will be estimated.

12.1 *Electron Multiplicity*

In some cases it is useful to either suppress or enhance the background of the lifetime histograms, which may be achieved by cutting on the number of electrons detected in the event window. A single-electron (1e) cut removes any events with both signal and background electrons, so backgrounds can only appear when the decay electron is missed. A double-electron (2e) cut instead forces every event to have at least one background electron, since there should only be one decay electron.

The detection probability for a decay electron signal (S) is about 2/3 due to the limited solid angle coverage of the detectors. Since there is only one decay electron, ignoring pileup, the corresponding probability to miss the decay electron is simply 1-S. The background detection probability (B) is equal to the beam rate multiplied by the length of the event window, and is typically at the percent level. A given event may potentially include an arbitrarily high number of background electrons, with the probability of n electrons given by Poisson statistics:

$$P(n) = \frac{1}{n!} B^n e^{-B}. \quad (12.1)$$

With an electron multiplicity cut requiring n electrons, there may either be n backgrounds




Figure 12.1: Time distributions from R2015 events with different electron multiplicity cuts. Electrons are only counted in the window $(-600, 24400)$ ns, resulting in discontinuities at the endpoints. A 10 ns coincidence veto has been applied to the 2e histogram.

or 1 signal and $n-1$ background electrons. Table 12.1 shows the resulting signal and background levels, as well as their ratio. In the ratio many of the factors cancel, so it simply increases by one with each subsequent electron. Multiplicities higher than 2 lead to a large reduction in statistics.

12.1.1 Lifetime Shift

The fitted disappearance rate should be insensitive to adding a constant background, but in practice the 2e cut produces a shift of several hundred s^{-1} as shown in figure 12.2. This could be explained by either a small time-dependent component to the background or by correlations between the two electrons selected by the 2e cut. Some correlations are expected and this possibility will be discussed first, but the time-dependent background appears to be the primary issue and will be the main focus of this chapter.

Afterpulsing in the individual eSC PMTs is a known issue, but appears to be largely

#e	Signal	Background	BG / Sig
All	S	B	$\frac{B}{S}$
1	$S \times e^{-B}$	$(1 - S)B \times e^{-B}$	$(1 - S) \frac{B}{S}$
2	$S \times B e^{-B}$	$(S + (1 - S)B) \times B e^{-B}$	$1 + (1 - S) \frac{B}{S}$
n+1	$S \times \frac{1}{n!} B^n e^{-B}$	$(nS + (1 - S)B) \times \frac{1}{n!} B^n e^{-B}$	$n + (1 - S) \frac{B}{S}$

Table 12.1: Signal and background probabilities for different electron multiplicity cuts. S is approximately 0.7 due to the solid angle coverage of the detectors, while B varies but is typically about 0.02.

eliminated when requiring full 4-fold coincidences as discussed in section 10.4.1. To check if any residual afterpulsing is the cause of the observed lifetime shift, we can also eliminate them by selecting pairs of electrons in different eSC segments as shown in figure 12.4. The result is nearly identical to the fits including all 2e events, so this is not the problem. Crosstalk effects between the detectors would also show up in this analysis, and do not appear to be significant.

There are also physical processes that are expected to produce correlated signals in the eSC. The decay electron may either emit bremsstrahlung radiation or eject Møller electrons as it passes through the walls of the vessel. Cosmic rays also typically pass directly through the experiment and intersect different eSC segments at the entrance and exit points, producing two triggers. These events all produce nearly simultaneous signals as shown in figure 12.3, and may be removed with a 10 ns coincidence veto. None of these effects should have a time-dependent probability, and do not appear to contribute to the rate shift as shown in figure 12.4.




Figure 12.2: Fitted decay rates of 1e and 2e events, from R2014 (left) and R2015 (right) datasets. Rates are shown as shifts relative to the fit to all events.

12.2 Measurement

If the effect is truly the result of a time-dependent background, it should be verified by studying the background directly. Ideally one would characterize the background independently and use this information to derive a correction for the fitted rates, which could then be verified by applying it to the electron multiplicity cuts. In practice the small size of the effect makes this extremely difficult.

12.2.1 Slope Model

The background time dependence will be quantified by including a simple linear slope in the background model. This may seem somewhat arbitrary, but is justified for several reasons. Firstly, as the simplest possible model it introduces the fewest assumptions about the exact nature of the effect. The slope may also be thought of as the first term in a series expansion, and neglecting higher-order terms is reasonable for a sufficiently slowly varying signal. With




Figure 12.3: Time difference between electrons in multi-electron events. The slow component comes from the muon lifetime, while the large peak at $t=0$ is a combination of cosmic rays, secondary particles produced by the decay electron, and possibly electronic crosstalk effects.

limited statistics a single-parameter model also has the best hope of producing reliable fits.

Most importantly, the basic exponential lifetime fit only has three parameters: signal (N), background (B), and decay rate (R). N naturally depends mainly on the number of muon decays, while B is fixed by the average background at late times. This leaves R as the only remaining free parameter, so a single slope parameter should be sufficient to capture the majority of the effect on the rate. As shown in figure 12.5, small adjustments to either of these parameters have similar effects of changing the number of events seen in the middle of the fit range.

One might naively expect that an upward sloping background would increase the measured lifetime by adding more events at late times, but since the B parameter will largely compensate the background at late times it is in fact a downward slope that produces the desired effect. Because the background level may change significantly between datasets or when applying event selection cuts, it is helpful to measure the slope as a fraction of the




Figure 12.4: Fitted decay rates for double electron events with various selection criteria. Prompt coincidences behave similarly to single electrons and have small rate shifts, while otherwise the 2e events have similar events whether or not pulses on the same segment are included. Isolating pulses on the same segment may produce a slightly larger rate shift as a result of afterpulsing, but the effect is not significant.

total background. The equation for the background thus becomes

$$B(t) = B_0 [1 - dr(t - t_0)], \quad (12.2)$$

where dr is the relative slope and t_0 is the reference point for the 'true' background value, typically the end of the event window at $25 \mu\text{s}$. Because we are almost exclusively interested in negative slopes, dr is defined such that positive dr values indicate a downward slope. Although electron times are normally measured in ns, dr will be reported in units of μs^{-1} to better reflect the change in background over the length of an event. A rough estimate from the electron multiplicity fits indicates that dr values on the level of $10^{-4} \mu\text{s}^{-1}$ are to be expected. Section 12.4 will use simulated data to investigate to what extent the exact shape of the background influences the results.




Figure 12.5: Left: effects of varying the parameters in a standard exponential + background fit. The signal parameter N and the background parameter B primarily affect early and late times, respectively, while the rate parameter R affects intermediate times. Right: effect of adding a background slope to the fit, parameterized as $B \equiv B_0(1 + S(t - t_0))$, where S is the slope as a fraction of the constant background term B_0 . The line colors indicate the magnitude of the slope, while the sign of the slope is indicated by the arrows.

12.2.2 Muon Clock Fits

The muon clock is included in the experiment for the express purpose of studying the behavior of the background, and is the natural first choice for measuring this effect. In normal events the 2e cut with a prompt coincidence veto should typically select a combination of one decay and one background electron, while rejecting the coincident backgrounds due to cosmic rays. In the clock data there is no decay electron, so instead a 1e cut is applied to select events with one background electron and to once again eliminate the prompt coincidences. In principle $\approx 30\%$ of the 2e clock events should be added to reproduce the case where the 2e cut misses the decay electron and finds two backgrounds instead, but this is a percent level correction

and should have little effect. Furthermore, testing using the 1e cut is an additional check that this is a real slope and not related to unexpected correlations between the beam electrons. The resulting background time distributions are shown in figure 12.6.




Figure 12.6: Electron backgrounds observed in R2014 (left) and R2015 (right) for muon clock events. Linear fits find small but non-zero slope parameters (dr).

Linear fits to these histograms find non-zero downward slopes for both datasets, but the clock statistics are too low to extract much beyond the overall slope. The spike in background just after the kicker step is expected, see appendix B for a description of this effect. To ensure these slopes are real and not simply the result of random fluctuations or residual effects from the peak, the fit start time is scanned from 1 to 10 μs as shown in figure 12.7. This produces different behavior for the two production runs, with R2015 consistently finding slopes around the 10^{-4} level while R2014 starts significantly higher but trends towards lower slopes with increasing start time. The R2014 clock data is mainly from the end of the run, with nearly 2/3 coming from DSM5 which also had an unusually large

electron multiplicity discrepancy. It would seem that DSM5 in particular has a large slope, but this is not necessarily representative of R2014 as a whole.




Figure 12.7: Estimated background slopes from the μ CL data, as a function of fit start time. Note that nearby fits are highly correlated so the trend as a function of start time is more significant than it appears, and indicates a slight curvature in R2014.

In any case, these fits certainly support the existence of a time-dependent background component. The R2015 appears reasonably consistent with the linear slope approximation, while the trend in R2014 suggests some slight curvature. This difference may indicate that more than one effect contributes to the background slope; this will be discussed more in section 12.3 below. There is a small exponential tail following the kicker step, which is normally unimportant but may introduce a small bias to the background fits. To avoid any possible bias from this effect a fit start time of $2 \mu\text{s}$ will be used when comparing to other estimates of the slope.

12.2.3 Clean Track Analysis

A similar method of directly measuring the background slope uses normal muon entrances instead of the clock events. This significantly increases the statistics, but the presence of the decay electrons completely dwarfs any subtle changes in the background. The analysis therefore focuses mainly on suppressing the decay electron signal as much as possible.

The 2e cut already increases the background to signal ratio by nearly two orders of magnitude. Requiring a third electron would further increase the background, but this will dramatically reduce the statistics while only increasing the background by a factor of two. This analysis instead retains the 2e cut, but uses the ePC tracking information to identify high quality "clean tracks" which are likely to be the decay electron. The other electron in the event must then be background, and a new lifetime histogram is filled with only these background events. The signal to background ratio of the new histogram is given by the false positive rate of the clean track identification, while the statistics will scale with the decay electron detection efficiency.

There are two tracking parameters which have a large impact on the signal to background ratio. Impact parameter measures the minimum distance between an electron track and the estimated muon stop position, while $\cos(\theta)$ measures the Z component of the electron direction vector. Because most electrons originate either from the target or from the beamline, they are approximately radially symmetrical and these two parameters are sufficient to capture the majority of the dynamics. Figure 12.8 shows the distribution of signal and background electron tracks with respect to these parameters.

The muon decay electrons should ideally have zero impact parameter and be uniformly distributed with respect to $\cos(\theta)$. In practice electron scattering limits the tracking resolution to around 15 mm, while tracks with extreme $\cos(\theta)$ exit along the beam axis and cannot be seen by the cylindrical electron detectors. The background electrons instead tend to have much larger impact parameters, and predominantly travel in the downstream direction as expected for particles from the beam. The large arcing structures visible in the background

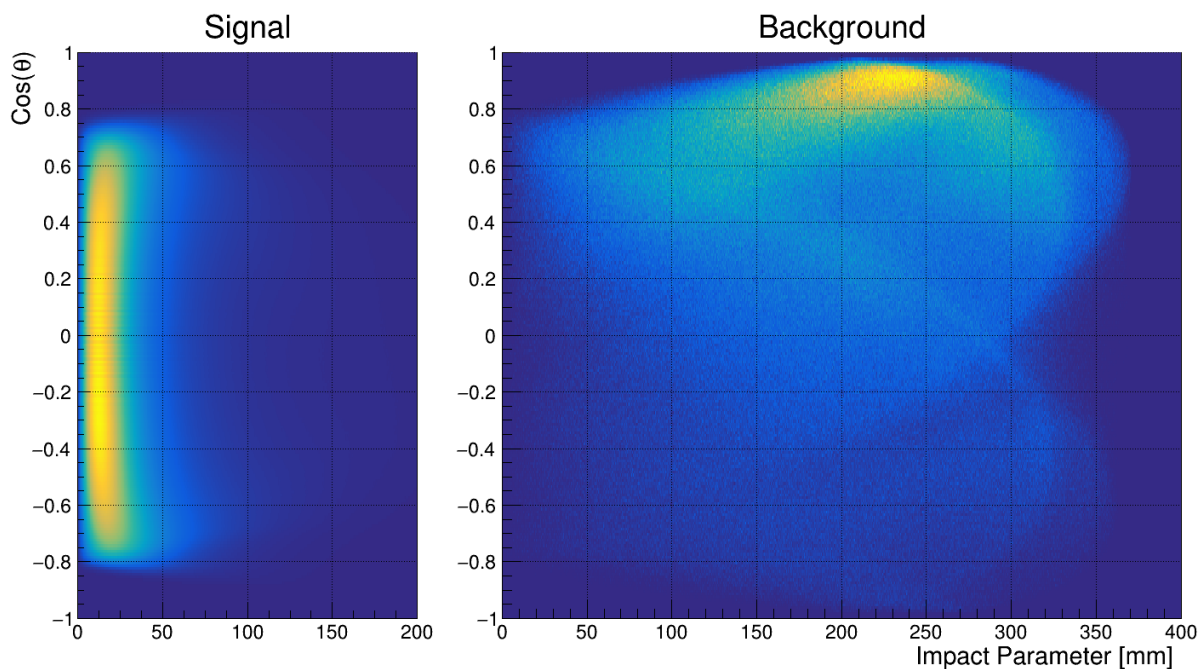


Figure 12.8: Impact parameter and track angle for the decay electron signal (left) and background (right) electrons. Decay electron tracks have low impact parameter and are emitted isotropically. The background structure results from selecting beam electrons that intersect both ePCs.

distribution result from the requirement that tracks intersect both ePCs. In fact the background should be higher outside of these regions, particularly in the region of low impact parameter and high $\cos(\theta)$, but again such tracks are not visible.

Dividing these two histograms results in the signal to background ratio distribution shown in figure 12.9. By selecting upstream-going electron tracks with a tight impact parameter cut it is possible to achieve signal to background ratios over 100, albeit with low detection efficiency. Restrictive electron tracking cuts are not normally used because they can distort the observed lifetime distribution, but in this case since the clean track is not the one being added to the histograms it should be safe to apply these cuts to them. Several cuts with




Figure 12.9: Signal to background ratio for electron tracks as a function of impact parameter and angle. Background is minimized by selecting upstream-going electrons with low impact parameters.

different signal to noise ratio targets are defined, and no problems with the strict cuts are observed. We avoid imposing any tracking cuts on the background electron to mimic the standard eSC-only electron definition.

One concern is that the clean track selection could influence the other background electron. This case is unusual in that we are comparing eSC electron detections with and without associated ePC tracks, whereas typically one would choose a consistent electron definition for both. These comparisons finally uncovered the bug in the electron tracking software discussed in appendix A, which produced unphysical correlations when comparing electrons with and without associated tracks. With these tracking issues now largely resolved the two electrons should be approximately independent, since long-range correlations have already been shown to have little effect.

In addition to the geometrical cuts the clean tracks are also required to occur within a short time window at the beginning of the event, with the background fits beginning after

this window. This further increases the signal to background ratio for the track, while preventing problems caused by coincidences between the two electrons. An 800 ns window allows standard fits starting at 1 μ s, and improves the signal to background ratio by an additional order of magnitude at the cost of a 60% reduction in statistics. Figure 12.10 shows an example of several background histograms produced via this method, using increasingly strict geometric cuts.




Figure 12.10: Example background distributions produced by the clean track analysis for R2015, with an 800 ns window. Stricter geometric cuts can reduce the signal fraction N/B , at the cost of reducing the size of the background B . The clock background is shown in black for comparison.

The clean track analysis can reduce the fraction of muon decay events in the background histogram below the percent level, which is impressive but still an obstacle to measuring a 10^{-4} slope. This residual signal must be accounted for when fitting the slope, typically by adding a fit parameter for the size of the signal contribution while keeping the decay rate fixed. Moving from the two-parameter linear fit to this three-parameter fit limits the power of this method, resulting in dr estimates with similar or somewhat lower precision compared

to the clock data. The signal parameter could instead be fixed by estimating the signal to background ratio for the clean tracks themselves, but this has been avoided out of concern it might bias the fits.




Figure 12.11: Background fits for the early R2014 datasets with the misaligned collimator (left) and the later datasets with the corrected collimator (right). Despite the inverted kick polarity and much different background rates the background slope appears similar, consistent with the earlier rate shift measurements.

This discussion has focused primarily on the R2015 data, where the analysis can achieve significantly higher statistics than the clock data. For R2014 the situation is rather different, with a 2.5 times higher clock rate making it difficult for this analysis to compete unless loose cuts and later fit start times are used. However, since the clock was only active for a few datasets mainly towards the end of R2014, the clean track analysis is useful as a substitute for the clock data when attempting to study the earlier datasets. These are particularly interesting, since the issues with the upstream collimator during DSM1-3 produce a much higher electron background and reverse the direction of the kicker step. Despite these differences, within the kicker window itself there appears to be a somewhat similar downward slope regardless of beam configuration as shown in figure 12.11.

Finally, because the background time distributions produced by this analysis are independent from the clock data it is possible to get higher statistics by simply adding both histograms together. This still has the downside of including a small muon lifetime component, but can produce more precise slope estimates which is useful for R2015 where the slope is relatively small. Of course this only works for datasets including the clock and cannot be applied consistently to R2014.

12.2.4 *Simultaneous Fit*

The most precise approach to constraining the background slope is to include it directly in the muon lifetime fit function. The direct studies of the background have independently corroborated the presence of a slope, and the linear model appears to be a reasonable approximation of the data. Thus, while the fit method is more model-dependent than the direct approach this should not be a problem.

Simply adding a linear term to the standard fits immediately runs into major problems due to the strong correlation between the decay rate and the slope. To constrain both parameters we simultaneously fit both the 1e and 2e histograms, constraining both fits to use the same R and dr parameters. In this way the rate estimate is dominated by the 1e histogram and the slope estimate is dominated by the 2e histogram, while automatically compensating for changes in the other fit parameters. Simultaneous fits are possible in ROOT by defining a custom global χ^2 function for both fits, although in practice the same effect may be achieved more easily by simply concatenating the data into a single large histogram. An example fit is shown in figure 12.12.

The simultaneous fit approach has several advantages. First, by leveraging the full 2e background signal as well as the differences between the 1e and 2e histograms, it achieves substantially more precise dr estimates than the direct fit methods. It can also be applied to any dataset regardless of the clock, making this method useful for checking the consistency of the early R2014 datasets similar to the clean track analysis. Figure 12.13 shows the fitted slopes and also the 1e signal to background ratio for each dataset. Evidently the reason why




Figure 12.12: Simultaneous fit to 1e and 2e histograms for R2015, including both background slope and blinded capture rate parameters. The resulting slope parameter is similar to those obtained via direct background measurements. The reduced χ^2 value is 1.03(4), compared to 0.95(6) for the standard 1e fit and 1.10(6) for a fit to all electrons.

DSM5 has the largest electron multiplicity effect is actually because there is an unusually low constant background term, resulting in a much larger relative slope. The other four R2014 datasets surprisingly have fairly consistent slope parameters despite the large changes in background rate due to the collimator changes.

Another major advantage of the simultaneous fit is that, because the slope is directly included in the fit, the resulting muon lifetime estimate has already been implicitly corrected to account for its effects. In contrast the other slope estimates would require an additional step of correcting the standard lifetime fits based on the measured slope value. The rate estimates produced by the simultaneous fit are shown in figure 12.14. Although these histograms are conventionally drawn relative to the all-electron fit, the vertically offset is arbitrary since these are all blinded values. In fact the simultaneous fit should give the best estimate of the true disappearance rate, and the other two fits should be understood as being shifted to




Figure 12.13: Slope and 1e background parameters found by the simultaneous fit, for each dataset.

lower decay rates due to the influence of the background. Interestingly the uncertainty in the fitted rate only increases slightly when using this new fit, from 8.5 to 8.8 s^{-1} for R2014 and from 6.5 to 6.7 s^{-1} for R2015. This corresponds to correction uncertainties of 2.1 and 1.5 s^{-1} for R2014 and R2015, respectively. The simultaneous fit finds a reduced χ^2 value of $1.03(4)$, which is not much different from the values of $0.95(6)$ for the standard 1e fit or $1.10(6)$ for a fit to all electrons.

12.2.5 Slope Fit Comparison

In summary, three separate methods have been devised to measure the background slope. The first method is simply to directly look at the muon clock data, which provides an unbiased look at the background but suffers from relatively low statistics and is not available for all datasets. A second method uses electron tracking cuts to isolate the background from normal events, which can be used to study datasets with no clock but includes some




Figure 12.14: Fitted decay rates found by the simultaneous fit for each dataset, with the 1e and all-electron fits shown for comparison. Rates are shown as shifts relative to the fit to all events.

residual muon decay signal which must be accounted for. These two methods also produce independent background histograms, which may be added together to reach higher statistics.

The final method defines a specialized fit function including the slope, and uses it to simultaneously fit both the 1e and 2e lifetime histograms. Not only does this method produce the most precise estimate of the background slope, but it also automatically corrects for its effect on the muon lifetime. To derive a correction using the more direct slope measurements from the first two methods would involve an additional step of estimating the rate shift as a function of slope, further increasing the uncertainty. These two methods thus serve primarily as a test of the reliability of the background slope model, while the simultaneous fit method is best for obtaining a quantitative correction.

A comparison of the estimated slope values is given in figure 12.15. The clean track analysis and the simultaneous fit produce very consistent slope estimates, although it is




Figure 12.15: Comparison of the various dr measurements.

important to keep in mind that the clean track analysis can vary quite a bit depending on the specific quality cuts and fit settings being used. In R2014 the clock analysis finds a higher slope, which is expected since it is dominated by DSM5 as discussed previously. The R2015 clock slope also appears somewhat low, although it would increase with a later fit start time as shown in figure 12.7 above. In general the slope appears to be nearly twice as large in R2014 as in R2015, presumably due to differences in the beam tune changing the makeup of the background.

12.3 Potential Causes

Several models have been proposed that might explain the origin of this time-dependent background. These models are all somewhat speculative, since with such imprecise measurements there is little information to work with. Despite this they are still quite interesting, as each model has different implications regarding the behavior of the background as seen by other detectors or with other cuts applied. If the source of the background slope can be identified, it may be substantially easier to constrain by measuring in a more direct manner. This section will discuss each of the leading models as well as additional measurements designed to test their predictions.

12.3.1 *Beam Fluctuations*

One possible source of a time-dependent background would be an oscillating in the incoming beam rate, which would result in muons preferentially entering at the peak of the oscillations and could produce a slope across the event window if it occurred with the correct frequency. This cannot explain all of the observed effects, as it would only couple to real muon entrances while any oscillations should be averaged out in the μ CL histograms. It may be one of multiple effects, however, so a Fourier analysis was used to check the beam stability as described in appendix H. A small 5 kHz oscillation was identified but it would produce a δr value of only $2 \times 10^{-7} \mu s^{-1}$, far below the level required to produce a measurable shift in the fitted muon decay rate.

12.3.2 *Detection Efficiency Changes*

Concerns have also been raised about the stability of the electron detection efficiency. The most obvious source of such an effect would be afterpulsing, which has already been ruled out, but there could also be a global change in efficiency possibly caused by problems with the electronics or the DAQ system. This would need to somehow be correlated with the kicker, since the slope is seen in the clock events with no real particles in the detector. It should also affect the decay signal as well as the background, meaning an increase in detection efficiency at late times would probably be needed to increase the lifetime rather than the downward slope that we observe. It is difficult to construct a model for the detector efficiency that explains all the observed behaviors, so this is not considered a likely cause of the effect.

12.3.3 *Kicker Instability*

Another way the kicker could produce a time-dependent background would be if the voltage changes. As shown in section 10.3.1, when the kicker diverts the beam it surprisingly produces an increased beam electron background in most datasets. This additional background is thought to result from beam muons that are deflected just far enough to graze the edge

of the beam beam collimator positioned after the kicker. The collimator is made of copper blocks with significant thickness along the beam axis, and as shown in figure 12.16 the background would come from muons that impact the inner face of the collimator rather than the front. This effect would be very sensitive to small changes in beam deflection angle, so if the kicker HV drifts over the length of the event window it might produce a background slope effect.

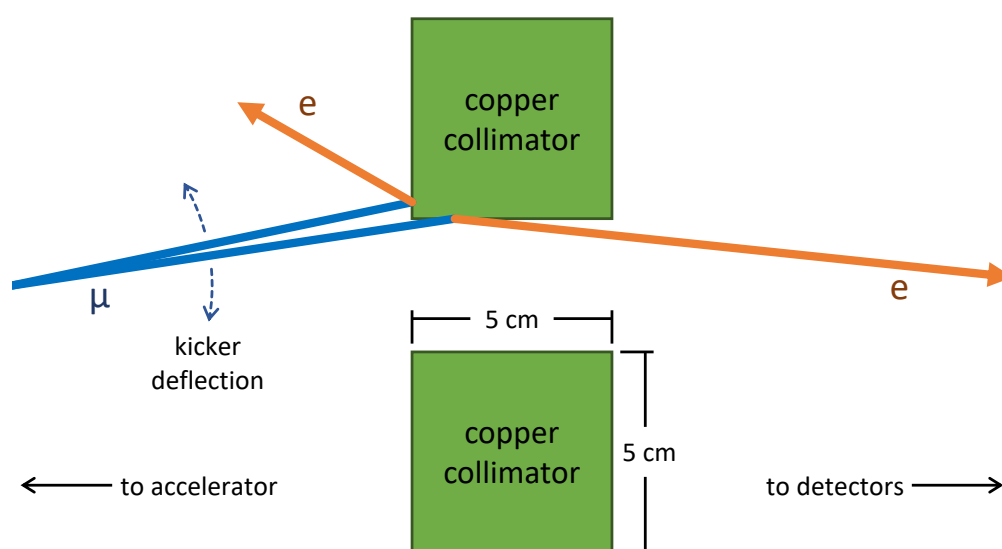


Figure 12.16: Diagram of the hypothesized mechanism whereby partially deflected muons impact the inner face of the beam collimator blocks, emitting secondary electrons that increase the beam background in the kicked configuration.

A major issue with this model is that in R2014 DSM1 and DSM2 the first beam collimator was misaligned, while in DSM3 it was removed entirely. There was a second upstream collimator at another beam focus point in the middle of the beam pipe, which was able to partially compensate for the issues with the first collimator and still achieve acceptable extinction factors for these datasets. However, compared to the later datasets with a corrected collimator these first three datasets have much higher beam electron backgrounds and kicker

steps that go in the opposite direction. If the background slope is truly the result of small changes in the alignment of the beam and the collimators then it is surprising to see such similar slopes for these different datasets. The model cannot be ruled out though, as the process of tuning the collimators to optimize the extinction without blocking the beam may naturally select such a configuration where the beam is just barely blocked.




Figure 12.17: Muons (left) and electrons (right) detected by the μ SC in R2015 μ CL events. A fit start time of 10 μ s is used to avoid the decay electrons produced by pileup muons, since pileup protection is not applied for this analysis. No slope is observed for electrons, while the muon fit finds a slope similar to that seen in the eSCs but with over 100% uncertainty.

If the beam deflection is changing then an obvious test is to study the muon entrance detectors. A large slope in the μ SC background would be strong evidence in favor of the model, but as shown in figure 12.17 this is not the case. Once again this does not rule out the beam deflection model, as the cylindrical electron detector arrangement means that electrons traveling along the beam axis will not be detected. The mechanism discussed above may

produce primarily off-axis electrons that are more likely to intersect the eSCs, while having a negligible effect on the on-axis beam rate seen by the entrance detectors. In principle muons impacting the horizontal face of the collimators should also produce a vertical asymmetry in the background, but the low statistics make testing this difficult. Splitting the clock background into upper and lower eSC segments results in dr estimates of $(3 \pm 5) \times 10^{-5} \mu s^{-1}$ and $(6 \pm 5) \times 10^{-5} \mu s^{-1}$ respectively, which is somewhat suggestive but not a statistically significant difference.

In summary, the kicker instability model is somewhat disfavored due to the similarity of the slope estimates for datasets with very different beam conditions, but this may be a result of the beam tuning process. Attempts to monitor the beam deflection using the entrance detectors or to observe an angular asymmetry from the effect are inconclusive, so it cannot be ruled out. The dedicated beam background testing performed during R2016 is also not particularly helpful for constraining this model, as discussed in the next section.

12.3.4 Long-Lived Daughter Nuclei

Finally, wall stop events may produce a small tail of very slow decays. At first this would seem to be impossible, since wall stops typically capture very quickly and certainly should not decay more slowly than the free muon lifetime. However, the initial muon capture event may also leave the nucleus in a metastable excited state, some of which have lifetimes on the scale of the MuSun event window. Gamma rays emitted by the decay of these isotopes could then scatter and produce detectable electrons, or the decay could occur via a process known as internal conversion that directly emits an electron. The internal conversion probability depends on the specific isomer, but they can be a large fraction of the total decays for low-energy transitions in high-Z atoms[43].

Evidence for such an effect was observed during the R2016 beam studies, where a 1 mm lead plate after the normal entrance detectors as described in section 6.4.3. A relatively large background slope was observed for this dataset, shown in figure 12.18, but comparison with the μ CL data indicates this slope is associated with the muon stops in the lead rather




Figure 12.18: Decay time distribution for R2016 test data with the muon beam stopped in a lead sheet, with the expected lead lifetime and background removed. A double-exponential fit is performed including a component for the muon lifetime in carbon (N_c and R_c) to account for stops in the scintillator itself, and an additional slow exponential decay.

than the beam background. The R2016 signal appears to be exponential, and a fit finds a decay rate of $(6.4 \pm 1.7) \times 10^{-5} \text{ ns}^{-1}$. This is equivalent to a half-life of $10.8 \pm 2.9 \mu\text{s}$, and may correspond to an excited isomer of thallium (^{203m}Tl) with a half-life of $7.7 \pm 0.5 \mu\text{s}$ [41]. This isotope should be produced by roughly 4% of muon captures in lead, although the probability of producing the long-lived isomer is less clear.

In real data the only source of stops in lead is the collimator included in the entrance detector stack. These stops can be minimized by applying μSCa pileup protection, and appear to account for roughly 10% of the total effect at most. However, the observation of this signal from lead raises the concern that stops in other materials could have similar effects. In particular, approximately 2/3 of muon entrances have no associated TPC track, in large part due to stops in the front flange of the pressure vessel or the front part of the

TPC frame. Even a small fraction of slow decay signals resulting from these stops could therefore be responsible for a background slope at the 10^{-4} level, with nickel (in the stainless steel) and tungsten as the most likely contributors.

Tungsten is particularly promising because it can produce two isotopes of tantalum ($^{180}_{73}Ta$ and $^{181}_{73}Ta$) that both have multiple isomers with half-lives ranging from 17 to 31 μs . Fitting the background of the R2016 calibration data for wall stops in tungsten finds a dr value of $(3.4 \pm 0.7) \times 10^{-3} \mu s^{-1}$. When accounting for an estimated 12% muon stop fraction in tungsten components and a delay of approximately one half-life before the next clean muon stop, this slope is reduced to $2 \times 10^{-4} \mu s^{-1}$. A final factor of two reduction to account for reduced detection efficiency in the real experimental setup would put this estimate in good agreement with the R2015 slope measurements.

Nickel could produce a cobalt isomer ($^{58}_{27}Co$) with a shorter 10.4 μs half-life, which would be consistent with the time dependence in the R2014 μCL fits and may explain the larger overall slope in that run. Differences in the beam tune between R2014 and R2015 could change the muon stop distribution, potentially increasing the stops in the flange during R2014. Thus, while not definitive it would seem that these long-lived isomers may be able to produce the observed background slopes. See appendix I for a more detailed description of this effect, including an inventory of the various materials included in the MuSun experiment and a comparison to known long-lived isomers.

12.4 Simulation

Finally, we return to the question of the reliability of the linear slope approximation. The two leading models are both non-linear, with the long-lived isomers producing exponential decays while the kicker instability would likely also be well approximated by an exponential. We will use a simulation of the beam background to test the accuracy of the slope measurements and their effectiveness when applied to a curved background.

12.4.1 Simulation Procedure

The simulation attempts to match the R2015 data, including known effects such as the 10 ns coincidence veto and the 100 ns artificial deadtime following an electron detection. The constant background term is set separately for each eSC segment, again based on the R2015 data. The background shape is controlled by two variables, the standard dr and an additional quadratic curvature parameter c. This quadratic component represents the next term in the Taylor series expansion of the true background shape, but is parameterized as

$$B(t) \equiv B_0 \left[1 + dr(t - t_0) \left(c \frac{t}{t_0} - 1 \right) \right], \quad (12.3)$$

with B_0 and t_0 being the nominal background and the length of the event window respectively.




Figure 12.19: Effect of the curvature term on the background shape.

Figure 12.19 shows an example of how the curvature parameter affects the background shape. The background at times 0 and t_0 are fixed by B_0 and dr, with $c = 0$ giving a linear slope while values of ± 1 produce parabolas with maxima at the beginning or end of the event. This allows the effect of the shape to be tested without altering the overall size of the effect, which is still entirely controlled by dr.

For a given set of input parameters the simulation models the theoretical time distribution of this background and a pure exponential decay signal, including the effects of the 1e or 2e cuts and clock histograms. For each 40 ns time bin it then generates a random bin content to reproduce the statistical fluctuations in the data. The result is a set of 1e, 2e, and μ CL histograms that can be analyzed using the various fit methods described previously, omitting the clean track analysis as that would require simulating the ePC information. This process is repeated many times while scanning over the slope or curvature parameters to produce scatter plots of the fitted parameters as a function of the input parameter.

12.4.2 Linear Slope

Testing was first carried out using a linear background, to check that the system was working properly and to study the accuracy of the different methods. As one would expect, all of the methods perform well in this case and the simulated software deadtimes have little effect. For small slopes the fitted rate shift is directly proportional to dr , with a dr of $10^{-4} \mu s^{-1}$ producing a $6.2 s^{-1}$ shift. The estimated slope to match the observed shift between the individual 1e and 2e fits is also consistent with the dr measurements from the data.

12.4.3 Curvature

Testing of the curvature was done while keeping dr fixed at a value of $10^{-4} \mu s^{-1}$, which should be fairly representative of the data. Once again we see approximately linear dependence of the rate shift on the input c parameter. The 1e rate shift changes from $6.2 s^{-1}$ at $c = 0$ to $4.5 s^{-1}$ at $c=1$, a change of 28%. The 2e fit is less sensitive to the curvature, only changing by 7.3% from $323 s^{-1}$ to $299 s^{-1}$. This difference is a consequence of the different times at which these two histograms transition from being signal-dominated to background-dominated, resulting in a different effective dr in the curved case.

The clock fit over the entire event window consistently returns the correct input dr of $10^{-4} \mu s^{-1}$, but does nothing to correct the rate shift due to curvature. A later start time for the clock fit can reproduce the smaller effective dr for the lifetime fits.




Figure 12.20: The decay rate shift (left) and estimated slope values (right) produced by the simultaneous fit for various background curvatures. The estimated dr value differs from the input of $10^{-4} \mu s^{-1}$, but in just the right way to cancel any effect on the fitted rate.

Finally, the simultaneous fit performs extremely well in this test, with average rate shifts much lower than $1 s^{-1}$ for any combination of input parameters tested. The estimated dr varies by $\pm 25\%$, which again matches the expectation that the curvature reduces the effective slope. This is able to almost perfectly correct for the any effect on the fitted decay rate, resulting in a shift of only $\pm 0.15 s^{-1}$ when testing with a dr value of $10^{-4} \mu s^{-1}$ as shown in figure 12.20. The insensitivity of the fit to the background curvature was not intentional and it is slightly unclear whether this would continue to hold true for cubic or higher-order terms, but given the flatness of the observed background it seems safe to neglect these additional terms. Thus, not only does the simultaneous fit produce the most precise background slope estimates and automatically correct the output decay rate, but it also appears to be extremely robust against any reasonable variations in the background shape.

12.5 Conclusion

Comparison of different electron multiplicity cuts has revealed the presence of a previously unknown contribution to the background which is slowly varying in time. The effect has been modeled using a linear slope term, which was measured with several methods all of which find fractional slopes on the order of $10^{-4} \mu s^{-1}$. Direct measurements using the muon clock as well as attempts to isolate the background from normal production data both support the presence of a small downward slope, and were unable to identify any large non-linearity. Another method uses a simultaneous fit to both 1e and 2e histograms to produce a more accurate slope estimate and decay rate correction.

A few hypotheses regarding the source of the background slope were presented, with two leading candidates. Instability in the kicker could cause the beam deflection angle to slowly drift, with differences in the alignment between the beam and the upstream collimators potentially producing a change in the observed beam electron background. There is also the surprising possibility of slow decays of excited nuclear isomers produced as a side effect of wall stops, with evidence for such an effect seen in some of the R2016 test data where the muon beam was deliberately stopped in high-Z targets.

Regardless of the cause of the effect, testing with simulated data indicates the simultaneous fit can reliably extract the correct muon decay rate even when using a slightly curved background. Compared to the standard single-electron exponential fit, the simultaneous fit finds shifts of $9.8 \pm 2.1 s^{-1}$ and $5.7 \pm 1.5 s^{-1}$ for R2014 and R2015 respectively.

Chapter 13

FUSION INTERFERENCE

Chapter 9 has already discussed how tracking errors in the TPC can lead to a distortion in the muon lifetime, and described a new optimized tracking algorithm that attempts to minimize this effect. However, even the best tracker cannot simply be trusted to reconstruct all events accurately. Monte Carlo simulation can be used to test the trackers to some extent, but it is not entirely representative of real data and the P-Dir tracker in particular may be very sensitive to any discrepancies due to its multiple finely tuned cuts. It is therefore necessary to evaluate the amount of tracking interference present in real data and use this to estimate an uncertainty and possibly a correction to the muon decay rate. This chapter focuses on fusion interference, but the analysis of electron interference discussed in chapter 14 relies on many of the same ideas.

Fusion interference is a long-standing issue for the experiment. The change in reconstructed muon stop position due to fusion products in the TPC is known as migration, and any net migration into or out of the fiducial volume cut can significantly bias the fitted muon disappearance rate. Previous analyses have mainly taken the approach of identifying p+t fusion events in the TPC and then determining a lifetime correction based on an estimate of the number of migrations. This chapter will cover some relevant aspects of those analyses, before introducing a new method based on a direct measurement of the lifetime distortion.

13.1 Migration Modeling

As mentioned above we use migration to describe the mis-reconstruction of the stopping position. A first simplified migration model was developed in the PhD thesis of M. Murray[54], expanding on the fusion formalism introduced in section 9.1 by estimating the probability

of any given fusion causing a migration. The net migration into the fiducial volume is the sum of migrations across each of the six fiducial boundaries, and as a first approximation transverse migrations may be ignored and only migrations across the upstream and downstream Z boundaries need to be considered. Using a slightly simplified model of the proton trajectory in the TPC, it is possible to calculate the probability of a migration as a function of the distance from a pad boundary. This probability may then be applied to the muon stop distribution to estimate the expected number of fusion migrations, and thus derive a lifetime correction.

13.1.1 Migration Probability




Figure 13.1: Diagram of the area reachable by a proton emitted by a muon stop at depth d within a pad, with the blue region indicating upstream migrations for the road tracker.

The two best trackers are affected by fusions in complicated and subtle ways, with the upstream tracker potentially selecting a different set of upstream pads to fit while the P-Dir tracker's cuts are sensitive to small differences in energy deposition. The migration probability was thus estimated for the much simpler road tracker, with the hope that the consistent single-direction migration behavior of the road tracker would lend itself better to

modeling and correction. Protons emitted by p+t fusions are emitted isotropically, resulting in a sphere of potential proton stop positions with radius equal to the proton range $r_p \approx 13.2$ mm as shown in figure 13.1. For a muon stop at depth d into the stop pad, the probability of the proton returning to the previous pad becomes

$$P_{mig}(d) = \frac{\Omega}{4\pi} = \frac{1}{2} \left(1 - \frac{d}{r_p} \right). \quad (13.1)$$

This estimate is based on a simple requirement that the proton reaches the target pad, but migration using the road tracker requires that the combined energy of the proton and the muon E_1 must exceed the road threshold of 1300 keV. This results in a slightly shorter effective range which varies as a function of muon stop depth, mainly at large stop depths where the maximum migration range is reduced to 12 mm. For stops very close to the front of a pad additional short-range effects also become important. The triton may also cause migrations similarly to the proton, and because these particles are emitted back-to-back the migration probability rises to nearly 100% near the boundary as one of the two is practically guaranteed to travel upstream to some extent. However, many of these events deposit too little energy on the stop pad to be detected and would have been reconstructed upstream even in the absence of the fusion, so they do not count as migration events. Finally, ^3He fusions may remain on the stop pad and deposit enough energy to push such events above threshold, resulting in a differential downstream migration of up to 50% near the boundary. The improved thresholds in the production data once again significantly reduce the size of this effect, but it as well as real ^3He migrations across the boundary remain an issue at extremely short ranges.

Some small but possibly important details have been omitted from these calculations. Fusions occurring after the end of the clustering window must also have the proton travel vertically downward far enough to reconnect to the drifting charge cloud. This will add a second horizontal boundary to figure 13.1, with only a wedge intersecting both the Z and Y boundaries able to cause migration. Even within the clustering window the assumption of spherical symmetry is not entirely accurate. Recombination effects can change the observed

fusion energy depending on angle of the proton track to the vertical, while pad threshold effects may come into play for protons crossing into different X rows.

13.1.2 Stopping Distribution

To actually apply this model the migration probability must be integrated over the stopping distribution to find the total number of migrations across the boundaries. This requires the original stop distribution for fusion events in the absence of migrations, which should be proportional to the non-fusion stop distribution. To minimize the effects of ^3He or electron migrations we may make a tight cut around the muon band in $E_0 - E_1$ space to select only clean muon stops, as shown in figure 13.2.




Figure 13.2: Narrow E_0 vs E_1 cut selecting only the primary stop band. The red line shows the curve predicted by a power law range-energy relation for the muon.

Originally a fit to the observed Z distribution was originally used as a proof of concept[54]. An improved stop distribution estimate may be obtained from the muon range-energy rela-

tionship, which may be approximated with a power law

$$E_0 = E_{max} \left(\frac{d}{l} \right)^\alpha. \quad (13.2)$$

Here d is the stop depth within the pad, l is the pad length, and $E_{max} \approx 1200$ keV is the energy deposition of an idealized straight muon track stopping precisely at the back of the pad. This equation only applies to E_0 because for E_1 and higher one must subtract the remaining energy of the muon when it leaves the pad:

$$E_1 = E_{max} \left[\left(\frac{d}{l} + 1 \right)^\alpha - \left(\frac{d}{l} \right)^\alpha \right] = E_{max} \left[\left(\frac{E_0}{E_{max}} \right)^{1/\alpha} + 1 \right]^\alpha - E_0. \quad (13.3)$$

SRIM¹ calculations predict the exponent should be about $\alpha = 0.54$, although a value of 0.53 appears to match the data slightly better and is shown in figure 13.2 above.

Using the range-energy relation one may estimate the stop depth for events in the primary muon band, resulting in the stopping distribution in figure 13.3. This procedure works well for stops in the middle of the pad, but produces some artifacts near the pad boundaries. At the front of a pad there are some losses due to the pulse-finding threshold as mentioned previously. At the back of a pad percent-level errors in energy calibration can cause either a gap or spike in the Z distribution where the predicted depths do not perfectly align with the pad lengths. To obtain a relatively smooth distribution the energy scales of certain rows have been adjusted slightly to correct for disagreement in the calibrations.

The Z distributions obtained with this method should be reasonably accurate, but they still include threshold effects and electron interference which overlaps with the main E_0 Vs E_1 band. An interesting idea to address these issues would be to instead look for events with a very delayed ^3He fusion, such that the fusion is not included in the stop cluster. The short-ranged helium ion should produce a localized energy pulse on the stop pad, helping to disambiguate the true stop pad. Furthermore, the muon decay must occur after the fusion so this would significantly reduce the amount of electron interference. However, since

¹www.srim.org




Figure 13.3: Approximate stop Z distribution obtained using the muon range-energy relation. Artifacts near the pad boundaries are produced by threshold effects and slight energy calibration differences.

this detailed migration model is not the preferred fusion interference correction method the delayed ${}^3\text{He}$ idea has not been pursued further.

13.1.3 Limitations

Some limitations of the migration model have already been discussed above. It is only really effective for the simpler tracking algorithms, meaning there is a relatively large number of migrations which must be corrected for. It also depends on a precise reconstruction of the true stopping distribution which has its own associated uncertainty. Another major problem is that only Z migration has been addressed, but while migration in the X and Y directions is generally smaller than in Z it is still a significant contribution to the net migration.

In general this method is somewhat backwards, with a simple tracker followed by a complex and model-dependent correction. It would be preferable to use a simpler data-driven method to determine the final correction, as this would produce a more robust error

estimate. A more sophisticated tracker such as the P-Dir could then be used while remaining fairly confident that the final correction step should account for any discrepancies between the model and reality. Thus, while the migration model is helpful for understanding the behavior it has since been superseded by alternate fusion interference correction methods.

13.2 Fusion Fraction

Rachel Ryan developed a more robust method of quantifying the net fusion migrations in her thesis[62]. This method involves tagging fusion products to determine the fraction of events with a fusion. For a given set of analysis cuts, the observed fusion fraction $\tilde{\epsilon}$ is equal to

$$\tilde{\epsilon} \equiv \frac{\tilde{F}}{\tilde{N}} = \frac{F + M}{N + M}. \quad (13.4)$$

This uses the notation developed in section 9.1, with N referring to the total number of muon stops, F referring to the number of fusions, and M denoting migrations. A tilde indicates an observable quantity, while variables without a tilde represent the underlying true distribution.

The fusions are only a few percent of the total number of events, and typically only a small fraction of fusions actually migrate across the fiducial boundary. M is therefore relatively insignificant compared to N, so the fusion fraction simplifies to

$$\tilde{\epsilon} \approx \frac{F + M}{N}. \quad (13.5)$$

If the true fusion fraction in the absence of migration, ϵ , is known, then the observed fusion fraction may be used to directly calculate the fraction of fusion migrations and thus the overall decay rate correction:

$$\Delta\lambda = \kappa \frac{M}{F} = \frac{\kappa}{\epsilon} (\epsilon - \tilde{\epsilon}). \quad (13.6)$$

The true fusion fraction as well as the proportionality constant κ must be determined empirically.

13.2.1 Fusion Tagging

This approach focuses on p+t fusions, as they deposit approximately 4 MeV of additional energy in the TPC making them clearly distinct from non-fusions or ^3He fusion events. Within the fiducial volume detecting this additional energy is easily achieved by selecting events with high cluster E_9 (The total energy on the stop pad and the 8 adjacent pads, see section 5.5.2). A threshold value of 3850 keV is normally used, but as shown in figure 13.4 there are only a few neutron scattering events occupying the range between 3000 and 4000 keV. R. Ryan's original analysis was somewhat more sophisticated and attempted to subtract out the neutron contribution, which was a relatively small correction[62].




Figure 13.4: Monte Carlo E_9 distribution, tagged by event type. A cut of 3850 keV is used to identify p+t fusions while minimizing the contribution from the neutron scattering tail.

Outside the fiducial volume the fusion tag begins to lose efficiency, as for stop positions near the edges of the TPC the proton may escape the active volume causing some energy to be lost. Appendix J discusses some modifications to the fusion tag designed to correct some of these losses. This issue is not a major concern for the main analysis, and primarily stop position scans where the edge pads will incorrectly show a reduced fusion fraction.

13.2.2 Lifetime Effect

By artificially selecting or rejecting events based on the fusion tag, one may easily obtain any desired fusion fraction. Performing lifetime fits with several different fusion fractions produces a plot of the rate shift with respect to ϵ , as shown in figure 13.5. The proportionality constant κ is then simply obtained via a linear fit to these rates. A stop position scan may also be used to ensure the tag is representative of real migrations, and while this method is much less controlled it produces similar results.




Figure 13.5: Decay rate shift as a function of observed fusion fraction $\tilde{\epsilon}$. The true fusion is roughly 2.6%, but is difficult to determine precisely.

13.2.3 Stop Distribution Shaping

A 1% error in estimating the number of fusion events would correspond to a $5 s^{-1}$ rate shift based on the slope from figure 13.5. This level of precision is achievable for the observed fusion fraction $\tilde{\epsilon}$, but recall that in equation 13.6 the rate shift was proportional to the difference between the observed and true fusion fractions. The true fusion fraction ϵ must

therefore be determined to a fractional uncertainty of roughly 0.5%, and this step is the main difficulty for the fusion fraction approach.

One way to extract the true fusion fraction would be to return to the migration modeling approach and attempt to correct for the observed migrations. On its own this would seem to provide no benefit, as if the migration model reached the desired accuracy there would have been no reason to introduce an alternate fusion interference correction method. However, something that the migration model makes quite clear is that the number of migrations is strongly dependent on the original number of muon stops near the boundary. Even with a large fractional error in the migration estimate, if the number of events migrating across the boundary is low then the absolute error in the fusion fraction will be small. If the stopping distribution could be altered such that all muons are in the center of the TPC and the regions near the boundary are depleted, it would become impossible for fusion products to cross the boundary at all and thus no fusion migrations could occur!




Figure 13.6: First row X vs stop X (left) and first row Y vs stop Y (right). The red line indicates the expected 45° line for perfectly straight tracks, with the actual distribution indicating a slight spread in track angles and generally downward slope.

It is not possible to achieve such a shaping effect by directly using the reconstructed stop position, since that is affected by interference and would be identical to the existing fiducial cut anyway. The trick is to instead look at the upstream Z rows, making use of the fact that fusion products have a range of less than one pad length so any pad two or more rows upstream from the stop should be clean. In particular, because the fiducial volume starts on the third row that means the first row should be clean for all fiducial stops. The muon tracks are approximately aligned with the Z axis, so the first row X and Y (FrX and FrY) are well correlated with the stop positions as shown in figure 13.6. Cuts on the first row position thus act like a virtual collimator, removing large-angle scattering from the entrance window and focusing the beam into the center of the TPC.




Figure 13.7: First row energy vs stop Z , with an approximation of the range-energy relation shown by the red curve.

The stop Z position may also be predicted from only first row information by using the

first row energy (FrE) to estimate the muon range. This relationship is shown in figure 13.7, and while it is less precise than X and Y it is able to produce a substantially narrower stop Z distribution and significantly deplete events near either Z fiducial boundary. This approach is rather similar to the upstream tracker, which could also be used to obtain a more precise shaped Z distribution. The advantage of the FrE cut is in its simplicity, as well as the fact that it will work even for events where the upstream tracker fit fails for some reason.

Strict cuts on FrX and FrY can nearly eliminate X and Y fusion migrations, and while Z migrations cannot be completely prevented they can be suppressed to a level low enough to be accurately corrected for by the migration model. The shaped distribution may then be used to extract a much more precise estimate of the true fusion fraction, which can in turn be used to correct for fusion interference in the general case. The tight cuts required for this shaping reduce the statistics by roughly 85% and thus cannot be applied to the final analysis, but there has also been some interest in adopting a looser set of first row cuts as part of the standard analysis.

Unfortunately, while the stop distribution shaping idea is reasonable in theory and also works perfectly well on the Monte Carlo data, it has always been difficult to obtain consistent results using real data. As one applies an increasingly strict series of shaping cuts the population near the boundary should be depleted, causing the migrations to fall to zero while the observed fusion fraction converges to the true value. The observed behavior is quite different, particularly for the downstream Z boundary where the number of migrations from row 7 into row 6 appears much higher than one would expect or is even physically possible. This is shown in figure 13.8, where if tight FrE cuts are applied to depopulate row 7 the estimated migrations into the TPC can be more than double the total number of fusions in row 7 to begin with. In fact these mysterious excess fusions at the back of the TPC persist even if the fiducial volume cut is modified to completely remove the downstream cut, including stops in the final Z row and even those exiting the TPC. Also, while the analysis was performed with the road tracker similar behavior is also observed with the P-Dir tracker, so it is clear that these are not normal migrations due to simple tracker errors.




Figure 13.8: Fusion fraction in the fiducial volume ($3 \leq Z \leq 6$) as a function of last row ($Z=7$) population. The Monte Carlo has a higher fusion fraction than the data, and has been scaled by 91% to better align the values for the least restrictive cuts on the right side of the plot. The number of migrations from $Z=7$ to $Z=6$, denoted M , is calculated from the slopes of the lines and is displayed as a percentage of the total estimated fusion population in the last row. The simulation shows a consistent 25% migration fraction as expected for the P-Dir tracker, but in data this fraction is 50% for loose FrE cuts and rises to 200% for very strict cuts. Such high numbers are unphysical, and are an indication that FrE couples directly to fusion events. Vertical grid lines correspond to a rate shift of approximately 3.7 s^{-1} .

The cause of this inconsistent behavior has recently been traced to the high voltage supply for the TPC Frisch grid. Recall that this grid of wires sits just above the anode pads, resulting in all anode pads being capacitively coupled to the grid. When a charge cloud arrives the grid voltage can sag, causing global cross-talk signal to be seen in all TPC pads. This allows the high energy pulses from p+t fusions to affect the first row energy regardless of stop position, and thus invalidates the core assumption of a clean first row signal. The cross-talk effect is now fairly well understood and a model developed by Duncan Prindle is able to reproduce most of the observed effects in Monte Carlo data, but additional development would be required to accurately correct for the problem in real data. First row X and Y cuts remain a useful tool, but this unexpected sensitivity of the FrE cut to fusions makes it a poor choice for this analysis.

13.3 Early Time Deviation




Figure 13.9: Decay time distribution for p+t fusion events, with a muon lifetime exponential fit for reference. After the fit start time at $t_{fit} = 1 \mu\text{s}$ the two curves are somewhat similar, but they completely diverge at early times.

To address the inconsistencies with the fusion fraction determination, a new independent method has been developed to quantify the amount of fusion interference. This method aims to minimize reliance on the TPC data, which may be subtly influenced by the presence of fusions, and instead extracts a fusion interference estimate directly from the final muon lifetime histograms. The key insight is that, while fusion interference has a rather subtle effect using the standard $1 \mu\text{s}$ fit start time, at early times the distortion of the lifetime is much stronger due to the non-exponential fusion time distribution shown in figure 13.9.




Figure 13.10: Decay time distribution for p+t fusion events, split up by migration direction as determined by the P-Dir tracker. The time distribution for upstream-going fusions is very slightly delayed relative to the others due to this cut's tendency to also select events with both a p+t and a ${}^3\text{He}$ fusion.

As discussed in section 9.1, this decay time distribution comes from integrating over the fusion time distribution multiplied by the migration probability function η . The fusion rate is determined by the muon kinetics, while η is a property of the clustering algorithm and is mostly determined by the fusion type and choice of tracker. Thus, to a good approximation the decay times for p+t fusion events always follow this characteristic functional form inde-

pendent of stop position or whether the fusion caused migration. Figure 13.10 verifies this by comparing non-migrating fusions with upstream and downstream migrations, as determined by the P-Dir tracker.




Figure 13.11: Data / fit ratios for all fiducial volume events, and with excess or missing fusions. The shaded regions indicate the integration range for quantifying the deviation, selected to avoid the entrance time and the kicker step at time t_k .

To quantify the extent to which the lifetime histograms are distorted by this fusion time distribution, we start by performing the normal exponential lifetime fit starting at $1 \mu\text{s}$. This exponential curve is then extrapolated backwards to $t=0$, and the ratio of the data to this extrapolated fit is calculated. Ideally this ratio should always be 1, but if fusions are lost or gained by the TPC cuts then the ratio will curve up or down as shown in figure 13.11. We thus define the early time deviation D as the amount that the ratio differs from 1, using a time window from 40 and 720 ns to select the region of flat background between the muon entrance and the kicker step near 800 ns.

To simplify the error estimate D is obtained by performing the integration first and then taking the ratio:

$$D \equiv \frac{\int_{t_1}^{t_2} data(t)}{\int_{t_1}^{t_2} fit(t)} - 1, \quad (13.7)$$

with t_1 and t_2 being the aforementioned 40 and 720 ns respectively. In an idealized model with a universal fusion time distribution and no other backgrounds, this quantity is directly proportional to the distortion of the lifetime at later times as well as to the observed fusion fraction:

$$D(\tilde{\epsilon}) \approx \frac{\tilde{\epsilon} - \epsilon}{1 - \epsilon} D(1). \quad (13.8)$$

Here $D(1)$ is the deviation measured for a pure fusion distribution, and simply acts as a proportionality constant that must be determined empirically.

Using D to quantify the amount of fusion migrations has several advantages over previous methods. As stated above it cannot be influenced by cross-talk or any other unexpected TPC behavior, making it useful as an independent check of the TPC performance. Similarly, the calculation of D is independent of the stop position in X and Z , unlike the fusion tagging which becomes less reliable near the edges of the TPC. Because an ideal pure exponential corresponds to $D=0$ by definition we can also skip the difficult step of estimating the true fusion fraction.

The deviation method is also preferable to performing a full kinetics fit, because as discussed previously the time dependence for migrating fusions is modified by the migration probability $\bar{\eta}(t_f)$. A full fit to the shape at early times would need to account for these effects as well as the precise shape of the kicker background, raising concerns about both the accuracy of the model and the reliability of the fit. The deviation method instead avoids using the non-constant portions of the background while $D(1)$ accounts for $\bar{\eta}(t_f)$, although care must be taken to ensure a representative sample of fusion migrations is used for calibration. $D(1)$ is approximately 0.6, with a D of 10^{-4} producing a 3 s^{-1} shift in the fitted muon disappearance rate.

Because of the insensitivity to X and Z stop position, a loose fiducial volume definition

has been adopted to minimize migration effects. The open downstream Z cut that was experimented with when studying the fusion fraction has now been made standard, although the old fiducial Z range of [2,6] is still used in some cases to eliminate punch-through events if accurate TPC energies are required. In addition the fiducial X cut has been removed, instead relying solely on a first row X cut. These cuts have been shown to produce minimal additional wall stops based on the neutron analysis, and reduce the possible migrations to only those passing through the upstream Z or the Y fiducial boundaries.

13.3.1 Background

The biggest challenge when attempting to accurately determine D is that because it only uses the lifetime histograms those spectra must be free of any other systematics or backgrounds. The presence of the beam background modifies the calculated D value slightly to

$$D = \frac{N_d + B_d}{N_f + B_f} - 1, \quad (13.9)$$

where N and B are the integral of the signal and background components for the data (d) and fit (f), respectively. With the integration window from 40 to 720 ns the ratio of B/N is approximately 0.5% when including all electrons, and lower when using a single-electron cut or full ePC tracks. Expanding in orders of B/N and introducing D_0 to refer to the original background-free deviation, we find

$$D = D_0 \left(1 - \frac{B_f}{N_f} \right) + \frac{B_d - B_f}{N_f} + \mathcal{O} \left(\left(\frac{B}{N} \right)^2 \right). \quad (13.10)$$

The first term is not particularly important, as it just modifies the empirical proportionality constant, although it does mean each electron definition must be calibrated separately. The terms of order $(B/N)^2$ or higher are also safe to ignore, but the second term indicates that any discrepancy between the real and fitted background will contribute to an offset from the correct D value. This is concerning, as the standard fit is only sensitive to the background after the kicker step. The background before the kick normally differs by 15%

or more, causing a 10^{-4} level shift in D and adding an uncertainty of 2-3 s^{-1} to the final correction.

To address this concern, a specialized fit function is introduced which includes the background step. It fits the lifetime histogram starting at 1 μs as usual, but also fits to times before -500 ns to measure the beam rate. The area from -500 ns to 1 μs is excluded from the fit to minimize the difference from the standard fits, and also to ensure that D is measured with an independent set of data points from those used in the fit. Because the kicker step itself is avoided by both the fit and the integration windows, a simple step function approximation is used which often produces a small bump in the ratio histograms but does not affect the results.

At this point the behavior of the background is quite well understood, and the kicker-induced background has been described in section 10.3.1 while the downward slope of the background was discussed in chapter 12. However, one new observation from this analysis that has not been discussed in previous chapters is the presence of a small dip in the background just before the muon entrance, which was traced to issues with the μPC pileup protection cut described in appendix K. This effect may be eliminated with application of a strict μPC pileup protection cut on the single X and Y clusters in addition to coincidences. To ensure there are no problems from this or any other anomalies in the beam background, figure 13.12 shows the D estimates using several electron definitions with various combinations of the single-electron cut, full electron tracks, and this strict pileup protection cut.

The D values obtained using all eSC hits with no multiplicity cut are shifted slightly relative to the lower-background single-electron or full track versions, indicating that the background is not being modeled perfectly. Most likely the fit to the size of the kicker step is slightly off, which is supported by the fact that this shift goes in the other direction for R2014. The consistency between the single-electron and full track measurements is a good sign that they are working reliably, and the eSC-only single-electron version will be treated as the standard for the purposes of this analysis.




Figure 13.12: Deviation values obtained using either eSC-only electrons or full electron tracks, and possibly using strict μ PC pileup protection (PP) or a single-electron cut (1e).

13.3.2 Wall Stops

Capture in wall materials in particular pose a potentially serious problem, as they produce excess early decays which could be misinterpreted as a large loss of fusion events. Fortunately this has been constrained by the neutron analysis and wall stops are not a concern with the standard fiducial Y cut, but for studying high or low Y the fusion fraction method is a better approach. As an extra check that punchthrough muons allowed by the loose X and Z cuts are not a problem, we can apply a minimum S-Energy cut of 1800 keV to eliminate them. This cut reduces statistics by 0.58% and the fit finds $D = -(7.8 \pm 2.8) \times 10^{-5}$, compared to the standard 1e fit result of $-(5.1 \pm 2.8) \times 10^{-5}$.

The S-Energy cut is effective for non-fusion muons, but fusion events can now migrate back into the TPC if an upstream-going proton deposits enough energy to exceed minimum S-Energy threshold. Assuming the cut only has a 3/4 efficiency for punchthrough fusion events yields a predicted change in D of approximately 2.3×10^{-5} , almost entirely accounting for the observed discrepancy of 2.7×10^{-5} . The remaining 4×10^{-6} difference in D values would

produce an error of roughly 0.1 s^{-1} in the final rate correction, so even with this sensitive analysis of the early decay signals the X and Z wall stops clearly have minimal effect. In fact the fitted rate rises by 0.8 s^{-1} when applying the minimum S-Energy cut, which is the wrong direction for both fusion migrations and wall stops and likely results from a similar electron migration effect instead.

13.3.3 Lifetime Calibration

Combining equations 13.6 and 13.8, we find that the rate shift is proportional to D:

$$\Delta\lambda = \frac{\kappa}{\epsilon} \frac{1 - \epsilon}{D(1)} D. \quad (13.11)$$

$D(1)$, κ , and ϵ are all properties of the muon kinetics and the clustering behavior, so for a given choice of fusion type and tracking algorithm the entire prefactor is simply a constant. This proportionality constant is determined empirically by once again studying the fusion events identified in the TPC. A pure fusion distribution would likely be outside the valid range of the linear rate shift approximation, so instead the tagged events are subtracted from the full dataset to produce fusion-veto histograms. These still have large rate shifts and D values that can be used to determine the proportionality constant, and tests subtracting fractions of the tagged fusions show no significant non-linearity.

Because the deviation method is very sensitive to the exact shape of the time distribution, the fusion tagging must be performed more carefully than in the fusion fraction analysis. We select p+t fusion events using the tag discussed above, but also distinguish migrating fusions and identify them as upstream or downstream-going according to the P-Dir cuts. A higher E_9 cut of 7.5 MeV is used to identify double-fusion events, and additional E_0 vs E_1 cuts are defined at lower energies to select ${}^3\text{He}$ and $\mu^3\text{He}$ migrations, as shown in figure 13.13. The ${}^3\text{He}$ fusion cuts are similar to those used by the P-Dir tracker, but here the goal is to select the entire population of helium ions crossing the pad boundary. The non-migrating portions of the ${}^3\text{He}$ bands are not included, both out of concern that they may have a different time dependence and because they are not well separated from the muon band.




Figure 13.13: E_0 vs E_1 cuts to select ${}^3\text{He}$ and $\mu^3\text{He}$ fusion migrations.




Figure 13.14: Decay time distributions for single and double p+t fusions and for ${}^3\text{He}$ fusions.




Figure 13.15: Rate shifts and D values obtained by rejecting various types of fusion events. The point near the origin is the fit to all events, and all other points have vertical error bars representing the uncertainty relative to this fit. The $p+t$ fusion tag is split into migrating (M) and non-migrating ($!M$) fusions, with the former further split into upstream (U) and downstream (D) migrations. The red line was fit to the sum of all migrating fusions including ${}^3\text{He}$ and double-fusions, and the dashed lines show a $\pm 10\%$ slope variation.

Figure 13.15 shows the D values and rate shifts obtained by vetoing each category of tagged event. The trend is measured by fitting the sum of all migrating fusions, including ${}^3\text{He}$ and double-fusions. Non-migrating $p+t$ fusions are not included in the calibration, but since they nearly lie on the line their inclusion would make little difference. The scatter in the different $p+t$ migration tags is attributed to their handling of events with an additional ${}^3\text{He}$ fusion or electron energy deposition, both of which will tend to be identified as upstream proton migrations.

The pure ${}^3\text{He}$ sample would give a slope about half that of the nominal slope, while double $p+t$ fusions have more than twice the slope. This difference comes from a difference in the characteristic decay time distributions for these events, shown in figure 13.14. Double-fusions

are much more delayed, which is to be expected, but the ^3He distribution has a sharper rise than the p+t fusions which is somewhat surprising as it occurs well before the $2\ \mu\text{s}$ effective clustering time window. This may be because when the ^3He energy is split across two pads the resulting low-energy pulses are more likely to be detected when they directly overlap with the muon pulse. This effect will become very important for electrons, which typically do not deposit sufficient energy to exceed the pulse finding threshold on their own.




Figure 13.16: Rate shift estimates corresponding to the D values reported in figure 13.12.

It is not easy to distinguish these different event types using the deviation method, making it possible for a pathological case to occur where one of these event types is dominant leading to a very different slope. However, in general they should only constitute a few percent of the migrations and produce about a 10% change in the slope. To be conservative the uncertainty in the slope is therefore inflated to 20% to account for any reasonable variation in the relative populations of these events. The error on the slope is anyway subdominant compared to the uncertainty in D, at least for the very small D values measured for the full fiducial volume. Figure 13.16 shows the final rate shifts derived from the measured D values from figure 13.12, where the eSC single-electron version is again considered the most reliable.

13.3.4 Road Tracker

So far the discussion has focused on fiducial volume stops determined by the P-Dir tracker, resulting in very low deviations and rate shift estimates. The deliberate subtraction of tagged fusion events is a valid way to perform the calibration, but a more realistic test case for the fusion interference correction is obtained when using the road tracker. Surprisingly, when the same fiducial volume cut is applied using the road tracker the deviation is almost exactly zero. Table 13.1 shows a comparison of the fitted rates for the two trackers and the rates after correcting for the observed deviation. There is an estimated 0.2 s^{-1} correction for the road tracker, while the P-Dir tracker estimates a 1.5 s^{-1} correction which is in good agreement with the difference in the fitted rates.

Tracker	Fitted R [s^{-1}]	Fitted D [10^{-6}]	Corrected R [s^{-1}]
P-Dir	0	-51 (28)	1.55 (91)
road	0.86 (23)	-6 (28)	1.05 (89)

Table 13.1: Original and corrected rate shifts relative to the P-Dir tracker.

This result disagrees with the prediction of the basic migration model. The road tracker has purely upstream fusion migrations, so with the downstream Z cut removed there should be a significant number of outward migrations through the upstream Z cut. The solution comes from the fact that the muon beam diverges, so the upstream migrations of the road tracker also tend to reduce the spread in Y. This effect produces inward migrations through the Y fiducial boundary, balancing the outward migrations in Z.

Using the P-Dir tracker as a reference, we can count the number of events that are reconstructed differently by the two trackers. Table 13.2 shows the number of migrations observed through each fiducial boundary. Evidently there is significant outward migration in Z, but nearly 3/4 of this is balanced by the Y migrations. From the observed D values

it appears that in fact the balance is almost perfect for the road tracker, either due to additional direct Y migrations or because the P-Dir has some migrations itself and is not a completely reliable reference. Some followup tests with Monte Carlo truth might be useful for understanding this behavior better, but in any case this is a good demonstration of the importance of including X and Y effects in the migration model.

Fiducial Boundary	Migrations [10^6]			Effect		
	In	Out	Net	ϵ [%]	D [10^{-6}]	R [s^{-1}]
Upstream Z	0	5.11	-5.11	-1.18	172	5.5
Low Y	3.02	0.035	2.99	0.69	-100	-3.2
High Y	0.913	0.120	0.79	0.18	-27	-0.8
All	3.94	5.26	-1.33	-0.31	45	1.4

Table 13.2: Observed migrations for the road tracker relative to the P-Dir tracker, with estimates of the corresponding deviations and rate shifts. Only Y migrations which occur as a side effect of Z migrations are included, cases in which fusion products directly affect the reconstructed Y position are also possible but should be a much smaller contribution.

The P-Dir tracker becomes beneficial when looking at smaller sub-volumes of the TPC, as it has consistently low migration probability regardless of the muon stop position. This balancing of in and outward migrations with the road tracker is entirely coincidental, and will disappear with a different choice of spatial cuts. Figure 13.17 shows a scan over Z stop position, which is a difficult test for the trackers due to large changes in the muon stop depth distributions for different pad rows. The deviation method agrees well with the estimated fusion fraction, apart from the final pad row where the fusion tag loses efficiency due to punch-through muons. The road tracker appears to have a slight downward slope that is not reflected by the p+t fusion fraction, which may be the result of double fusions or of downstream ^3He migrations as mentioned in chapter 9.




Figure 13.17: Measured values of D (points) and $\tilde{\epsilon}$ as a function of Z . Punch-throughs in $Z=7$ cause the fusion tag to lose efficiency, which also results in some P-Dir tracker errors as the energy cuts become invalid for these events. In the other Z rows the observed fusion fraction correlates well with the deviation for the road tracker, although they disagree slightly for the P-Dir tracker which may be due to ${}^3\text{He}$ migrations.

13.4 Conclusion

Previous attempts to constrain TPC fusion interference have focused on estimating the number of fusion migrations by either modeling the migration probability or by comparing the observed fusion fraction to the expected value. These approaches were quite informative, but struggled to produce precise quantitative results due to the difficulty of properly modeling the migrations in 3D as well as unexpected cross-talk effects in the TPC. The source of the cross-talk has recently been identified and a newly developed model of this effect appears to reproduce the puzzling behavior of the fusion fraction analysis, although it has yet to reach the point of providing a useful correction for the data.

To address the issues with these previous analyses, a new approach was introduced which

relies solely on the observed muon lifetime distribution. Using the fact that fusion events have a characteristic non-exponential time distribution, it is possible to look for the relatively large deviation from the fit which occurs at early times and use this to infer the much more subtle effects within the fit window. This approach is sensitive to other backgrounds and does not differentiate between the different types of fusion which may have somewhat different effects, but the extremely small deviation values observed for fiducial volume fits are a good indication that the net migration is very low. The deviation method also appears to agree fairly well with the fusion fraction method for stops within the fiducial volume where the fusion tagging has good efficiency. Fusion-induced rate shifts of -1.6 ± 0.9 and $-0.2 \pm 0.9 \text{ s}^{-1}$ are estimated for the P-Dir and road trackers respectively, where the small shift in the latter case results from a coincidental cancellation between outward migrations in Z and inward migrations across the Y fiducial boundaries.

The analysis in this chapter has focused on the R2015 data. Fitting the R2014 data finds a D value of $-(0.6 \pm 3.8) \times 10^{-5}$ for the P-Dir tracker with a single-electron cut. The full calibration procedure has not yet been carried out on the R2014 data, but using the R2015 calibrations this predicts a rate shift of $-0.2 \pm 1.2 \text{ s}^{-1}$. With such a small deviation value any differences in the calibration should have little effect, so this value is not expected to change significantly.

Chapter 14

ELECTRON INTERFERENCE

Electron interference generally has a much weaker effect on the muon decay rate than fusion interference does. However, electron interference also lacks many of the features that have been used to study and correct for fusion interference, such as the distinctive large energy deposition or characteristic time distribution. Furthermore, while only a few percent of events involve a fusion reaction nearly all muons decay and emit an electron, making it difficult to identify any sort of clean reference sample analogous to the non-fusion events. Thus, with fusion interference now well understood it is in fact electron interference which is the more concerning source of TPC tracking errors.

A characteristic feature of electron interference is a strong dependence on the Y component of the decay electron direction vector. Electrons emitted upwards travel away from the drifting charge cloud in the TPC, so only very early decays can cause interference while decays after the end of the clustering window have no effect. Electrons emitted downwards will instead catch up to the charge cloud, resulting in a tail to much later times. The maximum migration probability for downward-going electrons occurs for decays near the end of the clustering window, as that allows the electron track to cross the entire width of the window rather than starting in the middle.

This chapter will describe several ways in which electron migration can occur. It will then discuss a method for measuring the vertical asymmetry produced by electron interference, building on the early time deviation method developed in chapter 13. Migration probability modeling approaches are also used to constrain specific migration types and to check the plausibility of the asymmetry estimate.

14.1 Mechanism

Recall from chapter 5 that the ionization produced by a Michel electron in the TPC may be modeled as a continuous dE/dx contribution with the addition of occasional high-energy delta electrons. Here we will use ΔE_μ to denote the energy deposited by the dE/dx approximation, with ΔE_δ referring to the delta electron energy. To be detected and reconstructed by the MuSun analysis the electron must produce an energy pulse that exceeds a minimum amplitude threshold. This pulse-finding threshold varies between pads due to differences in pedestal and energy calibrations, but typically rejects pulses with energies below 60 to 100 keV. Since the decay electron is approximately a MIP it should deposit about 16.5 keV/cm, for a ΔE_μ of approximately 26 keV for tracks crossing a TPC pad. This is well below the pulse-finding threshold, so isolated electron pulses are only possible with the additional ΔE_δ contribution.




Figure 14.1: Diagrams of the track configurations that produce downstream, upstream, and pileup migrations.

The impact of these effects becomes more complicated in combination with the muon track and other possible pulses in the TPC. We identify four separate categories of electron interference, with somewhat different properties:

Downstream Electrons that produce an extra pulse downstream of the muon stop, extending the track. In this case the electron pulse may occur any time within the clustering window, but require a ΔE_δ contribution and are thus relatively rare.

Upstream Electrons which travel upstream can add to E_1 , potentially increasing it above the road threshold or the P-Dir ${}^3\text{He}$ fusion cut. There is no need for the additional ΔE_δ in this case, as ΔE_μ can add to the pre-existing muon pulse instead. This requires the electron track to directly overlap the muon pulse, so these events are expected to have an effective coincidence window similar to the muon pulse width rather than the clustering window.

Pileup Electrons on the stop pad may add to E_0 , which typically does not affect the event reconstruction for normal muon stops. However, in the special case where the muon stops just after passing the pad boundary the true E_0 pulse may be too small to pass the pulse-finding threshold. These events would typically be reconstructed in the previous pad row, but if the electron and muon pulses coincide the combined pulse can become large enough to detect. The time dependence of these events is expected to be similar to the upstream migrations since they both require the electron track to overlap with the existing muon pulse, although the solid angle acceptance is somewhat different.

Fusions A combination of fusion and electron interference may also occur, with the electron slightly increasing the energy of a normal fusion event. This is mainly a concern for the P-Dir tracker, where the proton direction cut is in close proximity to the migrating fusions and a small extra energy deposition could cause downstream-going p+t fusions to be misidentified as upstream-going. It would have the effect of making the separation between the proton direction bands time dependent.

Events with both a fusion and an electron migration appear to be fairly rare, and the agreement between the P-Dir and road tracker fits in chapter 13 indicates the possible electron migrations across the proton direction cut are not a major concern. Pileup electrons completely overlap the normal stop band in $E_0 - E_1$, as do a significant fraction of the downstream migrations, and so these events cannot be easily eliminated.

14.2 Event Tagging

Upstream and downstream electrons may be selected with E_0 vs E_1 cuts similar to those used to select fusions, shown in figure 14.2. The behavior of downstream electron migrations has already been described in chapter 9, and they are identified by the same low-energy cut used by the P-Dir tracker. This tag is only expected to include about half of the total downstream migration population.

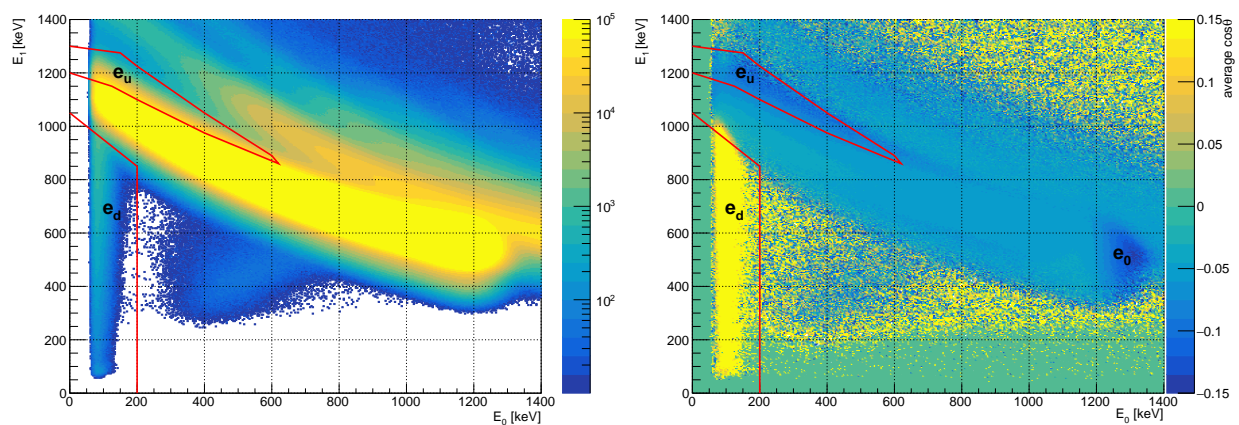


Figure 14.2: E_0 vs E_1 cuts to select upstream and downstream-going electrons, labeled e_u and e_d . On the left the color scale shows the number of events, while on the right it shows the average electron $\cos\theta$ value for decays within the first μs . Electrons adding extra energy to the stop row are also visible in the right panel, these are labeled e_0 .

The upstream electron cut aims to select events on the top edge of the muon stop band, which are likely to have an electron contributing to their high E_1 value. It focuses on low E_0 values as this corresponds to stops early in the pad which have a higher probability of an electron overlapping the E_1 pulse. This also avoids the ^3He fusion band, which starts to overlap the muon band at higher E_0 . The lower edge of the cut is tuned based on the second plot in figure 14.2, where the color scale shows the average electron $\cos\theta$ value for decays within the first μs . This makes the electron migrations clearly stand out from the

non-migrating events, and the tag is adjusted to pick out the region with a significant bias towards upstream-going tracks.

Electrons contributing energy to the stop row itself are also visible on the far right side of figure 14.2, labeled e_0 , where they can push E_0 to higher values than what is normally possible for a muon stop. Although these events appear quite different from the other tagged events, note that the addition of a small pulse on the pad row would shift these events to the top left side of the plot and fill in the low E_0 portion of the upstream tag region. The e_0 and e_u events are thus nearly identical, with many of the former in fact being upstream migrations where the final muon stop pulse was below the pulse-finding threshold. However, because of the poor separation between the muon band and the ^3He fusion band at the right side of the plot, these events are not used in the electron interference analysis.

The primary concern is actually the pileup case mentioned above, where the electron energy adds to a very small muon E_0 pulse and pushes it over the TPC pulse-finding threshold. In the E_0 vs E_1 plots these events originate from the cut off portion of the muon band at the far left, but are shifted to the right by the extra electron energy. They should therefore have a higher than average E_1 for events with the same observed E_0 value, but the shallow slope of the muon band means that the majority of these events are indistinguishable from normal muon stops. Pileup electrons are expected to exhibit similar behavior to upstream migrations, since both work by directly overlapping with the muon pulse and should thus have similar constraints on the decay electron time. However, since the TPC pulse finding involves an amplitude threshold rather than an energy threshold it may be more sensitive to horizontal electron tracks that deposit their energy in a single coincident pulse.

14.3 Time Distributions

Selecting events in the e_u and e_d tag regions and plotting the decay times for individual eSC segments reveals the expected vertical dependence, as shown in figure 14.3. The e_u region in particular has a large contribution from normal muon decays, so to show the effects of electron migration the expected muon lifetime component has been subtracted based on a




Figure 14.3: Decay times for upstream (left) and downstream (right) electron migrations, with the expected muon lifetime component subtracted to isolate the electron migration effect. The electron times are split by eSC segment as indicated. Upward-going electrons have little effect on the standard lifetime fits starting at $1 \mu\text{s}$.

fit starting at $10 \mu\text{s}$. This subtraction is only for visualization purposes and to estimate the number of real migrations in the tag region, and when performing rate fits the non-migrating events are included to avoid biasing the results.

For downstream migrations the signal in the upper eSC segments falls to zero around the end of the $2 \mu\text{s}$ effective clustering window, with the signal in the lower segments peaking at a similar time as expected. For upstream migrations these points are closer to $1 \mu\text{s}$, roughly consistent with the typical width of a muon pulse. This observation supports the model that upstream migrations primarily contribute energy to an existing pulse, rather than adding a separate pulse to the cluster. When looking at the sum over all eSC segments, the fast decay of the upward-going electrons largely cancels with the buildup of the downward-going electrons. Muon lifetime histograms will thus appear nearly exponential even with relatively large amounts of electron migration, but the disappearance rate will increase since the electron migrations are skewed towards early times on average. The deviation measurement discussed in chapter 13 is therefore insensitive to electron interference.

14.4 Early Time Asymmetry

Several attempts were made to develop a new measurement similar to the deviation, but optimized for detecting vertical asymmetry. The best results were obtained by simply reusing the deviation calculation, but with separate fits to the upper eSC segments (1-3, 14-16) and the lower eSC segments (6-11). For the upper segments the migrations primarily occur before the 1 μ s fit start time, so the fits find an accurate lifetime and the migration peak generates a positive deviation. For the lower segments the fitted rate is increased due to the long electron migration tail, with the buildup at early times then causing a negative deviation. We thus define the asymmetry A as the difference between these deviation measurements:

$$A \equiv D_{up} - D_{down}. \quad (14.1)$$

The horizontal eSC segments (4,5,12,13) are not included as they were found to add uncertainty without meaningfully increasing the difference between the fits.

This procedure is a bit odd, as it relies on an inaccurate initial fit to the downward-going electrons. Each fit also uses only 3/8 of the data, making the asymmetry substantially less precise than the deviation. It is possible that a more elegant procedure could be devised based on directly subtracting the original histograms, although a naive approach encounters difficulties due to the TPC frame and support structure blocking some electrons and reducing detection efficiency for downward-going electrons. Because the deviation is a ratio it is not particularly sensitive to these efficiency differences, allowing the values for upper and lower segments to be subtracted without any additional scaling. The subtraction also naturally cancels any uniform D value caused by fusion migrations, which is a nice property that would not necessarily be shared by an alternate measure of the asymmetry.

Figure 14.4 shows the results of applying this procedure to the electron tag regions. For the upstream migrations where the non-exponential behavior is mostly confined to the first microsecond, the standard fits starting at 1 μ s perform well without any modification. However, the more extended time distributions of the downstream migrations require a delayed fit start time of about 3 μ s to properly capture the asymmetry. Fitting with the




Figure 14.4: Fit ratio histograms using only upper (1-3, 14-16) or lower (6-11) eSC segments, for upstream (left) and downstream (right) electron migrations. Normal events have been added to stabilize the fits and adjust the rate shift to 1000 s^{-1} in both cases. A $3 \mu\text{s}$ fit start time is used for the downstream migration asymmetry fits.

standard start time would instead force the two curves to converge near $1 \mu\text{s}$, after which they would overshoot and oscillate back and forth. In contrast using a later start time for the upstream migrations achieves little except increasing the fit uncertainties.

To test whether the asymmetry can produce a consistent correction estimate for the muon decay rate, figure 14.5 shows the calibration values obtained for different asymmetry fit start times. The rate shift is determined using a standard $1 \mu\text{s}$ start time and is constant for all points, it is only the estimated asymmetry that changes. Once again $3 \mu\text{s}$ appears to be the optimal time, producing a consistent correction that can be applied to both types of electron migration. Unfortunately such a delayed start time excludes nearly 60% of the events included in the standard fits, and combined with only using a subset of the eSC segments this means the fits only have about 15% of the normal statistics. As a result the asymmetry values have large error bars, producing an uncertainty of 7 s^{-1} in the rate shift estimate even when ignoring any uncertainty in the calibration.




Figure 14.5: Rate shift vs asymmetry calibration as a function of fit start time. A single correction factor can account for both upstream and downstream migrations using a fit start time of $3 \mu\text{s}$.

14.5 Downstream Migration Correction

The issues with the asymmetry primarily stem from the late start time needed to consistently treat all electron migrations. To achieve a more precise correction, it is better to instead treat the downstream migrations separately. If these could be eliminated then the remaining upstream and pileup electrons could be corrected using the asymmetry method with the standard $1 \mu\text{s}$ start time. The total rate shift from subtracting all of the tagged forward electron migrations is only $12.3 \pm 0.1 \text{ s}^{-1}$, so it should not be hard to estimate their effect to the required accuracy.

To constrain the downstream migrations, we return to the idea of modeling the migration probability as a function of position. Recall from chapter 9 that the observed energies \tilde{E} for downstream migrations are

$$\tilde{E}_0 = E^e \qquad \tilde{E}_1 = E_0^\mu, \qquad (14.2)$$

where superscripts e and μ indicate electron and muon energies, respectively. We can there-

fore derive the E_1 distribution of migration events by multiplying the E_0 distribution from the previous pad with an estimate of the migration probability.




Figure 14.6: Extrapolated electron E_1 distributions and estimated uncorrected P-Dir migrations, for Z rows 2 (left) and 6 (right). Calibration is performed using the clean range between the high-energy muon stop band and the low-energy peaks due to decay in flight events or double-pad electron migrations.

To match the observed shape of the electron energy distribution, the migration probability as a function of stop depth is approximated with a two-parameter model including a solid angle correction factor and an overall efficiency term α . A description of this model is given in appendix L, but it may be tuned to give good agreement with the observed part of the forward electron migration band as shown in figure 14.6. The low-energy peak in the $Z = 2$ histogram is the result of residual decay in flight events, while the one for $Z = 6$ may be due to the electron migrating two pads in Z . These are avoided when calibrating α by focusing on the interval from 400 to 800 keV, which is typically clean. Using the estimated electron E_1 distribution we can estimate the number of events missed by the P-Dir electron migration cut, shown by the shaded regions in figure 14.6. This procedure is repeated for each Z row, with individual α 's to account for threshold variations between pads. Figure 14.7 shows the estimated α and the fraction of events tagged by the P-Dir $E_0 - E_1$ cut as a function of Z .

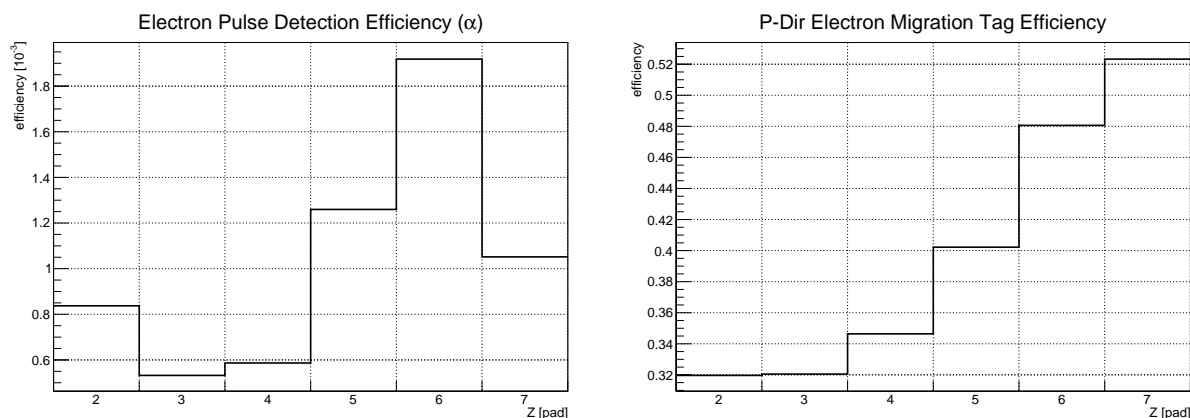


Figure 14.7: Estimated electron pulse detection efficiency α (left) and the P-Dir electron migration cut efficiency (right), as a function of Z .

Once again it is also necessary to include Y migrations in these models, as they can be just as important as the Z migrations. For the tagged downstream migrations it is straightforward to recalculate the stop position while ignoring the final electron pulse, and then simply count the number of events migrating across the fiducial boundaries. The untagged migrations are estimated by scaling the number of tagged events by the estimated tag efficiency for each Z row. Table 14.1 shows the number of tagged migrations and the estimated untagged migrations. Rate shifts are estimated by taking the ratio of the number of migrations to the total number of tagged events in the fiducial volume, again assuming a consistent time distribution for all events. Interestingly in this case there are actually more Y migrations than Z migrations, leading to a net outward migration opposite to what one would naively expect.

Quite a few simplifications were made in the migration model, namely that all forward electron migrations produce a similar lifetime shift and the correction may be estimated by simply scaling the shift caused by all the tagged events. The model also used a single-parameter solid angle correction, and although it appears to reproduce the observed energy distributions quite well it does not account for the different clustering probabilities or lifetime

	Fiducial Boundary	Migrations [10^3]			Effect	
		In	Out	Net	M / N_{tag} [%]	R [s^{-1}]
Tagged Events	Upstream Z	32	0	32	1.1	0.13
	Low Y	11	78	-66	-2.2	-0.27
	High Y	11	31	-20	-0.7	-0.08
	All	54	109	-55	-1.8	-0.23
Untagged Events	Upstream Z	67	0	67	2.2	0.28
	Low Y	14	97	-82	-2.7	-0.34
	High Y	15	40	-25	-0.8	-0.10
	All	97	137	-40	-1.3	-0.16

Table 14.1: Estimated forward electron migrations for the P-Dir tracker, split into tagged and untagged events. Rate shifts are estimated based on the number of migrations as a fraction of all tagged events in the fiducial volume.

shifts caused by different electron track angles. It is therefore somewhat difficult to evaluate the uncertainty for these results, and it would be reasonable to assign a large uncertainty estimate. However, because the final rate shift predictions are very small, even a 100% error is still fairly insignificant compared to the other MuSun uncertainties.

The downstream electron migration model also indirectly constrains the effects of events with both fusion and electron interference, which have so far been omitted from any of this analysis. The P-Dir and road trackers behave identically in the case of the pileup electrons or untagged forward electron migrations, only differing in their handling of the tagged downstream migrations and of fusion events. After the fusion correction in the previous chapter there was only a $0.50 s^{-1}$ discrepancy between the P-Dir and road trackers, and the estimated $0.23 s^{-1}$ from the tagged electron migrations reduces this to $0.27 s^{-1}$. Thus, while no attempt has been made to correct events with both fusions and electrons, this is clearly a small effect as well.

14.6 Pileup Electron Asymmetry Correction

The forward electron migrations had a small effect mainly because of the roughly 10^{-3} detection efficiency factor shown in figure 14.7, resulting from the fact that a high-energy delta electron is required to exceed the pulse-finding threshold. For pileup electrons on the stop pad the situation is very different, as the muon provides much of the required energy. With no need for a delta electron, the electron detection efficiency could potentially be increased by two orders of magnitude or more. The solid angle acceptance for pileup electrons should also be higher as they do not need to reach the adjacent pad, although they must overlap with the muon pulse so they must occur in a narrower time window. Fortunately the pileup electrons are only a problem for stops within about the first 0.5% of a pad ($80 \mu\text{m}$), as past that point the muon will generally exceed the pulse finding threshold by itself any additional electron energy no longer affects the tracking.

Now that the downstream electron migrations are fairly well understood, a portion of the tagged downstream migrations may be added back to the standard histogram to cancel the estimated net out-migration. Calculating the asymmetry for this sample and for the upstream tag veto results in the rate shifts and asymmetry estimates in table 14.2. The upstream veto tag also affects the measured D by selecting some events at the bottom edge of the ^3He fusion band, and so this is corrected using the rate shift vs deviation relation observed for the tagged helium events. To be conservative the uncertainty has been inflated to account for up to 50% non-migrating ^3He fusions that may have effects more similar to the p+t fusions instead, which still results in a relatively low fractional uncertainty.

Dividing the corrected rate shift by the measured asymmetry for the upstream veto fit yields a slope of 43.5 ± 2.5 . Applying this slope to the asymmetry of the normal sample yields an estimated rate shift of $3.3 \pm 2.8 \text{ s}^{-1}$. This uncertainty is still a bit large given a total uncertainty target of 4 s^{-1} , but it is much better than the earlier version with the $3 \mu\text{s}$ fit start time.

It is possible that pileup electrons behave somewhat differently than upstream migrations.

Sample	Fitted R [s^{-1}]	D [10^{-6}]	Corrected R [s^{-1}]	A [10^{-6}]
Standard	0	-51 (28)	1.55 (91)	74 (65)
e_d Corrected	0.16	-51 (28)	1.71 (91)	76 (65)
e_u Veto	-71.10 (47)	-388 (28)	-64.7 (3.9)	-1453 (65)

Table 14.2: Fitted rates and asymmetries for the upstream electron migration calibration. The corrected sample refers to a sample with 1.3% of tagged downstream migrations added back in to correct for the estimated net out-migration from these events. The e_u veto sample has a relatively large D due to contamination from the ^3He band, this is used to correct the fitted rates before performing the asymmetry calibrations.

Perfectly vertical tracks have the strongest asymmetry effect and are included in the pileup case, while upstream migrations require tracks to be at least somewhat slanted in order to reach the next pad row. The difference should not be particularly large though, as the majority of the tagged events should have stop depths below about 2 mm so electron tracks only need to be tilted a few degrees from vertical in order to cross the pad boundary. The interaction between the electron and muon pulses is also different, with upstream migrations being sensitive to any electron energy added to the muon pulse while the pileup events are only really sensitive to the peak amplitude. This may instead reduce the average asymmetry for pileup electron migrations, although this is also expected to be a minor difference because of the charge integration and shaping performed by the TPC amplifier chain. Similarly, the TPC amplifiers include baseline restorer systems which will have different effects on small or large pulses, but this works on a relatively long time scale and should make little difference for electron pulses coincident with the muon. The uncertainty of the asymmetry method is already completely dominated by the measurement of the asymmetry itself rather than the calibration, so while it would certainly be better if the calibration were performed using real pileup electrons this is a relatively minor concern.

14.7 Pileup Electron Modeling




Figure 14.8: Observed muon E_0 distributions and estimated losses at low energy based on a quadratic fit from 150 to 400 keV, for Z rows 2 (left) and 6 (right).

Because the asymmetry method is a bit complicated, it is also helpful to test these results against the migration modeling approach. In this case estimating the migration probability is difficult, as it would require knowledge of the muon and electron energy distributions as well as how the overlapping pulses interact with the pulse-finding algorithm. However, it is relatively straightforward to at least estimate the number of muon stops with a final pulse below threshold, giving an upper bound on the pileup electron effects. This is achieved by extrapolating the observed muon E_0 distribution for each Z row to zero energy, using a simple quadratic fit to approximate the shape as shown in figure 14.8.

To account for Y migrations we assume the missing events behave similarly to observed events with E_0 below 200 keV, justified by the fact that this energy corresponds to a muon range of only 0.5 mm ($<4\%$ of a pad length). Recalculating the stop position for these events while ignoring the final low-energy pad allows us to count the number of migrations across the fiducial boundaries. Scaling these results by the estimated fraction of missing low-energy events in each Z row yields the numbers of candidate migration events listed in table 14.3

Fiducial Boundary	Migrations [10^6]			Effect		
	In	Out	Net	M / N_{tag} [%]	A [10^3]	R [s^{-1}]
Upstream Z	8.9	0	8.9	134	2053	89
Low Y	0.2	6.8	-6.6	-99	-1515	-66
High Y	0.7	1.8	-1.1	-16	-251	-11
All	9.8	8.6	1.2	19	286	12

Table 14.3: Candidate electron pileup events, and associated asymmetries and rate shifts assuming a 100% migration probability. These values were obtained by comparing to the number of upstream migrations estimated to be within the tag region, and assuming similar time dependence between the upstream and pileup migrations.

The tagged upstream migration events are used to estimate the effect of a given number of migrations, again assuming the pileup electrons behave similarly. However, based on the time distribution of the tagged upstream migration events only about a quarter of them are actually electron migrations, giving a total of 6.6 million tagged upstream migration events responsible for producing the rate shift in table 14.2. It appears that if all the missing events had a 100% migration probability due to pileup electrons then it should produce about a 12 s^{-1} shift. Comparing this to the asymmetry estimate, it would seem that approximately one quarter of these candidate events actually migrate. The close proximity to the upstream pad boundary means half of the decay electrons should immediately cross into the next row, so if these events are neglected then half of the remaining events would need to cause a migration. This seems rather high, but may be possible considering electron tracks can deposit about 30 keV on the stop pad, accounting for about half of the required energy to reach the threshold. The large cancellation between the estimated in- and out-migrations also means that the observed effects could be produced by a significantly smaller fraction of migrations biased in one direction.

Currently this model is only meant to provide a very rough estimate of the maximum size of the effect, but there are some ideas that may improve this approach enough to provide an independent test of the asymmetry method. Studying the E_0 distributions from figure 14.8 while imposing electron track angle cuts may allow for an estimate of the actual number of migrations, since only downstream-going electrons should cause the effect. Alternatively, if the electron pulse amplitude distribution were known then it could be added to the extrapolated E_0 distributions to determine the number of events crossing the threshold, although again the dependence on amplitude rather than energy makes this somewhat difficult to estimate. Some followup Monte Carlo analysis is planned to test these ideas and also evaluate the consistency with the asymmetry measurements.

14.8 Conclusion

Building on the early time deviation analysis for fusion interference developed in chapter 13, a new approach to quantifying electron migrations has been developed. By dividing the decay time distribution out by eSC segment, one observes very different shapes for upward or downward-going electron tracks. The asymmetry between these two signals provides a measure of the migration fraction, and the effect has been calibrated by tagging likely migration events using a set of E_0 vs E_1 cuts similar to those used in the fusion analysis.

Several types of electron interference have been identified, with the two most important types being electrons which add an additional pulse to the end of the track and electrons which add to a small muon E_0 pulse which would otherwise be below the pulse-finding threshold but is made detectable due to the electron energy. These two migration types have different associated time distributions, so while it is possible to correct for both with a single asymmetry measurement this requires a late fit start time resulting in large uncertainties. Instead the downstream migrations are treated separately by modeling the migration probability as a function of position, and result in a very small correction term as it is quite rare for an electron to independently produce a detectable pulse. The remaining pileup electrons are then estimated using the asymmetry method, with a standard $1 \mu\text{s}$ fit start time producing

much more precise results. The asymmetry was also compared against an estimate of the number of events with low-energy muon pulses that could potentially be affected by pileup electrons, finding that approximately 1/4 of these events appear to migrate which seems high but not entirely outside the realm of possibility. Table 14.4 shows the final estimated rate shifts due to electron interference.

Tracker	Rate Shift [s^{-1}]		
	Downstream	Pileup	Total
P-Dir	-0.16	3.28 ± 2.84	3.1 ± 2.9
road	-0.39	3.28 ± 2.84	2.9 ± 2.9

Table 14.4: Estimated R2015 rate shifts due to electron migration. The pileup events dominate the total uncertainty, even when assuming 100% uncertainty in the downstream migration model.

In the R2014 dataset we find a larger asymmetry of $(1.88 \pm 0.88) \times 10^{-4}$. The full calibration procedure and migration modeling has not been performed for R2014, but applying the R2015 calibrations and assuming negligible downstream migrations produces an estimated rate shift of $8.2 \pm 3.9 s^{-1}$. The large cancellations estimated in table 14.3 mean the difference in estimated rate shift for R2014 and R2015 could be explained a relatively small shift in the muon stop distribution.

Chapter 15

CONSISTENCY CHECKS

Individual MuSun lifetime fits typically produce reasonable results, with reduced chi squared values close to 1 and no apparent problems with the fit residuals as shown in figure 15.1. However, it is possible for very subtle effects to cause fitted rate shifts at the 10^{-5} level precision required for the MuSun experiment without any obvious effect on these basic diagnostics. For example, in chapter 12 the background slope was shown to produce a rate shift of hundreds of s^{-1} in 2e events, but the 2e residuals still appear completely flat. This chapter will cover several important consistency checks used to ensure the stability and reliability of the fit results, including dataset comparisons, variations in the fit range, and scans over the muon stop position and decay electron direction.




Figure 15.1: Fits to the R2014 1e and 2e lifetimes (left), and the normalized fit residuals (right). No obvious problems are visible in the residuals, even for the 2e fit with a significantly enhanced background slope.

As discussed in previous chapters, a loose fiducial stop definition has been adopted involving a Y cut and upstream Z cut, but no X cuts or downstream Z cut. Instead, a first

row X and Y cut are used to focus the beam and limit the number of muons escaping the sides of the TPC. Unless otherwise stated, all fits use the P-Dir tracker for the fiducial cut, impose a single-electron cut, and use a standard three parameter exponential fit function in the range $[1, 24] \mu\text{s}$.

15.1 Dataset Consistency

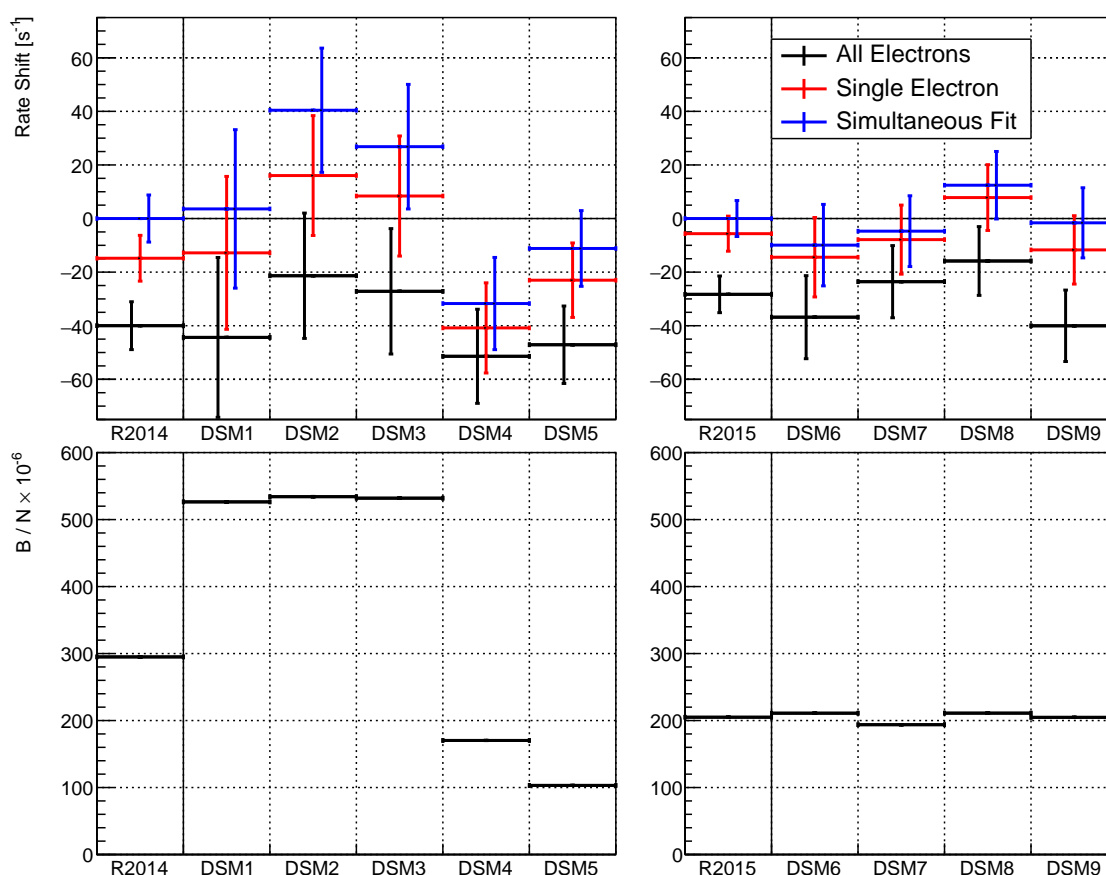


Figure 15.2: Fitted rate and background parameters for each of the large datasets. The background parameter, B , has been normalized to the estimated number of muons, N .

Recall that the R2014 data was split into five sub-datasets, while the R2015 data was split into four. The R2015 datasets are all quite similar, but the R2014 datasets had different

Fit	χ^2/NDF	
	R2014	R2015
All Electrons	1.59/4	2.14/3
Single Electron	5.68/4	1.81/3
Simultaneous Fit	8.32/4	1.54/3

Table 15.1: Chi squared values for fitting the individual sub-dataset lifetimes to a constant.

beam conditions and in particular DSM1-3 had problems with the upstream beam collimator resulting in much higher backgrounds. This issue has already been studied fairly extensively in chapter 12, but figure 15.2 shows a comparison between fits to all nine datasets. Because the R2014 and R2015 datasets are relatively blinded, the results are shown as shifts relative to the full-statistics fits.

The fit including all electrons appears relatively consistent, with the high-background DSM1-3 having moderately higher fitted rates. The single-electron fit and the simultaneous fit to 1e and 2e events shift the fitted rates up, but have a larger effect on the high-background datasets and exacerbate the discrepancy. As a result there is actually poorer consistency between the R2014 datasets when using the improved fits meant to be insensitive to the background slope, as shown in table 15.1. Muon pileup tends to increase the fitted rate, so this discrepancy may be a result of additional beam muons and worse kicker extinction during DSM1-3.

15.2 Fit Range Variations

The standard fit range of $[1,24] \mu\text{s}$ is chosen to be as large as possible while avoiding the kicker step at 800 ns and the end of the event window at $25 \mu\text{s}$. However, these boundaries are somewhat arbitrary, and ideally the fit results should not depend on the exact choice of fit range. Expanding the fit range would require hardware changes to decrease the kick

delay or increase the length of the window, but we can easily shrink the fit range to test for consistency.

Because this procedure uses a subset of the full data, the fit results are highly correlated and cannot be compared using their statistical error bars. A separate figure of merit must be used to determine whether the scan points are statistically consistent. Assuming two lifetime fits obtain rates R_1 and R_2 , the relative uncertainty between the two rates may be written

$$\delta(\Delta R) = \sqrt{\left(\frac{\partial \Delta R}{\partial R_1}\right)^2 \delta_1^2 + \left(\frac{\partial \Delta R}{\partial R_2}\right)^2 \delta_2^2 + 2 \frac{\partial(\Delta R)}{\partial R_1} \frac{\partial(\Delta R)}{\partial R_2} \text{cov}(R_1, R_2)}, \quad (15.1)$$

where δ_i are the statistical uncertainties given by $\lambda/\sqrt{N_i}$, and $\Delta R = R_1 - R_2$. If the second start time is later, such that N_2 is a subset of N_1 , then the covariance becomes $\text{cov}(R_1, R_2) = \delta_1^2$, and the uncertainty simplifies to

$$\delta(\Delta R) = \sqrt{\delta_1^2 + \delta_2^2 - 2\delta_1^2} = \sqrt{\delta_2^2 - \delta_1^2}. \quad (15.2)$$

The allowed 1σ deviation from the standard fit will be shown as a red shaded region in the following plots.

15.2.1 Start Time

Several systematic effects are sensitive to the fit start time, and in particular the background slope makes a large difference as later start times effectively reduce the signal to background ratio of the fit. Figure 15.3 shows start time scans for all electrons and for single-electron events, as well as the results from the simultaneous fit method. In the R2014 dataset the fitted rate for all electrons trends downward with later start times, and the single-electron cut and simultaneous fits correct this to produce relatively flat trends. However, because the R2015 dataset has a smaller slope and less background it is fairly flat to begin with, and the single-electron cut and simultaneous fit instead result in an upward trend.

The TPC track interference correction has not been included in figure 15.3. The estimates from chapters 13 and 14 combine to an 8 s^{-1} correction for R2014 and 1.5 s^{-1} for R2015.




Figure 15.3: Rate shift vs fit start time for R2014 (left) and R2015 (right). The fits in row (a) include all electrons, (b) imposes a single-electron cut, and (c) uses the simultaneous fit to 1e and 2e histograms introduced in chapter 12. The red shaded area indicates the allowed statistical 1σ deviation from the standard fit starting at $1 \mu\text{s}$, and all fits stop at $24 \mu\text{s}$.

Tracking interference is primary relevant at early times, so these positive shifts would result in a downward slope in the start time scans as the interference becomes weaker. Correcting this effect would thus shift the R2014 start time scan further upward, putting it more in line with the R2015 results.

Some degree of upward slope is expected in the start time scans, as the quartet hyperfine state is depleted. At $1 \mu\text{s}$ the population should be approximately 6.5% quartet state and 93.5% doublet state, for an instantaneous capture rate of roughly 375 s^{-1} . This would be a -25 s^{-1} shift from the expected capture rate, and could potentially explain the majority of the observed trend. However, it is only a rough estimate and a more careful calibration is required to estimate the effect of the muon kinetics on the simple exponential fit function.

15.2.2 Stop Time




Figure 15.4: Rate shift vs fit stop time for R2014 (left) and R2015 (right). The red shaded area indicates the allowed statistical 1σ deviation from the standard fit ending at $24 \mu\text{s}$, and all fits start at $1 \mu\text{s}$.

The end of the fit window is intended to be after nearly all of the muons have decayed, and the fit in this region is dominated by accidental backgrounds. Figure 15.4 shows the

effects of varying the fit stop time, and as one would expect there is little effect on the fitted rate. At fit stop times before about $16 \mu\text{s}$ the R2015 scan does drift a bit outside the expected 1σ region, but such a short fit window would not be advisable in any case because the background-dominated late times are excluded.

15.3 TPC Position

Ideally the fitted muon decay rate should be independent of the muon stop position as long as it stops within the deuterium gas. As shown in previous chapters this is not the case, with tracker interference effects making the fits very sensitive to the number of migrations across the fiducial boundary. This effect is exacerbated in fine-grained scans, which impose a series of much narrower cuts passing through the middle of the muon stop distribution. Note that migrations across these artificial internal boundaries have no effect on the full fits, so while these scans provide a useful diagnostic to quantify tracker performance they have relatively little bearing on the final result. Similar effects may also be seen in the Monte Carlo tracker testing from section 9.5, where there were a large number of migrations inside the fiducial volume but a much lower net effect.

15.3.1 Stop X

Figure 15.5 shows the fit results for each TPC X row in the R2015 data. Fiducial volume cuts have been applied in the Y and Z dimensions to isolate the effects of the X position. The middle four rows appear to show a slight upward trend in the fitted rate, this most likely indicates the muon beam is somewhat angled in X resulting in a coupling between the X and Z stop positions. The edge rows in X find higher fitted rates, particularly row 5. It is tempting to associate this with wall stops, but recall from chapter 11 that the number of wall stops in these rows is only moderately higher than in the center of the TPC. This is instead an artifact of the first row X cut, which requires all tracks to start in the middle four rows. As a result there is a relatively sharp cutoff in the true stop X distribution, making these edge bins very sensitive to interference effects that broaden or shift the distribution.

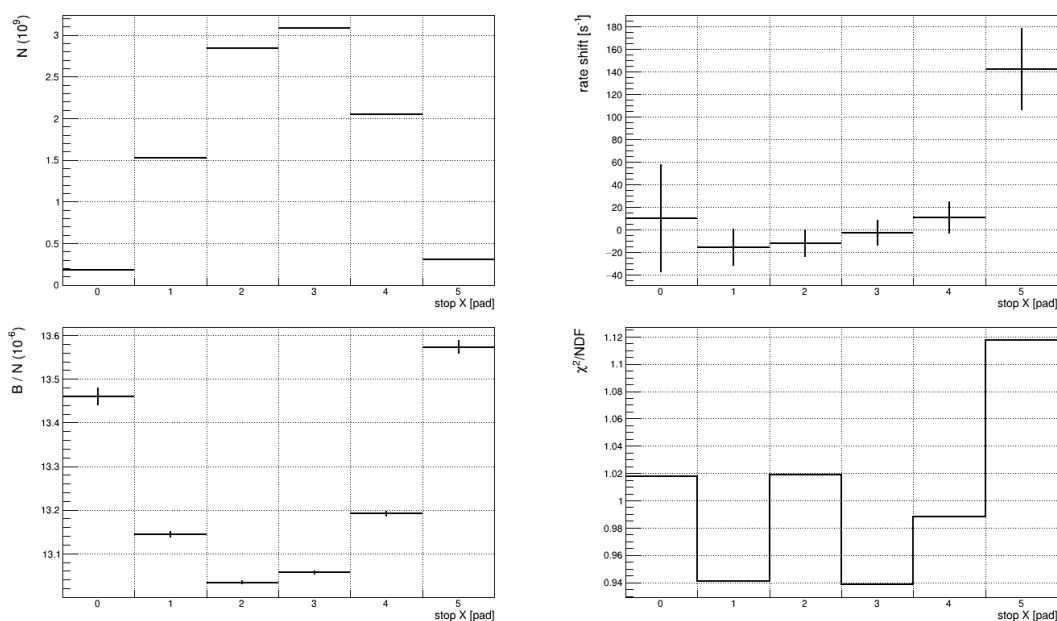


Figure 15.5: Lifetime fit results versus muon stop X row in the TPC. The fitted rate is displayed as shifts relative to the fit to all X rows. Fiducial volume cuts have been applied to the Y and Z axes.

15.3.2 Stop Y

The stop Y position is determined by the drift time, so there are no discrete pads as in X and Z. However, to obtain adequate fit statistics the Y position is split into 10 mm bins for the purpose of these scans. Figure 15.6 shows the fit results as a function of Y stop position for the R2015 data. As in the X scan, the edge bins outside of the 15-56 mm first row Y cut have large migration effects due to the lack of normal stops stabilizing the fits. The low Y bin is more strongly effected because the beam has a slight downward angle, so Z reconstruction errors will have a much larger Y effect near the bottom of the TPC. In this case the central bins have no obvious trend in the fitted rate, possibly because the tracking algorithm uses the Y slope to produce a more accurate stop Y estimate, while no such slope information is available in X.

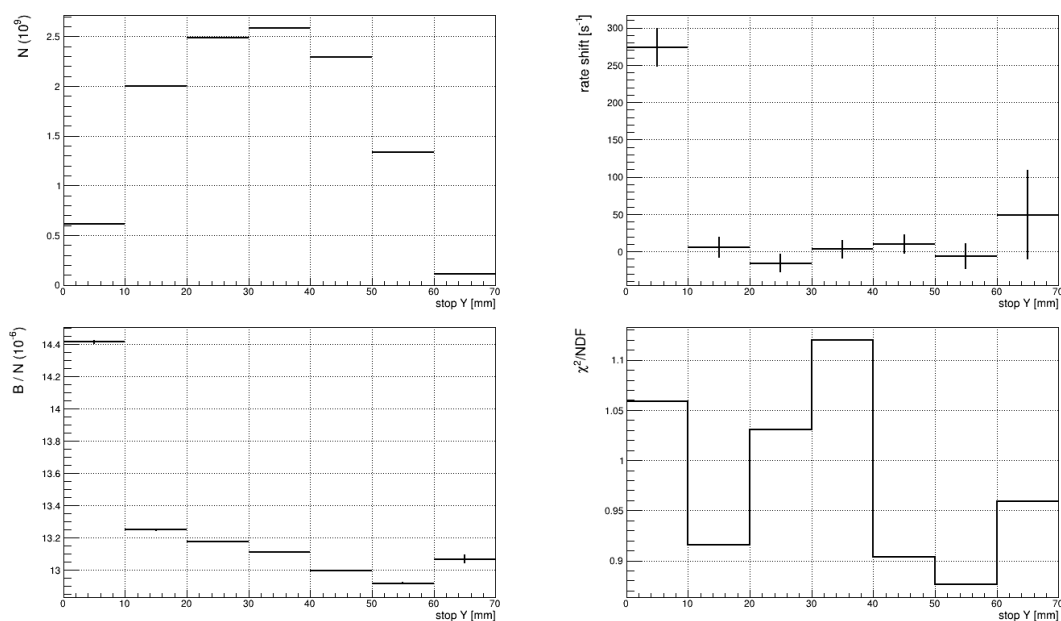


Figure 15.6: Lifetime fit results versus muon stop Y in the TPC. The fitted rate is displayed as shifts relative to the standard fiducial Y cut from 15-56 mm. Fiducial volume cuts have been applied to the X and Z axes.

15.3.3 Stop Z

Figure 15.8 shows the fit results as a function of Z stop position for the R2015 data, using the P-Dir tracker. The Z dimension is far more sensitive to tracking interference issues than X or Y, and in fact much of the observed X and Y behaviors are side effects of coupling to the stop Z position. Even with the improved P-Dir tracker there is still a shift of approximately $100 s^{-1}$ between Z rows at the front and back of the TPC. This primarily is only a problem with the R2015 dataset, while R2014 seems much more consistent as shown in figure 15.7.

As shown in the left panel of figure 15.9, the Z dependence of the P-Dir tracker is reduced by about a factor of 4 relative to the road tracker. Applying a fusion interference correction based on the early time deviation shifts the road tracker results into good agreement with the P-Dir results, but has little effect on the P-Dir tracker apart from the final Z row where

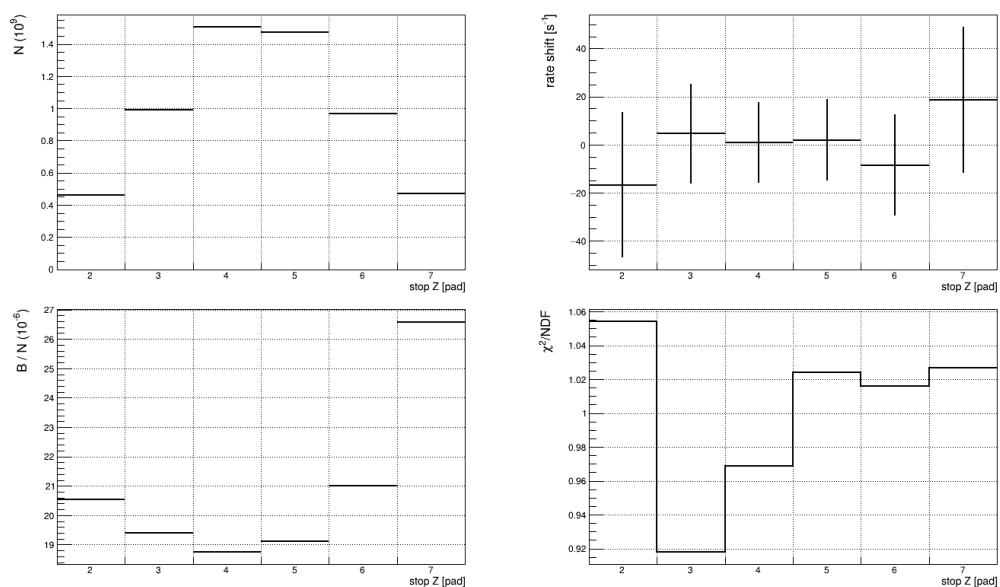


Figure 15.7: Lifetime fit results versus muon stop Z row in the TPC, for the R2014 dataset.

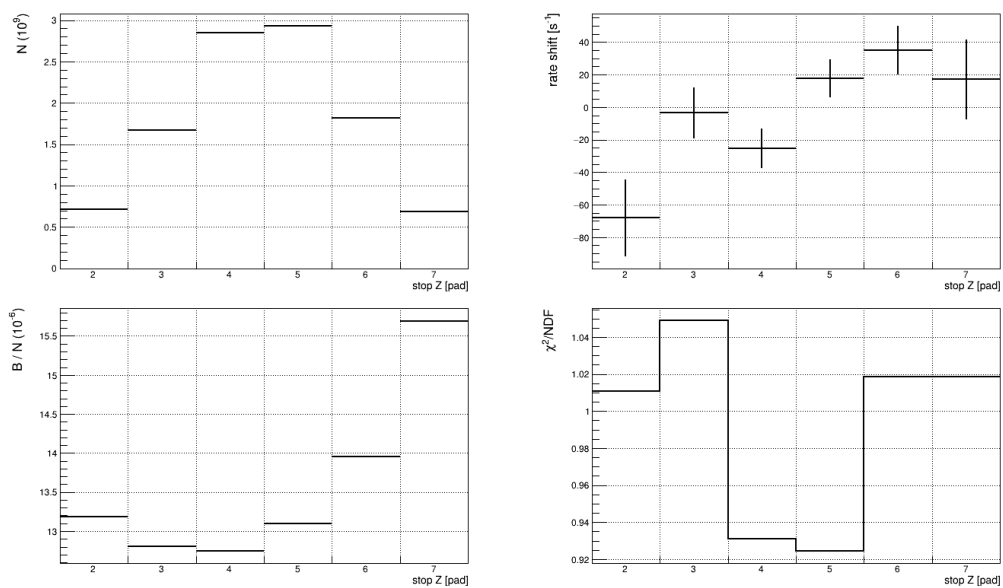


Figure 15.8: Lifetime fit results versus muon stop Z row in the TPC, for the R2015 dataset.

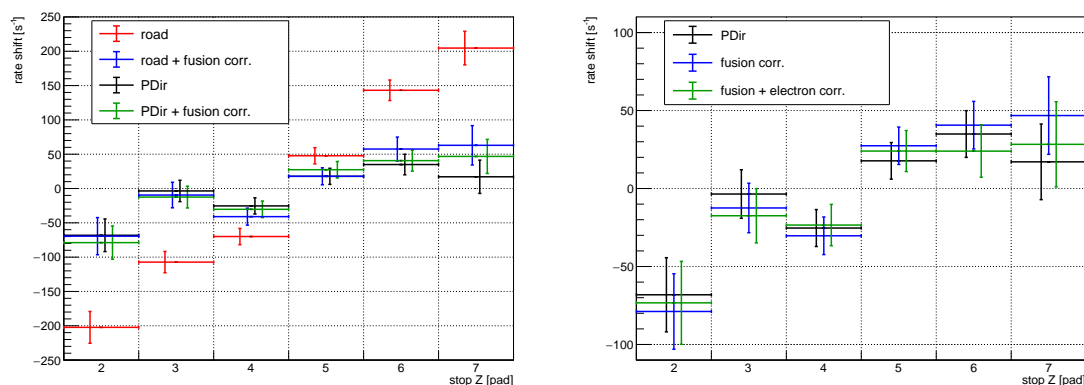


Figure 15.9: Left: Fitted rate as a function of Z stop for the road and P-Dir trackers, with and without fusion interference corrected based on the early time deviation method. Right: Fitted rate as a function of Z stop for the P-Dir tracker, with the fusion interference correction and an electron interference correction based on up/down asymmetry.

punch-through muons cause the P-Dir cuts to lose efficiency. Thus, it appears that the P-Dir tracker is correctly reconstructing fusion events, and the remaining trend is a separate issue. The right panel of figure 15.9 shows the electron interference correction based on the asymmetry measurement, which also has relatively little effect.

The upward slope in the R2015 data is thus fairly puzzling, as both the fusion and electron interference estimates indicate there should be little migration. The trend would be consistent with pileup electron migrations, and may indicate a problem with the assumption that these events behave similarly to the upstream electron migrations. Alternatively, there could be cancellation between different types of fusion events, leading to a significant rate shift without a corresponding early time deviation. This can be checked using the TPC fusion tags, which indicate there should be less than a 5 s^{-1} shift across the TPC due to changing ^3He and double-fusion fractions. Further testing comparing the R2014 and R2015 results would be helpful for identifying the cause of this effect.

15.4 Electron Direction

Finally, plotting the fit results as a function of decay electron emission direction may reveal effects related to the beam background, electron interference, or detector inefficiencies. The standard MuSun analysis uses only the eSCs and not the full electron tracks, but some limited directional information is still available.

15.5 eSC Segment




Figure 15.10: Fitted rate shift for each eSC segment relative to the fit to all segments, for the R2014 (top) and R2015 (bottom) datasets. A constant fit to the rate shifts is shown in red, and are close to $\Delta R = 0$ indicating good agreement with the full-statistics fits. However, the chi squared values of the constant fits are fairly poor, as indicated on the plots.

The azimuthal angle of the electron track may be obtained by fitting individual eSC segments, as shown in figure 15.10. These produce somewhat high chi squared values when fitting to a constant, but this appears to come largely from random scatter rather than a strong asymmetry effect. For instance, the up-down asymmetry from electron interference should produce a higher fitted rate in the lower eSC segments (6-11). While segments 8 and 9 have the highest fitted rates in R2015 and R2014, respectively, the behavior is not particularly consistent either between nearby segments or between runs.

15.6 eSC Z Position

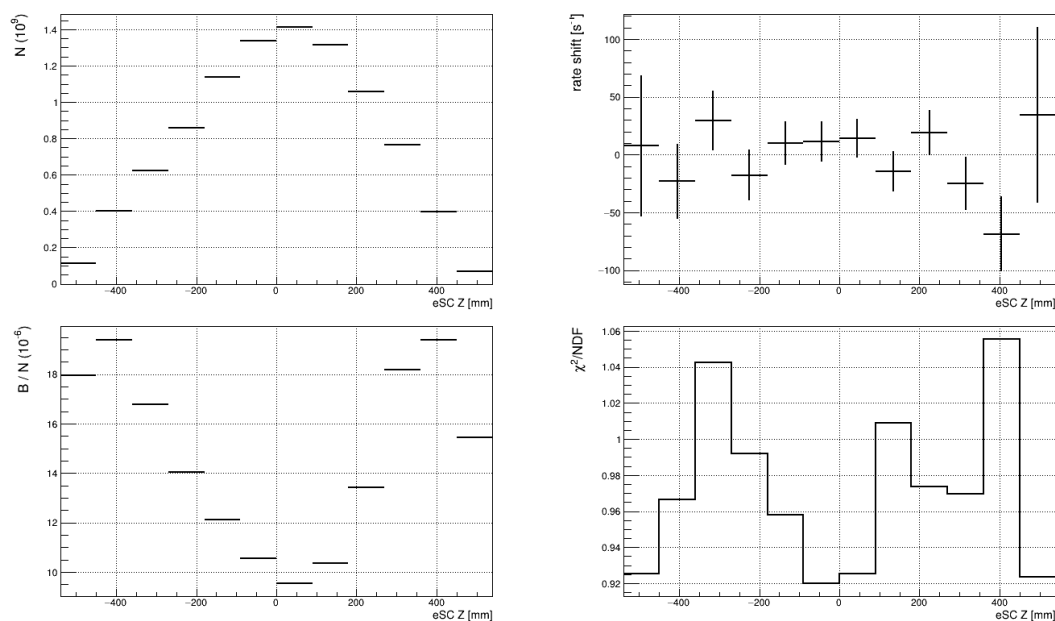


Figure 15.11: Lifetime fit results versus eSC Z position.

The Z position of the electron detection in an eSC segment may also be determined by the time of flight delay between the upstream and downstream PMTs. Figure 15.11 shows the fit results as a function of eSC Z for the R2015 dataset. The scintillators are only 900 mm in length so stop positions beyond ± 450 mm are unphysical, although a small number of events

are mis-reconstructed in this range. The eSC Z coordinate may in some cases be used as a proxy for the polar angle θ of the electron track, but this is only valid for tracks originating from the TPC. Beam background electron tracks predominantly point in the downstream direction, but may impact the scintillator at any point resulting in a fairly symmetrical B/N distribution. The N and B plots indicate that the TPC is well centered within the electron detectors, and eSC Z appears to have no particular effect on the fitted rate.

15.7 Summary

Despite the efforts in previous chapters to constrain the important systematic effects, some of the standard parameter scans do not appear to be completely consistent. The start time scans are most concerning, as these represent rate shifts in the full fit rather than small subsets of the data as with most of the other scans. The drift in the start time scans may be largely explained by the expected change in hyperfine state population at early times, but some additional followup analysis is required to properly quantify this effect.

There also appears to be some residual tracking interference in the R2015 muon stop Z scans, even after the tracker improvements and migration corrections discussed in previous chapters. This does not occur in R2014, and some additional testing should be performed comparing the two datasets. However, once again it is important to remember that these fine-grained TPC scans greatly exaggerate the size of the migration effects, and the agreement between the road and P-Dir tracker results is a good sign that the migrations across the fiducial boundary are minimal.

Finally, the comparisons between the R2014 sub-datasets and between fits in different eSC segments showed somewhat poor chi squared values. This may require some error inflation for the final results, and the pileup protection estimates must be double-checked to see if they are related to the rate shifts seen in the first three high-background sub-datasets.

Chapter 16

CONCLUSION

This chapter summarizes the key systematics studies presented in this thesis, and discusses the outlook for producing a final result from the experiment.

16.1 Systematics

The MuSun experiment uncertainty target of 1.5% corresponds to $6 s^{-1}$, which may be split into $4 s^{-1}$ statistical and $4 s^{-1}$ systematic uncertainties. The main systematic effects and their estimated rate shifts and uncertainties are summarized in table 16.1. Note that the analyses presented in this thesis were primarily focused on the R2015 dataset, while the R2014 data has been studied in less detail. This thesis has established uncertainties close to the target level for both production datasets, with the R2015 dataset constrained to below the uncertainty target.

Overall, the analysis of the MuSun experiment has made substantial progress, with many long-standing issues having been resolved by the author. The difference between the eSC-only electrons and full tracks including the ePCs was traced to a software bug in the electron tracking, and tracking improvements have corrected the majority of the discrepancy. The rate shift when selecting events with different electron multiplicities has been identified as a background slope and appears to be well constrained, although the source of the slope is somewhat uncertain as discussed in chapter 12. The fusion interference problem has also been largely resolved via tracking improvements described in chapter 9 and the new lifetime deviation method for measuring fusion migrations described in chapter 13.

Two more major issues have also been addressed recently by other collaboration members. The improved in-situ impurity analysis performed by Duncan Prindle has produced a very

Effect		Shift (s^{-1})		Uncertainty (s^{-1})	
		R2014	R2015	R2014	R2015
Beam	Pileup	1.5		1.5	
	MORE	≈ 0		0.25	
	Slope	-9.8	-5.7	2.1	1.5
TPC	Fusion	-0.2	-1.6	1.2	0.9
	Electron	8.2	3.1	3.9	2.9
Purity	Wall Stops	≈ 0		≈ 0	
	Isotopic	0.8		0.2	
	Chemical	5.2	4.1	0.6	0.5
Electrons	Afterpulses	-0.54		0.64	
	Tracking*	≈ 10		-	
Total		5.16	1.66	4.9	3.8

Table 16.1: Summary of the major systematic effects in the MuSun experiment. Note that the electron tracking uncertainty is not included in the total as it does not apply to the main eSC-only analysis. Some R2014 estimates are preliminary and subject to change, in particular the the electron migration rate shift.

precise estimate of this contribution. Finally, the puzzling effects produced by the TPC first row energy cut have been traced to crosstalk between TPC pads, and a model of the effect is now well developed.

Several other potential systematic errors have been shown to have minimal effect. The TPC tracking nearly eliminates wall stop effects within the fiducial volume, and as discussed in chapter 11 any remaining contribution is negligible compared to the other experimental uncertainties. A pileup protection analysis developed by Rachel Ryan and repeated by the author for R2015 shows very good pileup detection efficiency, and produces a relatively

small estimate for the uncertainty in the fitted rate. The isotopic purity requirements are much less stringent than the chemical purity requirements, and are well constrained by the gas chromatography measurements taken during the production runs. Finally, effects such as eSC afterpulsing or a fast decay tail following the kicker step both appear to have no significant effect, with less than 1 s^{-1} uncertainty.

16.2 Limitations

With so many of the major systematic effects now well controlled, electron interference appears to be the primary remaining issue. This has always been difficult to disentangle from similar fusion interference effects, and the pileup electron migrations in particular are very challenging to study directly. Chapter 14 presented an electron interference analysis based on the vertical asymmetry characteristic of electron interference, which appears to give reasonably precise results. However, this relies on the assumption that pileup electron migrations have similar properties to the upstream migration events that may be selected by TPC energy cuts, which has not been rigorously demonstrated.

Some of the results from the consistency checks in chapter 15 also require further attention. The drift in the start time scans must be compared against the full muon kinetics model, and if any inconsistency remains it must be studied in more detail. Additional comparisons between the muon tracking in R2014 and R2015 should also be performed to better understand the rate shifts observed in the R2015 stop position scans. Finally, error inflation may be required to account for the somewhat poor chi squared values obtained when fitting some scan results.

16.3 Outlook

Additional testing is required to investigate these remaining questions. In particular there are several ideas to better constrain electron migration, by testing the reliability of the asymmetry measurement using Monte Carlo data and by improving the migration model to provide an independent estimate of the effect. Performing fits with alternate fiducial

volume cuts would also provide an independent test of the number of net migrations, and in particular a looser fiducial Y cut should be safe to use as shown in chapter 11. There was also a previous analysis by Rachel Ryan attempting to evaluate the effect of artificially increasing the E_0 threshold, but this was difficult to interpret due to also producing a large amount of fusion interference. An improved version of this analysis using E_0 and E_1 cuts to apply the threshold specifically to non-fusion events may be quite informative, and should be relatively simple to implement using the existing migration analysis machinery. These tests will be performed this summer shortly after the completion of this thesis.

One final pass of the full analysis software chain is currently planned, incorporating various minor software improvements that have been implemented but not yet run on the full statistics. Similarly, this final pass would be an opportunity to adjust some analysis parameters, in particular tuning the P-Dir tracker parameters to hopefully eliminate the remaining interference effects observed in R2015. We have the supercomputer allocation available to perform this pass, and as the planned changes are either already implemented or involve simple parameter adjustments this should require minimal additional software development work. This pass is not expected to significantly change most of the results presented above, instead serving primarily as a confirmation that the experiment is not sensitive to subtle analysis details. The pass should also be performed this summer, after the completion of the electron interference and other followup tests.

Some dataset comparisons were presented in chapter 15, but the different blinding of the R2014 and R2015 made direct comparison impossible and instead only offsets from the relevant full-statistics fits were shown. The next step for the experiment would be to perform a relative unblinding between the R2014 and R2015 datasets, allowing comparisons between all nine sub-datasets while still maintaining blinding of the results as a whole. It may be beneficial to perform the relative unblinding before the planned analysis pass, in case there are any unexpected discrepancies that should be addressed in the final analysis.

If all goes well with the remaining analysis and no major discrepancies are revealed by the relative unblinding, a full unblinding of the results will be in order. After the completion of

these final checks, the analysis reports will be sent to the collaboration members for review. An unblinding meeting is tentatively planned for this fall, along with a PRL publication announcing the result. A more detailed publication discussing the experimental methods will follow shortly afterward.

16.4 Summary

MuSun will be the first experiment to measure the doublet muon capture rate on the deuteron with enough precision to provide a useful constraint of the weak axial current in the two-nucleon system. This will result in an estimate of the L_{1A} and d^R LECs relevant to pionless EFT and chiral perturbation theory, respectively, and will provide an important check on the current EFT paradigm used to study light nuclei. State of the art EFT calculations produce capture rate predictions with precision commensurate with the expected MuSun experimental uncertainty, and if a significant discrepancy is found between these results it would have serious theoretical implications.

The experiment requires both high statistics and careful attention to constrain all systematic errors in the muon lifetime to the 10 ppm level. The author joined the MuSun collaboration before the second main production run in 2015, and contributed to some hardware upgrades for this and for the final systematics testing run in 2016. R2015 run collected $7 \times 10^9 \mu^-$ events after all analysis cuts, and together with the previous R2014 production run this resulted in a total of $1.3 \times 10^{10} \mu^-$ events, exceeding the statistics target.

The majority of the author's work focused on establishing a complete table of systematic uncertainties and ensuring each of them was constrained to below the target level. This thesis details improvements to the analysis of several key effects, including wall stops, the background slope, and fusion and electron interference in the TPC. The systematic uncertainty goal has now largely been achieved, although a few additional tests are planned to study some remaining discrepancies in the basic consistency checks. These tests should be finished within the next few months, and final unblinded result is expected this year!

BIBLIOGRAPHY

- [1] B. Acharya, A. Ekström, and L. Platter. Effective-field-theory predictions of the muon-deuteron capture rate. *Phys. Rev. C*, 98:065506, 2018.
- [2] P Ackerbauer, D.V Balin, V.M Baturin, G.A Beer, W.H Breunlich, T Case, K Crowe, H Daniel, J Deutsch, J Govaerts, Yu.S Grigoriev, F.J Hartmann, P Kammel, R King, B Lauss, E.M Maev, V.E Markushin, J Marton, M Mühlbauer, C Petitjean, Th Petitjean, G.E Petrov, R Prieels, W Prymas, W Schott, G.G Semenchuk, Yu.V Smirenin, A.A Vorobyov, N.I. Voropaev, and P Wojciechowski. A precision measurement of nuclear muon capture on ^3He . *Physics Letters B*, 417(3):224–232, 1998.
- [3] J Adam Jr, M Tater, E Truhlik, E Epelbaum, R Machleidt, and P Ricci. Calculation of doublet capture rate for muon capture in deuterium within chiral effective field theory. *Physics Letters B*, 709(1-2):93–100, 2012.
- [4] B. Aharmim, S. N. Ahmed, A. E. Anthony, N. Barros, E. W. Beier, A. Bellerive, B. Beltran, M. Bergevin, S. D. Biller, K. Boudjemline, M. G. Boulay, B. Cai, Y. D. Chan, D. Chauhan, M. Chen, B. T. Cleveland, G. A. Cox, X. Dai, H. Deng, J. A. Detwiler, M. DiMarco, P. J. Doe, G. Doucas, P.-L. Drouin, F. A. Duncan, M. Dunford, E. D. Earle, S. R. Elliott, H. C. Evans, G. T. Ewan, J. Farine, H. Fergani, F. Fleurot, R. J. Ford, J. A. Formaggio, N. Gagnon, J. T.M. Goon, K. Graham, E. Guillian, S. Habib, R. L. Hahn, A. L. Hallin, E. D. Hallman, P. J. Harvey, R. Hazama, W. J. Heintzelman, J. Heise, R. L. Helmer, A. Hime, C. Howard, M. Huang, P. Jagam, B. Jamieson, N. A. Jelley, M. Jerkins, K. J. Keeter, J. R. Klein, L. L. Kormos, M. Kos, C. Kraus, C. B. Krauss, A Kruger, T. Kutter, C. C. M. Kyba, R. Lange, J. Law, I. T. Lawson, K. T. Lesko, J. R. Leslie, J. C. Loach, R. MacLellan, S. Majerus, H. B. Mak, J. Maneira, R. Martin, N. McCauley, A. B. McDonald, S. R. McGee, M. L. Miller, B. Monreal, J. Monroe, B. G. Nickel, A. J. Noble, H. M. O’Keeffe, N. S. Oblath, R. W. Ollerhead, G. D. Orebi Gann, S. M. Oser, R. A. Ott, S. J. M. Peeters, A. W. P. Poon, G. Prior, S. D. Reitzner, K. Rielage, B. C. Robertson, R. G. H. Robertson, R. C. Rosten, M. H. Schwendener, J. A. Secrest, S. R. Seibert, O. Simard, J. J. Simpson, P. Skensved, T. J. Sonley, L. C. Stonehill, G. Tešić, N. Tolich, T. Tsui, R. Van Berg, B. A. VanDevender, C. J. Virtue, H. Wan Chan Tseung, D. L. Wark, P. J. S. Watson, J. Wendland, N. West, J. F. Wilkerson, J. R. Wilson, J. M. Wouters, A. Wright, M. Yeh, F. Zhang, and K. Zuber. Combined analysis of all three phases of solar neutrino data from the sudbury neutrino observatory. *Phys. Rev. C*, 88:025501, Aug 2013.

- [5] T. Albahri, A. Anastasi, A. Anisenkov, K. Badgley, S. Baeßler, I. Bailey, V. A. Baranov, E. Barlas-Yucel, T. Barrett, A. Basti, F. Bedeschi, M. Berz, M. Bhattacharya, H. P. Binney, P. Bloom, J. Bono, E. Bottalico, T. Bowcock, G. Cantatore, R. M. Carey, B. C. K. Casey, D. Cauz, R. Chakraborty, S. P. Chang, A. Chapelain, S. Charity, R. Chislett, J. Choi, Z. Chu, T. E. Chupp, S. Corrodi, L. Cotrozzi, J. D. Crnkovic, S. Dabagov, P. T. Debevec, S. Di Falco, P. Di Meo, G. Di Sciascio, R. Di Stefano, A. Driutti, V. N. Duginov, M. Eads, J. Esquivel, M. Farooq, R. Fatemi, C. Ferrari, M. Fertl, A. T. Fienberg, A. Fioretti, D. Flay, E. Frlež, N. S. Froemming, J. Fry, C. Gabbanini, M. D. Galati, S. Ganguly, A. Garcia, J. George, L. K. Gibbons, A. Gioiosa, K. L. Giovanetti, P. Girotti, W. Gohn, T. Gorringer, J. Grange, S. Grant, F. Gray, S. Haciomeroglu, T. Halewood-Leagas, D. Hampai, F. Han, J. Hempstead, A. T. Herrod, D. W. Hertzog, G. Hesketh, A. Hibbert, Z. Hodge, J. L. Holzbauer, K. W. Hong, R. Hong, M. Iacovacci, M. Incagli, 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, Y. I. Kim, B. King, N. Kinnaird, E. Kraegeloh, 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, A. Lusiani, A. L. Lyon, B. MacCoy, R. Madrak, K. Makino, F. Marignetti, S. Mastroianni, J. P. Miller, S. Miozzi, W. M. Morse, J. Mott, A. Nath, H. Nguyen, R. Osofsky, S. Park, G. Pauletta, G. M. Piacentino, R. N. Pilato, K. T. Pitts, B. Plaster, D. Počanić, N. Pohlman, C. C. Polly, J. Price, B. Quinn, N. Raha, S. Ramachandran, E. Ramberg, J. L. Ritchie, B. L. Roberts, D. L. Rubin, L. Santi, C. Schlesier, A. Schreckenberger, Y. K. Semertzidis, D. Shemyakin, M. W. Smith, M. Sorbara, D. Stöckinger, J. Stapleton, C. Stoughton, D. Stratakis, 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, T. Walton, A. Weisskopf, L. Welty-Rieger, P. Winter, A. Wolski, and W. Wu. Measurement of the anomalous precession frequency of the muon in the fermilab muon $g - 2$ experiment. *Phys. Rev. D*, 103:072002, Apr 2021.
- [6] I. Alekseev, Ev. Arkhipov, S. Bondarenko, O. Fedorchenko, V. Ganzha, K. Ivshin, P. Kammel, P. Kravtsov, C. Petitjean, V. Trofimov, A. Vasilyev, T. Vasyanina, A. Vorobyov, and M. Vznuzdaev. Cryogenic distillation facility for isotopic purification of protium and deuterium. *Review of Scientific Instruments*, 86(12):125102, 2015.
- [7] L. W. Alvarez, H. Bradner, F. S. Crawford, J. A. Crawford, P. Falk-Vairant, M. L. Good, J. D. Gow, A. H. Rosenfeld, F. Solmitz, M. L. Stevenson, H. K. Ticho, and R. D. Tripp. Catalysis of nuclear reactions by μ mesons. *Phys. Rev.*, 105:1127–1128, Feb 1957.
- [8] S. Ando, T.-S. Park, K. Kubodera, and F. Myhrer. The μ -d capture rate in effective field theory. *Physics Letters B*, 533(1):25–36, 2002.

- [9] V. A. Andreev, T. I. Banks, R. M. Carey, T. A. Case, S. M. Clayton, K. M. Crowe, J. Deutsch, J. Egger, S. J. Freedman, V. A. Ganzha, T. Gorringer, F. E. Gray, D. W. Hertzog, M. Hildebrandt, P. Kammel, B. Kiburg, S. Knaack, P. A. Kravtsov, A. G. Krivshich, B. Lauss, K. R. Lynch, E. M. Maev, O. E. Maev, F. Mulhauser, C. Petitjean, G. E. Petrov, R. Prieels, G. N. Schapkin, G. G. Semenchuk, M. A. Soroka, V. Tishchenko, A. A. Vasilyev, A. A. Vorobyov, M. E. Vznuzdaev, and P. Winter. Measurement of muon capture on the proton to 1determination of the pseudoscalar coupling g_P . *Phys. Rev. Lett.*, 110:012504, Jan 2013.
- [10] Baldini, A. M., Bao, Y., Baracchini, E., Bemporad, C., Berg, F., Biasotti, M., Boca, G., Cascella, M., Cattaneo, P. W., Cavoto, G., Cei, F., Cerri, C., Chiarello, G., Chiri, C., Corvaglia, A., de Bari, A., De Gerone, M., Doke, T., D’Onofrio, A., Dussoni, S., Egger, J., Fujii, Y., Galli, L., Gatti, F., Grancagnolo, F., Grassi, M., Graziosi, A., Grigoriev, D. N., Haruyama, T., Hildebrandt, M., Hodge, Z., Ieki, K., Ignatov, F., Iwamoto, T., Kaneko, D., Kang, T. I., Kettle, P.-R., Khazin, B. I., Khomutov, N., Korenchenko, A., Kravchuk, N., Lim, G. M. A., Maki, A., Mihara, S., Molzon, W., Mori, Toshinori, Morsani, F., Mtchedilishvili, A., Mzavia, D., Nakaura, S., Nardò, R., Nicolò, D., Nishiguchi, H., Nishimura, M., Ogawa, S., Ootani, W., Orito, S., Panareo, M., Papa, A., Pazzi, R., Pepino, A., Piredda, G., Pizzigoni, G., Popov, A., Raffaelli, F., Renga, F., Ripiccini, E., Ritt, S., Rossella, M., Rutar, G., Sawada, R., Sergiampietri, F., Signorelli, G., Simonetta, M., Tassielli, G. F., Tenchini, F., Uchiyama, Y., Venturini, M., Voena, C., Yamamoto, A., Yoshida, K., You, Z., Yudin, Yu. V., and Zanello, D. Search for the lepton flavour violating decay $\mu^+ \rightarrow e^+\gamma$ with the full dataset of the meg experiment - meg collaboration. *Eur. Phys. J. C*, 76(8):434, 2016.
- [11] D. Balin, V. Ganzha, S. Kozlov, Evgeny Maev, G. Petrov, M. Soroka, G. Schapkin, G. Semenchuk, V. Trofimov, A. Vasiliev, A. Vorobyov, N. Voropaev, Claude Petitjean, B. Gartner, B. Lauss, Johann Marton, Johann Zmeskal, T. Case, K. Crowe, and M. Faifman. High precision study of muon catalyzed fusion in d2 and hd gas. *Physics of Particles and Nuclei*, 42:185–214, 03 2011.
- [12] G. Bardin, J. Duclos, A. Magnon, J. Martino, A. Richter, E. Zavattini, A. Bertin, M. Piccinini, A. Vitale, and D. Measday. A novel measurement of the muon capture rate in liquid hydrogen by the lifetime technique. *Nuclear Physics A*, 352(3):365–378, 1981.
- [13] G. Bardin, J. Duclos, J. Martino, A. Bertin, M. Capponi, M. Piccinini, and A. Vitale. A measurement of the muon capture rate in liquid deuterium by the lifetime technique. *Nuclear Physics A*, 453(4):591–604, 1986.
- [14] M.J. Barnes and G.D. Wait. A 25-kv 75-khz kicker for measurement of muon lifetime. *IEEE Transactions on Plasma Science*, 32(5):1932–1944, 2004.

- [15] V. Bernard, N. Kaiser, and Ulf-G. Meissner. Qcd accurately predicts the induced pseudoscalar coupling constant. *Phys. Rev. D*, 50:6899–6901, Dec 1994.
- [16] A. Bertin, A. Vitale, A. Placci, and E. Zavattini. Muon capture in gaseous deuterium. *Phys. Rev. D*, 8:3774–3793, Dec 1973.
- [17] E. J. Bleser, E. W. Anderson, L. M. Lederman, S. L. Meyer, J. L. Rosen, J. E. Rothberg, and I T. Wang. Muonic molecules in liquid hydrogen. *Phys. Rev.*, 132:2679–2691, Dec 1963.
- [18] J. Bonilla, B. Acharya, and L. Platter. Muon capture on the deuteron in chiral effective field theory. *arXiv:2212.08138*, 2023.
- [19] O. Bunemann, T. E. Cranshaw, and J. A. Harvey. Design of grid ionization chambers. *Canadian Journal of Research*, 27a(5):191–206, 1949.
- [20] Paul Büttiker and Ulf-G. Meißner. Pion–nucleon scattering inside the mandelstam triangle. *Nuclear Physics A*, 668(1):97–112, 2000.
- [21] M. Cargnelli, W.H. Breunlich, H. Fuhrmann, P. Kammel, J. Marton, P. Pawlek, J. Werner, J. Zmeskal, W. Bertl, and C. Petitjean. Measurement of the muon capture rate in gaseous deuterium. In *Proceedings of the XXIII Yamada Conf. on Nuclear Weak Processes and Nuclear Structure*. World Scientific, Osaka, Japan, 1981.
- [22] Carl E. Carlson. The proton radius puzzle. *Progress in Particle and Nuclear Physics*, 82:59–77, 2015.
- [23] Jiunn-Wei Chen, Takashi Inoue, Xiangdong Ji, and Yingchuan Li. Fixing two-nucleon weak-axial coupling $l_{1,A}$ from μ^-d capture. *Phys. Rev. C*, 72:061001, Dec 2005.
- [24] CDF Collaboration^{†‡}, T. Aaltonen, S. Amerio, D. Amidei, A. Anastassov, A. Annovi, J. Antos, G. Apollinari, J. A. Appel, T. Arisawa, A. Artikov, J. Asaadi, W. Ashmanskas, B. Auerbach, A. Aurisano, F. Azfar, W. Badgett, T. Bae, A. Barbaro-Galtieri, V. E. Barnes, B. A. Barnett, P. Barria, P. Bartos, M. Bauce, F. Bedeschi, S. Behari, G. Bellettini, J. Bellinger, D. Benjamin, A. Beretvas, A. Bhatti, K. R. Bland, B. Blumenfeld, A. Bocci, A. Bodek, D. Bortoletto, J. Boudreau, A. Boveia, L. Brigliadori, C. Bromberg, E. Brucken, J. Budagov, H. S. Budd, K. Burkett, G. Busetto, P. Bussey, P. Butti, A. Buzatu, A. Calamba, S. Camarda, M. Campanelli, B. Carls, D. Carlsmith, R. Carosi, S. Carrillo, B. Casal, M. Casarsa, A. Castro, P. Catastini, D. Cauz, V. Cavalieri, A. Cerri, L. Cerrito, Y. C. Chen, M. Chertok, G. Chiarelli, G. Chlachidze, K. Cho, D. Chokheli, A. Clark, C. Clarke, M. E. Convery, J. Conway, M. Corbo, M. Cordelli, C. A. Cox, D. J. Cox, M. Cremonesi, D. Cruz, J. Cuevas, R. Culbertson, N. d’Ascenzo,

M. Datta, P. de Barbaro, L. Demortier, M. Deninno, M. D'Errico, F. Devoto, A. Di Canto, B. Di Ruzza, J. R. Dittmann, S. Donati, M. D'Onofrio, M. Dorigo, A. Driutti, K. Ebina, R. Edgar, A. Elagin, R. Erbacher, S. Errede, B. Esham, S. Farrington, J. P. Fernández Ramos, R. Field, G. Flanagan, R. Forrest, M. Franklin, J. C. Freeman, H. Frisch, Y. Funakoshi, C. Galloni, A. F. Garfinkel, P. Garosi, H. Gerberich, E. Gerchtein, S. Giagu, V. Giakoumopoulou, K. Gibson, C. M. Ginsburg, N. Giokaris, P. Giromini, V. Glagolev, D. Glenzinski, M. Gold, D. Goldin, A. Golossov, G. Gomez, G. Gomez-Ceballos, M. Goncharov, O. González López, I. Gorelov, A. T. Goshaw, K. Goulios, E. Gramellini, C. Grosso-Pilcher, J. Guimaraes da Costa, S. R. Hahn, J. Y. Han, F. Happacher, K. Hara, M. Hare, R. F. Harr, T. Harrington-Taber, K. Hatakeyama, C. Hays, J. Heinrich, M. Herndon, A. Hocker, Z. Hong, W. Hopkins, S. Hou, R. E. Hughes, U. Husemann, M. Hussein, J. Huston, G. Introzzi, M. Iori, A. Ivanov, E. James, D. Jang, B. Jayatilaka, E. J. Jeon, S. Jindariani, M. Jones, K. K. Joo, S. Y. Jun, T. R. Junk, M. Kambeitz, T. Kamon, P. E. Karchin, A. Kasmi, Y. Kato, W. Ketchum, J. Keung, B. Kilminster, D. H. Kim, H. S. Kim, J. E. Kim, M. J. Kim, S. H. Kim, S. B. Kim, Y. J. Kim, Y. K. Kim, N. Kimura, M. Kirby, K. Kondo, D. J. Kong, J. Konigsberg, A. V. Kotwal, M. Kreps, J. Kroll, M. Kruse, T. Kuhr, M. Kurata, A. T. Laasanen, S. Lammel, M. Lancaster, K. Lannon, G. Latino, H. S. Lee, J. S. Lee, S. Leo, S. Leone, J. D. Lewis, A. Limosani, E. Lipeles, A. Lister, Q. Liu, T. Liu, S. Lockwitz, A. Loginov, D. Lucchesi, A. Lucà, J. Lueck, P. Lujan, P. Lukens, G. Lungu, J. Lys, R. Lysak, R. Madrak, P. Maestro, S. Malik, G. Manca, A. Manousakiskatsikakis, L. Marchese, F. Margaroli, P. Marino, K. Matera, M. E. Mattson, A. Mazzacane, P. Mazzanti, R. McNulty, A. Mehta, P. Mehtala, A. Menzione, C. Mesropian, T. Miao, E. Michielin, D. Mietlicki, A. Mitra, H. Miyake, S. Moed, N. Moggi, C. S. Moon, R. Moore, M. J. Morello, A. Mukherjee, Th. Muller, P. Murat, M. Mussini, J. Nachtman, Y. Nagai, J. Naganoma, I. Nakano, A. Napier, J. Nett, T. Nigmanov, L. Nodulman, S. Y. Noh, O. Norriella, L. Oakes, S. H. Oh, Y. D. Oh, T. Okusawa, R. Orava, L. Ortolan, C. Pagliarone, E. Palencia, P. Palni, V. Papadimitriou, W. Parker, G. Pauletta, M. Paulini, C. Paus, T. J. Phillips, G. Piacentino, E. Pianori, J. Pilot, K. Pitts, C. Plager, L. Pondrom, S. Poprocki, K. Potamianos, A. Pranko, F. Prokoshin, F. Ptohos, G. Punzi, I. Redondo Fernández, P. Renton, M. Rescigno, F. Rimondi, L. Ristori, A. Robson, T. Rodriguez, S. Rolli, M. Ronzani, R. Roser, J. L. Rosner, F. Ruffini, A. Ruiz, J. Russ, V. Rusu, W. K. Sakumoto, Y. Sakurai, L. Santi, K. Sato, V. Saveliev, A. Savoy-Navarro, P. Schlabach, E. E. Schmidt, T. Schwarz, L. Scodellaro, F. Scuri, S. Seidel, Y. Seiya, A. Semenov, F. Sforza, S. Z. Shalhout, T. Shears, P. F. Shepard, M. Shimojima, M. Shochet, I. Shreyber-Tecker, A. Simonenko, K. Sliwa, J. R. Smith, F. D. Snider, H. Song, V. Sorin, R. St. Denis, M. Stancari, D. Stentz, J. Strologas, Y. Sudo, A. Sukhanov, I. Suslov, K. Takemasa, Y. Takeuchi, J. Tang, M. Tecchio, P. K. Teng, J. Thom, E. Thomson, V. Thukral, D. Toback, S. Tokar, K. Tollefson, T. Tomura, S. Torre, D. Torretta, P. Totaro, M. Trovato, F. Ukegawa, S. Uozumi, F. Vázquez, G. Velez, K. Vellidis, C. Vernieri, M. Vidal, R. Vilar, J. Vizán, M. Vogel,

- G. Volpi, P. Wagner, R. Wallny, S. M. Wang, D. Waters, W. C. Wester, D. Whiteson, A. B. Wicklund, S. Wilbur, H. H. Williams, J. S. Wilson, P. Wilson, B. L. Winer, P. Wittich, S. Wolbers, H. Wolfmeister, T. Wright, X. Wu, Z. Wu, K. Yamamoto, D. Yamato, T. Yang, U. K. Yang, Y. C. Yang, W.-M. Yao, G. P. Yeh, K. Yi, J. Yoh, K. Yorita, T. Yoshida, G. B. Yu, I. Yu, A. M. Zanetti, Y. Zeng, C. Zhou, and S. Zucchelli. High-precision measurement of the $|i_L w_i|/i_L$ boson mass with the cdf ii detector. *Science*, 376(6589):170–176, 2022.
- [25] Eugene D. Commins and Philip H. Bucksbaum. *Weak Interactions of Leptons and Quarks*. Cambridge University Press, 1983.
- [26] Andrzej Czarnecki, Matthew Dowling, Xavier Garcia i Tormo, William J. Marciano, and Robert Szafron. Michel decay spectrum for a muon bound to a nucleus. *Phys. Rev. D*, 90:093002, Nov 2014.
- [27] Hilla De-Leon, Lucas Platter, and Doron Gazit. Tritium β decay in pionless effective field theory. *Phys. Rev. C*, 100:055502, Nov 2019.
- [28] Arnaud Dupays. Isotopic effects in the muon transfer from $p\mu$ and $d\mu$ to heavier atoms. *Phys. Rev. Lett.*, 93:043401, Jul 2004.
- [29] Nadia Fettes, Ulf-G. Meißner, and Sven Steininger. Pion-nucleon scattering in chiral perturbation theory (i): Isospin-symmetric case. *Nuclear Physics A*, 640(2):199–234, 1998.
- [30] V. Ganzha, K. Ivshin, P. Kammel, P. Kravchenko, P. Kravtsov, C. Petitjean, V. Trofimov, A. Vasilyev, A. Vorobyov, M. Vznuzdaev, and F. Wauters. Measurement of trace impurities in ultra pure hydrogen and deuterium at the parts-per-billion level using gas chromatography. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 880:181–187, 2018.
- [31] V.A. Ganzha, P.A. Kravtsov, O.E. Maev, G.N. Schapkin, G.G. Semenchuk, V.A. Trofimov, A.A. Vasilyev, M.E. Vznuzdaev, S.M. Clayton, P. Kammel, B. Kiburg, M. Hildebrandt, C. Petitjean, T.I. Banks, and B. Lauss. A circulating hydrogen ultra-high purification system for the mucap experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 578(3):485–497, 2007.
- [32] Alex Gnech, Laura Elisa Marcucci, and Michele Viviani. Bayesian analysis of muon capture on deuteron in chiral effective field theory, 2023.

- [33] J. Goldstone. Field Theories with Superconductor Solutions. *Nuovo Cim.*, 19:154–164, 1961.
- [34] Jan Govaerts and J.I Martinez. Nuclear muon capture on the proton and ^3He within the standard model and beyond. *Nuclear Physics A*, 678:110–146, 09 2000.
- [35] I.H. Hashim, H. Ejiri, F. Othman, F. Ibrahim, F. Soberi, N.N.A.M.A. Ghani, T. Shima, A. Sato, and K. Ninomiya. Nuclear isotope production by ordinary muon capture reaction. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 963:163749, 2020.
- [36] Richard J Hill, Peter Kammel, William J Marciano, and Alberto Sirlin. Nucleon axial radius and muonic hydrogen—a new analysis and review. *Reports on Progress in Physics*, 81(9):096301, jul 2018.
- [37] P. Kammel, W. H. Breunlich, M. Cargnelli, H. G. Mahler, J. Zmeskal, W. H. Bertl, and C. Petitjean. First observation of muonic hyperfine effects in pure deuterium. *Phys. Rev. A*, 28:2611–2622, Nov 1983.
- [38] Peter Kammel and Kuniharu Kubodera. Precision muon capture. *Annual Review of Nuclear and Particle Science*, 60(1):327–353, 2010.
- [39] David B. Kaplan, Martin J. Savage, and Mark B. Wise. A new expansion for nucleon-nucleon interactions. *Physics Letters B*, 424(3-4):390–396, apr 1998.
- [40] David B. Kaplan, Martin J. Savage, and Mark B. Wise. Two-nucleon systems from effective field theory. *Nuclear Physics B*, 534(1):329–355, 1998.
- [41] F.G. Kondev, M. Wang, W.J. Huang, S. Naimi, and G. Audi. The nubase2020 evaluation of nuclear physics properties *. *Chinese Physics C*, 45(3):030001, mar 2021.
- [42] CJC Kruger. Constrained cubic spline interpolation for chemical engineering applications.
- [43] Michael F. L’Annunziata. *Handbook of Radioactivity Analysis, Second Edition*. Academic Press, 2003.
- [44] R. Machleidt and D. R. Entem. Chiral effective field theory and nuclear forces. *Phys. Rept.*, 503:1–75, 2011.
- [45] Ruprecht Machleidt. Chiral Symmetry and the Nucleon-Nucleon Interaction. *Symmetry*, 8(4):26, 2016.

- [46] L. E. Marcucci. Muon capture on deuteron and ${}^3\text{He}$: A personal review. *Int. J. Mod. Phys. A*, 27:1230006, 2012.
- [47] L. E. Marcucci, A. Kievsky, S. Rosati, R. Schiavilla, and M. Viviani. Chiral effective field theory predictions for muon capture on deuteron and ${}^3\text{He}$. *Phys. Rev. Lett.*, 108:052502, Jan 2012. [Erratum: *Phys. Rev. Lett.* **121**, (2018) 049901].
- [48] L. E. Marcucci, A. Kievsky, S. Rosati, R. Schiavilla, and M. Viviani. Erratum: Chiral effective field theory predictions for muon capture on deuteron and ${}^3\text{He}$ [*phys. rev. lett.* 108, 052502 (2012)]. *Phys. Rev. Lett.*, 121:049901, Jul 2018.
- [49] L. E. Marcucci, M. Piarulli, M. Viviani, L. Girlanda, A. Kievsky, S. Rosati, and R. Schiavilla. Muon capture on deuteron and ${}^3\text{He}$. *Phys. Rev. C*, 83:014002, 2011.
- [50] D. F. Measday. The nuclear physics of muon capture. *Phys. Rept.*, 354:243–409, 2001.
- [51] M. P. Mendenhall, R. W. Pattie, Y. Bagdasarova, D. B. Berguno, L. J. Broussard, R. Carr, S. Currie, X. Ding, B. W. Filippone, A. Garcí a, P. Geltenbort, K. P. Hickerson, J. Hoagland, A. T. Holley, R. Hong, T. M. Ito, A. Knecht, C.-Y. Liu, J. L. Liu, M. Makela, R. R. Mammei, J. W. Martin, D. Melconian, S. D. Moore, C. L. Morris, A. Pérez Galván, R. Picker, M. L. Pitt, B. Plaster, J. C. Ramsey, R. Rios, A. Saunders, S. J. Seestrom, E. I. Sharapov, W. E. Sondheim, E. Tatar, R. B. Vogelaar, B. VornDick, C. Wrede, A. R. Young, and B. A. Zeck. Precision measurement of the neutron β -decay asymmetry. *Physical Review C*, 87(3), mar 2013.
- [52] A. Mooser, S. Ulmer, K. Blaum, K. Franke, H. Kracke, C. Leiteritz, W. Quint, C.C. Rodegheri, C. Smorra, and J. Walz. Direct high-precision measurement of the magnetic moment of the proton. *Nature*, 509:596–599, 2014.
- [53] D. Mund, B. Märkisch, M. Deissenroth, J. Krempel, M. Schumann, H. Abele, A. Petoukhov, and T. Soldner. Determination of the weak axial vector coupling $\lambda=g_A/g_V$ from a measurement of the β -asymmetry parameter a in neutron beta decay. *Phys. Rev. Lett.*, 110:172502, Apr 2013.
- [54] Michael Murray. *Muon Catalyzed Fusion Effects in the Precision Measurement of Muon Capture on the Deuteron*. PhD thesis, University of Washington, 2017.
- [55] Yoichiro Nambu. Quasi-particles and gauge invariance in the theory of superconductivity. *Phys. Rev.*, 117:648–663, Feb 1960.
- [56] Ray Kreswell Neely. *A First Experimental Limit on the Relative Rates of Muon Capture on Deuterium from the Quartet and Doublet Hyperfine Spin States of the μd Atom*. PhD thesis, University of Kentucky, 2017.

- [57] C. Patrignani. Review of particle physics. *Chinese Physics C*, 40(10):100001, oct 2016.
- [58] Randolph Pohl, Aldo Antognini, François Nez, Fernando D Amaro, François Biraben, João M R Cardoso, Daniel S Covita, Andreas Dax, Satish Dhawan, Luis M P Fernandes, Adolf Giesen, Thomas Graf, Theodor W Hänsch, Paul Indelicato, Lucile Julien, Cheng-Yang Kao, Paul Knowles, Eric-Olivier Le Bigot, Yi-Wei Liu, José A M Lopes, Livia Ludhova, Cristina M B Monteiro, Françoise Mulhauser, Tobias Nebel, Paul Rabinowitz, Joaquim M F Dos Santos, Lukas A Schaller, Karsten Schuhmann, Catherine Schwob, David Taqqu, João F C A Veloso, and Franz Kottmann. The size of the proton. *Nature*, 466:213–216, 2010.
- [59] Randolph Pohl, François Nez, Luis M. P. Fernandes, Fernando D. Amaro, François Biraben, João M. R. Cardoso, Daniel S. Covita, Andreas Dax, Satish Dhawan, Marc Diepold, Adolf Giesen, Andrea L. Gouvea, Thomas Graf, Theodor W. Hänsch, Paul Indelicato, Lucile Julien, Paul Knowles, Franz Kottmann, Eric-Olivier Le Bigot, Yi-Wei Liu, José A. M. Lopes, Livia Ludhova, Cristina M. B. Monteiro, Françoise Mulhauser, Tobias Nebel, Paul Rabinowitz, Joaquim M. F. dos Santos, Lukas A. Schaller, Karsten Schuhmann, Catherine Schwob, David Taqqu, João F. C. A. Veloso, Aldo Antognini, and null null. Laser spectroscopy of muonic deuterium. *Science*, 353(6300):669–673, 2016.
- [60] Bruno Rossi and David B. Hall. Variation of the rate of decay of mesotrons with momentum. *Phys. Rev.*, 59:223–228, Feb 1941.
- [61] R A Ryan, F Wauters, F E Gray, P Kammel, A Nadtochy, D Peterson, T van Wechel, E Gross, M Gubanich, L Kochenda, P Kravtsov, D Orozco, R Osofsky, M H Murray, G E Petrov, J D Phillips, J Stroud, V Trofimov, A Vasilyev, and M Vznuzdaev. Design and operation of a cryogenic charge-integrating preamplifier for the musun experiment. *Journal of Instrumentation*, 9(07):P07029, jul 2014.
- [62] Rachel Ryan. *MuSun: A Precision Measurement of Nuclear Muon Capture in Deuterium with a Cryogenic Time Projection Chamber*. PhD thesis, University of Washington, 2019.
- [63] Martin J. Savage, Phiala E. Shanahan, Brian C. Tiburzi, Michael L. Wagman, Frank Winter, Silas R. Beane, Emmanuel Chang, Zohreh Davoudi, William Detmold, and Kostas Orginos. Proton-proton fusion and tritium β decay from lattice quantum chromodynamics. *Phys. Rev. Lett.*, 119:062002, Aug 2017.
- [64] T. Suzuki, D. F. Measday, and J. P. Roalsvig. Total nuclear capture rates for negative muons. *Phys. Rev. C*, 35:2212–2224, Jun 1987.

- [65] E. Tognelli, S. Degl’Innocenti, L.E. Marcucci, and P.G. Prada Moroni. Astrophysical implications of the proton–proton cross section updates. *Physics Letters B*, 742:189–194, 2015.
- [66] Timo van Ritbergen and Robin G. Stuart. On the precise determination of the Fermi coupling constant from the muon lifetime. *Nucl. Phys. B*, 564:343–390, 2000.
- [67] I-T. Wang. Muon capture by deuterons. *Phys. Rev.*, 139:B1539–B1544, Sep 1965.
- [68] D. M. Webber, V. Tishchenko, Q. Peng, S. Battu, R. M. Carey, D. B. Chitwood, J. Crnkovic, P. T. Debevec, S. Dhamija, W. Earle, A. Gafarov, K. Giovanetti, T. P. Gorringer, F. E. Gray, Z. Hartwig, D. W. Hertzog, B. Johnson, P. Kammel, B. Kiburg, S. Kizilgul, J. Kunkle, B. Lauss, I. Logashenko, K. R. Lynch, R. McNabb, J. P. Miller, F. Mulhauser, C. J. G. Onderwater, J. Phillips, S. Rath, B. L. Roberts, P. Winter, and B. Wolfe. Measurement of the positive muon lifetime and determination of the fermi constant to part-per-million precision. *Phys. Rev. Lett.*, 106:041803, Jan 2011.
- [69] R. L. Workman and Others. Review of Particle Physics. *PTEP*, 2022:083C01, 2022.

Appendix A

ELECTRON TRACKING

Ideally an electron passing through the ePCs should trigger small groups of wires in each layer. Coincidences between the three planes of a given eSC will then constrain the hit location to a small area where all the wires overlap, because of the different winding directions of the layers. Finally, it should be simple to construct an electron track vector from the two ePC hit locations, with the electron time determined by the eSC.

In practice the ePCs have several data quality problems that make accurate reconstruction of the tracks difficult. These problems can cause electrons to fail various track quality cuts, or result in the complete loss of tracks that should exist. If this happens in a time-dependent way it can produce a lifetime shift for the full tracks relative to the simpler eSC-only electron definition. The extremely high precision requirements of the MuSun experiment mean that even a 10^{-5} level change in efficiency across the event window may produce a significant lifetime shift, as seen in chapter 12. The author was responsible for a full rewrite of the electron tracking software from the ground up, focusing on improving the tracking reliability and avoiding mistakenly rejecting good electrons.

A.1 Data Quality

A.1.1 Sparks

One way the idealized tracking system breaks down is in the case of extremely large clusters, which are usually produced by sparking in the ePCs. These sparks can trigger a substantial fraction of the wires in a given plane, rendering the precise position ambiguous. While sparks can occur in isolation, they are often produced by real electron tracks as shown in figure A.1.

Historically there was a major discrepancy between the eSC-only electrons and the full

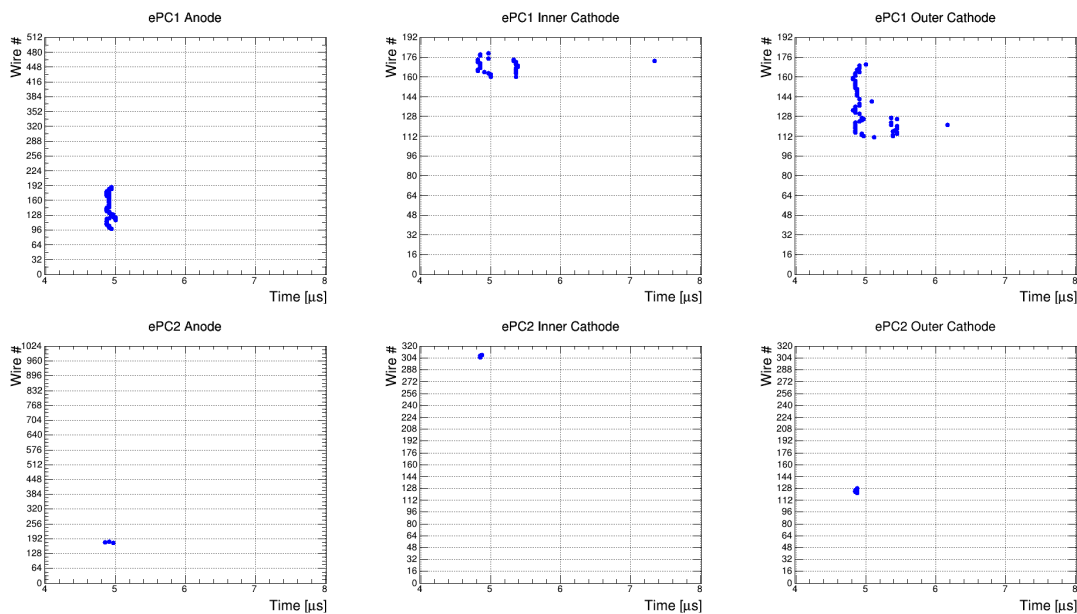


Figure A.1: Example of a spark in ePC1, resulting in a large number of wire triggers as indicated by the blue dots. The coincident triggers in ePC2 form well-localized clusters as one would expect, indicating that this is still a real electron track despite the ambiguity of the ePC1 hit position.

electron tracks, which was eventually traced to a bug in a pre-processing cut meant to eliminate spark events. This cut identified any cluster of 10 or more wires as a spark, and rejected all ePC data in a range of $(-1, +10) \mu\text{s}$ around the cluster for being unreliable. Originally the associated muon entrance should also have been flagged and removed from the final analysis, but this flag was lost in subsequent software updates resulting in the events being retained. It was clear that there was a problem with the tracking, but the deletion of the ePC data caused the spark tracks to appear as eSC hits with no associated track and made them very difficult to identify. The problem was eventually uncovered by the clean track analysis discussed in section 12.2.3, in which both eSC-only and full track electrons were used to distinguish between signal and background. As shown in figure A.2, the spark




Figure A.2: Time difference between full tracks and eSC-only electrons from the clean track analysis. The dip near 0 results from the coincidence veto, while the structure in the center of the plot comes from a $(1.5, 5) \mu\text{s}$ time window applied to the clean track. The sharp steps at -1 and $10 \mu\text{s}$ were unexpected, and are the result of the broken spark cut deleting track data within this window.

cut resulted in an unexpected discontinuity when studying the time difference between the electrons.

Fixing the spark cut was non-trivial. Simply disabling the cut would introduce excessively large clusters that the rest of the tracking system was not designed to handle. Reinstating the spark flag in the muon entrance objects was a more promising route, but applying the spark veto cut resulted in a 70 s^{-1} reduction in the fitted rate even when ignoring the tracks and using the eSC-only electron definition. This was presumably because some of the large clusters were actually the combination of two or more real electron tracks, which would be more likely at early times where there are more decay electrons. As shown in figure A.3, there is no clear size difference between good clusters and sparks or multi-electron clusters. Clearly discarding events based on a somewhat arbitrary size threshold is not ideal, and a




Figure A.3: Distribution of cluster sizes, for the ePC1 anode with 512 total wires. There is no clear distinction between good clusters and spark clusters.

more sophisticated approach is required to resolve this problem.

A.1.2 Afterpulsing

The ePCs have significant afterpulsing, and these afterpulses have a relatively high probability of forming valid coincidences due to the loose 300 ns time window used for ePC matching. It is also possible for afterpulses to interfere with the tracking of a second real electron, making it unclear which groups of wires to use. A 1 μ s artificial deadtime is applied to eliminate the majority of the afterpulsing, but in some cases it can still cause data quality problems.

Feedback

The most prominent issue occurs when the afterpulses are large enough to produce additional afterpulsing, resulting in a feedback loop and a chain of regularly spaced triggers. This may affect a single wire, but more commonly occurs because of electrical noise in an amplifier card and affects the entire 16-wire sector read out by the card. Figure A.4 shows such an event, where there may have been an initial real electron detection but the entire sector quickly




Figure A.4: Example of repeated afterpulsing affecting a sector of 16 wires in the outer cathode of ePC1. Blue dots show individual wire triggers, and several normal clusters are visible in addition to the large afterpulsing structure. The first six points shown in the inset do not align with the repeating pattern, and may be a real event that initiated the afterpulsing.

begins producing synchronous triggers at ~ 400 ns intervals. The feedback loop continues for about $75 \mu\text{s}$, although at later times some of the expected triggers are missing. Presumably in these cases the feedback continues to affect the analog portion of the electrons, but fails to cross the discrimination threshold to produce an output signal. The characteristic repetition frequency varies somewhat between amplifier cards, and the duration of the feedback loop can vary but is often tens or even hundreds of microseconds.

Masking

Typically the feedback loop problem was caused by poor grounding of the amplifier card in question, and would generally be fixed by reseating the card and ensuring all electrical connections were secure. However, the amplifier cards were difficult to physically access, making this process rather cumbersome and delicate. As a result, adjustments to the amplifier cards

were usually performed during scheduled periods of beam downtime and maintenance. Even after maintenance it was not possible to eliminate feedback effects in some cases. Individual noisy wires, known as hot wires, may also be a sign of problems internal to the ePC chamber that cannot be fixed by adjusting the amplifier cards.

Excessive noise triggers from a given wire or sector clearly make it difficult to identify real signals, but the biggest problem was that a large number of triggers could overwhelm the DAQ system. To ensure smooth operation of the rest of the experiment, software masks could be applied to eliminate signals from noisy wires before the data is passed to the DAQ. Often cards that became noisy would be temporarily masked until the next maintenance period, while the particularly bad wires were permanently masked.

Strangely, the masked wires are not completely devoid of hits. These signals appear to follow the expected muon decay time distribution, and most likely produced after the masks are applied as a result of crosstalk from signals in adjacent wires. Presumably similar crosstalk signals also occur for unmasked wires, although this has not been studied in detail.

The specific mask configuration for any given dataset was not well documented, and must be reconstructed from the data. The crosstalk triggers on the masked wires mean we cannot simply select channels with no data, but they still have reduced hit rates compared to unmasked channels as shown in figure A.5. For each data file a smoothed distribution of hit rate vs wire is produced, and any wires falling significantly below this average are flagged. The mask settings were changed infrequently, so an additional temporal smoothing step is applied to prevent a wire from being marked as masked in a single file. Once periods of stable mask settings were identified the masks were saved in the central MySQL database for access during the main analysis.

A.1.3 Bad Sector

Finally, recall from section 4.5.1 that one of the ePC2 anode sectors had a broken wire and needed to be ramped down to a significantly reduced voltage. This resulted in a near zero detection efficiency both for that sector and for the two adjacent sectors, as shown in figure




Figure A.5: Positions of triple coincidences in ePC1 (top) and ePC2 (bottom). Sharp vertical and diagonal stripes are caused by masked wires. The fuzzy vertical band in the center of the ePC2 plot is bad sector with reduced voltage.

A.5. In R2016 an additional wire broke, requiring a second sector to be ramped down and producing another region with minimal sensitivity.

These bands are different from the masking, which only affects a single plane while the other two will trigger normally. None of the three planes will be triggered by electrons in the bad sector, so there is no way to reconstruct tracks in this region. This essentially just reduces the effective solid angle coverage of the detectors when requiring full tracks. It can also couple to asymmetries in the electron emission direction, but fortunately the bad sector is on the side of the cylinder so it does not produce a strong vertical asymmetry effect as described in chapter 14.

A.2 Preprocessing

Before beginning the tracking some preprocessing is performed to clean the data. First, any stretch of 10 or more hits with at least 1 hit per μs on average and at most 2 μs between adjacent hits is identified as a feedback loop. The minimum hit requirement prevents false positives, since it is common for normal electrons to have one or two afterpulses. This test is applied to individual wires, and then nearby wires are grouped together using a simple clustering algorithm based on separate time and space coincidence tests.

The position and time of the beginning of a feedback stretch are saved, and are treated similarly to normal clusters since they may correspond to a real electron detection. They are flagged as being part of a feedback loop so that their effects may be tested at later analysis stages. The stop time is also recorded, although it is not clear if this has any physical significance. The individual hits involved in the feedback loop are then deleted, since any real detection in this period cannot be distinguished from the noise. These stretches are then treated similarly to masked wires for the remainder of the analysis.

The spark cut based on the sizes of individual clusters is no longer used, but there is a second cut based on the total number of hits in each plane. This works by breaking each data block up into 45 μs chunks, and any chunk with more than 200 hits for a given plane is identified as being excessively noisy. This cut should be less sensitive to double-electron

events, with the high threshold being primarily sensitive to sparking or any feedback loops that are not caught properly. There is also no action taken by default, instead muon entrances in noisy chunks are simply flagged for potential further study.

A.3 Clusters

Signals from individual ePC wires must be grouped into clusters associated with a given electron. Originally this used the same clustering algorithm described for the μ PC in section 7.6.3, where time-coincident clusters are created first and then split into groups of spatially adjacent wires. However, with additional noise in the ePCs as well as extended structures created by the sparks and feedback loop effect, this simple algorithm is liable to result in unrelated signals being clustered together.

The new clustering algorithm instead builds up clusters one wire at a time, simultaneously checking both the time and space coincidence windows for each pair of wires. By performing only local pair-wise comparisons this avoids any possibility of separate clusters being mistakenly merged together. There is also a final step where overly large clusters are broken apart into multiple reasonably sized clusters, preventing tracking problems caused by the large clusters without vetoing the events.

A.3.1 Clustering

The algorithm maintains an array of "active" clusters, which are stored as tree structures with nodes representing individual hits. The edges of the tree record the distance between hits in space and time, and the tree is arranged in such a way that the edges always connect the nearest possible neighbors. This structure is maintained when new hits are added to the cluster or when two clusters are merged together, and will be important for the fragmentation step.

The hits for a given ePC plane are sorted by time and processed in order. Each new hit is compared against every active cluster, and is added if it is within 3 wires and 300 ns of any hit already included in the cluster. If a match is found the remaining clusters are still

checked, and if the new hit is in range of any additional clusters they are merged with the first matching cluster. Alternatively, if no matches are found then the new hit becomes the root node of a new cluster. Finally, any clusters that fail the time coincidence check with the new hit will never be matched with any later hits, so they are saved and removed from the active list.

A 1 μs deadtime is applied after each hit, and any afterpulses occurring within this deadtime will extend the deadtime. This can potentially result in quite long deadtime periods, although typically chains of afterpulses lasting more than a few μs will be caught by the feedback loop filtering and removed during preprocessing. When a cluster becomes inactive the same 1 μs deadtime is also applied to all wires inside the cluster, regardless of whether a hit was actually detected on that wire. This prevents an issue where occasionally afterpulses are seen on wires that did not trigger initially, most likely due to crosstalk between channels in the amplifier cards.

Masked wires present a problem for the clustering, because an electron impact near a masked region will often create two separate groups of hits on either side. If these are coincident in time then they should be merged into a single cluster even if they are separated by several masked wires. When calculating the separation between two wires any masked wires in between are weighted by 50%, so as many as 5 masked wires are allowed in the middle of a cluster.

A.3.2 Fragmentation

Clusters more than 10 wires across or 450 ns in length are considered oversized, and likely include either multiple electrons or some combination of sparking and afterpulsing. These large clusters are therefore broken apart into smaller clusters, in a process known as fragmentation. Using the tree representation for the clusters, it is simple to identify the pair of hits in the cluster with the greatest separation and to split the cluster in two by removing that edge. In the case where there are multiple edges of the same length the one closest to the midpoint of the cluster is chosen for the split. The fragmentation repeats recursively,




Figure A.6: Number of ePC1 anode cluster fragments for an R2015 test file. The large majority of electrons produce a single unambiguous cluster, but about 1% produce large clusters that are split into multiple fragments.

until all output clusters are below the maximum size.

After the final clusters are formed the tree representation is no longer required, and a simplified bounding box representation is used for the output. The number of hits included in the cluster is saved, but the positions and times of individual hits are discarded. The number of fragments produced is also saved, and figure A.6 shows the distribution in the number of fragments. About 99% of the original clusters are below the size threshold to begin with, with a tail somewhat similar to that shown in figure A.3 above indicating that large clusters are being split into groups of about 10 wires each as expected.

A.4 Plane Coincidences

With the wires on each plane grouped into clusters, the next step is to find coincident clusters in different planes. Because the three planes have different angles, a coincidence constrains the location of the electron to the point where the clusters overlap. Ideally electrons should produce clusters in all three detector planes that overlap at a single location, and such 3-fold

coincidences may be identified as real electron detections with high confidence.

In practice one cluster is often missing, leaving only a 2-fold coincidence. Most pairs of wires on different planes will overlap at some point in the ePC, so without the additional consistency check from the third plane lining up at the same point these 2-fold coincidences are much more likely to be produced by unrelated noise triggers. The interpretation of these events depends on the status of the third wire plane, and if it is working properly then the lack of a matching detection is a sign that the detection may not be real. However, if the two location of the coincidence overlaps with a masked region in the third plane then the natural assumption is that the third cluster was simply not seen due to the masking. Such events are labeled as "2.5-fold" coincidences, and are considered higher quality than a normal 2-fold.

In the case of multiple simultaneous electron detections there will also be a number of spurious 2-folds where the clusters from different electrons overlap. Typically the real coincidences will be 3-folds or 2.5-folds, and can therefore be distinguished from the extra 2-fold coincidences. It is therefore natural to give less credence to any 2-fold coincidences that re-use clusters from a higher quality coincidence.

Figure A.7 shows a cartoon of the coincidence identification algorithm. First all possible coincidences between clusters are added to a list of preliminary pairs. These are then analyzed in several passes, with each pass identifying coincidences of progressively lower quality:

1. 3-fold coincidences
2. 2.5-folds with only unused clusters
3. 2-folds with only unused clusters
4. 2-folds sharing a cluster with a 2.5-fold
5. 2.5-folds sharing a cluster with a 3-fold
6. 2-folds sharing a cluster with a 3-fold

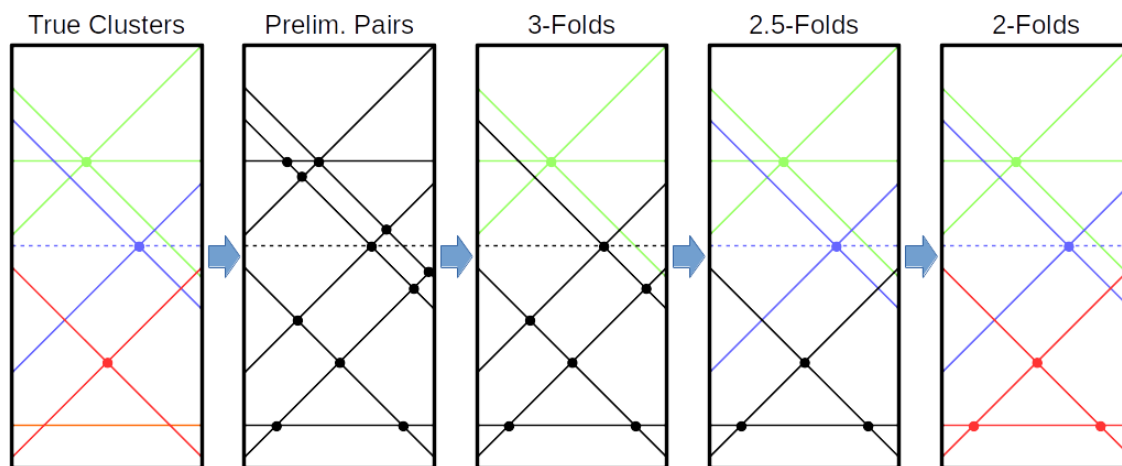


Figure A.7: Cartoon of the ePC plane matching algorithm applied to coincident 3-fold, 2.5-fold, 2-fold, and anode-only electrons, with the dashed line representing the masked wires in the 2.5-fold coincidence. Preliminary pairs are formed from all coincident clusters, which are then sorted into coincidences in order of decreasing quality. The 3-fold and 2.5-fold are both reconstructed correctly, but the algorithm cannot distinguish between the real 2-fold and the extra overlaps with the unused anode cluster. In the case of ambiguity all equally good coincidences are created.

All possible coincidences of a given class are created simultaneously, so for instance it is possible for two 3-folds to be created sharing a single wire. This typically occurs for 2-folds where creating a valid coincidence is relatively easy, while 3-folds and 2.5-folds are unlikely to be aligned in such a way that they share a cluster. Coincidences are only formed if at least one of the clusters in the pair has not yet been used in a previous step.

A 300 ns time window is used to identify coincident pairs of clusters, measured as the time from the end of the first cluster to the beginning of the next. For 3-fold coincidences only two of the pairs must pass the time coincidence check, but they are also characterized by the spatial gap between the wires as indicated in figure A.8. Coincidences are only formed if the gap is at most 1 wire wide, and ideally it should be a negative value indicating that




Figure A.8: Left: Cartoon showing how the misalignment of a 3-fold coincidence is measured using the gap between the clusters. Right: Observed gap distribution, with negative values instead measuring the amount of overlap. Most coincident groups of three clusters have some degree of overlap, as one would expect.

all three clusters overlap with each other.

Another consideration is that the anode wires tend to be more reliable and less noisy than the cathodes, and can also provide ϕ information on their own. Anode clusters are therefore always saved for further analysis, even if they do not form coincidences with clusters in either cathode plane. We also typically ignore 2-folds formed from cathode clusters with no anode, although 2.5-folds with a masked anode can be meaningful.

Finally, in the case where a large cluster has been fragmented there may be many possible matches between the planes. Sparks can produce large clusters in all three planes, resulting in a grid of 3-fold coincidences spanning the affected area as shown in figure A.9. However, if only one or two planes have large clusters then the electron position may still be fairly well localized, since it is constrained by the small clusters in the other planes.




Figure A.9: Diagram of the plane matching behavior for one, two, or three planes with groups of fragmented clusters. All possible 3-fold coincidences are formed, as well as additional 2-folds that are not shown in the diagram.

A.5 Tracks

At MTA the cluster coincidences from both ePCs are matched with eSC electron detections to create full tracks. This process starts by forming track candidates from all possible combinations of hits in the ePCs. A series of track quality cuts are then used to select a single best track associated with each eSC hit.

Hits in the two ePCs are first made into pairs, once again using a 300 ns coincidence window and also requiring a ϕ difference less than 1 radian (57°) to avoid selecting hits on opposite sides of the target. These ePC pairs are compared with the list of eSC electrons, and formed into tracks if the average ePC times is within 250 ns of the eSC time and the projected ϕ of the ePC pair matches either the correct eSC segment or an adjacent segment. Unpaired ePC hits are then checked against the eSC electrons, and are saved as "half tracks" if they pass the 250 ns coincidence window and have a ϕ value within 1 radian of the center of the eSC segment. A final eSC-only track candidate is created for each electron, acting as a fallback case if no better tracks are identified.

A.5.1 *Quality Cuts*

There may be a large number of track candidates associated with a single eSC hit, particularly in the case where there are many cluster fragments in the ePCs. Often many of these are physically implausible for various reasons, for example due to being at an odd angle. To select only high-quality tracks, the following cuts are applied in order:

1. Tracks should include both ePCs.
2. The projected ePC ϕ at the eSC should intersect the correct eSC segment.
3. Tracks should pass within 60 mm of the beam axis (half the TPC's diagonal width).
4. The projected ePC Z at the eSC should be within 150 mm of the estimated eSC Z.
5. 3-fold ePC hits are preferred over those with missing planes.
6. ePC planes within 2 time bins (60 ns) of the eSC time are preferred.
7. Misaligned ePC 3-folds have lower priority than those with at least 1 wire of overlap.

Full tracks are much more useful than half tracks, so the first priority is to select those if possible. The next three cuts ensure that the ePC track aligns with both the eSC detection position and the beam axis, since both decay electrons and beam backgrounds should generally be emitted outwards from the center of the cylinder. The final three cuts select higher-quality ePC clusters, but must be applied after the loose geometrical cuts to avoid the possibility of identifying a very high-quality ePC cluster that is part of a different electron track. The specific cut values are chosen based on the detector sizes and resolutions, with more restrictive cuts not being particularly meaningful.

These cuts are applied by sorting the track candidates based on the parameter in question, and selecting either those that pass the cut. If no tracks pass the cut then the track that came closest is selected, or in the case of a tie multiple tracks may be selected. For cuts 3,

4, and 6 tracks are also selected if they are sufficiently close to the highest quality track, essentially loosening the cut threshold if there are no good matches. This helps ensure that the later filtering steps still have multiple candidates to choose from, and prevents these cuts from having too much control over the choice of the final track. The tracks that pass all steps of this screening process are added to a list of "good tracks".

Best Tracks




Figure A.10: Example of a track display, with shaded regions representing the eSC, ePCs, and TPC. Track candidates are color coded based on how many cuts they passed, with good tracks colored blue and the best track colored green.

The first round of quality cuts is meant to eliminate track candidates that are clearly bad, but often retains a group of several good tracks as shown in figure A.10. This is by design, since part of the goal was to avoid unnecessarily restrictive or arbitrary cuts during the initial screening. Now that the set of good tracks has been identified, tighter cuts may be

applied to select a single "best track" without much fear of a biased cut picking out a track that excels in one parameter while failing in others. For example, historically tracks were selected based on their alignment with the center of the eSC segment but in figure A.10 that would clearly select one of the worst track candidates if applied without the first filtering step.

The best track selection is similar to the initial selection of good tracks, but now there is no threshold to pass the cut and instead only the highest quality tracks are selected. A list of several cuts is still required since there are often ties, and are ordered roughly by increasing arbitrariness:

1. Time alignment between the ePC planes and the eSC.
2. Cluster quality flags from the ePC plane matching analysis.
3. Agreement between the projected ePC Z and the eSC Z.
4. 3-fold ePC clusters with the greatest overlap area.
5. Distance of closet approach between the track and the beam axis.
6. Distance between the projected ePC ϕ and the center of the eSC segment.
7. Distance of closest approach between the track and the center of the TPC.

Cuts 1 and 3 are reasonable things to check, but here they are being compared below the resolution of the detectors. Cuts 2 and 4 are more restrictive versions of the ePC cluster quality checks originally applied to the good tracks, and these were left to the end out of concerns that they could introduce too much dependence on the ePC sparks and other data quality issues. The final 3 geometric cuts are applied to eliminate any remaining ties, with the ϕ cut in particular being entirely arbitrary while the distance to the center of the TPC

is somewhat biased depending on the muon stop position. Usually the eSC Z cut is enough to determine a single best track, and the remaining cuts are only required in very rare cases.

To test how sensitive the tracking is to the specific choice of final cuts, several alternate best track definitions are created that prioritize different parameters. These work by applying an extra quality cut to the list of good tracks, before checking the cuts listed above in the case of any remaining ties. The alternate best track types are:

Default No special sorting step is applied.

Impact Parameter Tracks must pass as close to the reconstructed stop position as possible. This should provide the best track choice for decay electrons, but could introduce some unwanted coupling to TPC track interference. It defaults to use the center of the TPC if there is no muon stop.

Zero Impact Parameter Tracks must pass as close to the center of the TPC as possible. This is similar to the normal impact parameter version, but attempts to remain independent of the TPC tracking details.

Phi Tracks must pass as close to the center of the eSC panel as possible. This version is not particularly meaningful, but was the way tracks were historically selected so it is retained for comparison purposes.

The impact parameter variations are intended for use in conjunction with a tight geometric cut on the associated measurement, and ensure decay electrons are included in the cut if possible. All variations produced fitted decay rates within $\pm 1 s^{-1}$ even with a 120 mm impact parameter cut applied, indicating that the specific choice of cuts applied at this stage has little effect.

Some additional special best tracks are created for specific use cases. Three tracks are created that require anode clusters from either ePC1, ePC2, or both ePCs. In most cases these will all be identical to the standard best track, but occasionally the only available

anode cluster will result in a highly angled track that will be rejected for failing to pass near the TPC. A set of ePC-only best tracks are also created from hits with no matching eSC, but these are hard to interpret and not currently used for anything.

A.5.2 High Multiplicity Events

In most cases the number of initial track candidates is less than 10, and the above procedure performs well despite its many sorting steps. However, the distribution of track candidates has a long tail, and in the worst case where both ePCs had a spark there may be more than 10,000 possible combinations of all the different cluster fragments. These can cause serious performance degradation despite being a very small fraction of events, so in such events additional cuts are applied to the individual ePC hits before they are grouped into pairs:

1. The ePC ϕ should align with the eSC.
2. The ePC Z should roughly align with the eSC Z.
3. 3-fold ePC clusters are preferred.
4. Closer time alignment between the ePC and the eSC is preferred.
5. 3-fold clusters with at least one wire of overlap are preferred.

These cuts are similar to the quality cuts for full tracks, but because only single ePC hits are being considered the cuts are less reliable for cases such as off-axis tracks. For normal events these cuts would risk mistakenly cutting the true best track, but with such high multiplicity there will still be a large number of relatively similar track candidates to choose from even after these cuts. After each cut is applied the product of the number of ePC1 and ePC2 clusters is rechecked, and if it falls below 10,000 we revert to the standard analysis procedure and omit the remaining pre-screening steps.




Figure A.11: Fitted decay rate shifts for events with a single eSC hit, when applying increasingly strict ePC track requirements. Cathode OR refers to ePC hits with an anode and at least one cathode, while cathode AND refers to ePC 3-folds. The points are all highly correlated, so apart from the first point the displayed error bars reflect the uncertainty in the difference from the eSC-only definition rather than the total fit uncertainty. Dashed lines indicate the original fit values, while the solid lines have been corrected for the expected background slope effect.

A.6 Results

Comparing the output of the ePC tracking to the standard eSC-only analysis is rather complicated. There are a large number of possible ePC track variations, ranging from very loose tracks with only a single anode cluster required to strict tracking requiring 3-fold coincidences in both ePCs and a tight impact parameter cut. Recall that the tracking can produce a lifetime shift due to time-dependent cut efficiencies, so stricter tracking with more cuts applied typically produces a larger effect. Figure A.11 shows a comparison of different track definitions, with cathode AND referring to ePC 3-folds while cathode OR refers to ePC

hits with an anode and at least one cathode. The error bars on the track results modified to show the uncertainty in the difference from the eSC-only fit:

$$\Delta R' = \sqrt{(\Delta R_{track})^2 - (\Delta R_{eSC})^2}. \quad (\text{A.1})$$

Even with the tracking improvements discussed in this appendix, there is still a roughly 8 s^{-1} rate shift when requiring any form of ePC information. This shift is surprisingly similar when using any combination of ePC1 and ePC2 anodes and when using ePC1 cathode clusters, and this behavior is not entirely understood. If requiring a single plane already produces a significant shift then the shift should increase when using more planes, but this is not observed. An additional shift is produced by the ePC2 cathodes, which is expected since they tend to have significantly more noise and afterpulsing than the other planes. This is also consistent with an analysis specifically looking at events including the beginning of an ePC feedback loop, which found greatly increased effective decay rates since electrons later in the events were less likely to form tracks. Although the shift caused by the ePC2 cathodes is understood in principle, it is hard to correct for quantitatively since it depends on the specific configuration of ePC wire masks and noise pulses in each event.

Another important consideration when comparing different ePC track definitions is that they impose different implicit impact parameter cuts on the electrons, as shown in figure A.12. The eSC is sensitive to electrons coming from any direction, but the requirement that an electron pass through both the eSC and ePC2 eliminates beam electrons with very shallow angles. Requiring an electron to pass through ePC1 imposes a much tighter angular cut, in which case it makes little difference whether ePC2 is required or not. We also impose artificial impact parameter cuts in some cases, typically the 200 mm or 120 mm cuts shown in the diagram. Any of these impact parameter cuts will tend to prioritize decay electrons coming from the TPC, while significantly reducing beam backgrounds.

While eliminating beam backgrounds is one of the primary benefits of using full tracks, the background slope effect described in chapter 12 means that suppressing the beam background also produces a rate shift relative to the eSC-only tracks. An attempt was made to correct




Figure A.12: Diagram of the implicit geometric cuts imposed when selecting ePC tracks, for background electrons originating upstream on the beam axis. Electrons from the red region always miss both ePCs, while those in the orange region can hit ePC2 but always miss ePC1. Explicit 200 mm or 120 mm impact parameter cuts (dashed lines) further limit the tracks to those from the blue and/or green regions, respectively.

for the background slope in figure A.11, but it is not obvious how to apply the electron multiplicity analysis to full tracks. Instead the slope is assumed to scale with the overall background level, but this may not be a valid assumption depending on the source of the slope. If the tracks have a significantly different background slope this may help explain the rate shift observed for ePC1 and anode-only tracks, but this has not been studied in much detail.

A.7 Conclusion

The redesigned electron tracking system described in this appendix has resolved some long-standing software issues, and reduces the discrepancy between the eSC-only electrons and the full tracks to about the $10 s^{-1}$ level. This is still above the few s^{-1} uncertainty required

for the final MuSun analysis, but it is perfectly sufficient for most other purposes. The improved tracking has been very useful for several systematic tests, such as the clean track analysis used to study the background slope in section 12.2.3 and electron migration tagging as mentioned in chapter 14.

Appendix B

MUON ON REQUEST

Section 10.3.1 briefly described the time-dependent background structure produced by the kicker, and presented a simple background model given by

$$B(t) = \begin{cases} B_k + (sB_0 - B_k)e^{rt} & t < 0 \\ B_0 & 0 \leq t < t_k, \\ B_k & t > t_k \end{cases}, \quad (\text{B.1})$$

where r is the nominal beam rate, t_k is the time of the kicker step, B_k and B_0 are the kicked and un-kicked background levels, and s is 0 for muons seen by the μSC and 1 for unseen muons. The model is sufficient for capturing most of the important behavior relevant to this thesis, but omits most of the fine details. This appendix will discuss some of the more subtle effects caused by the kicker, as well as its effects on the backgrounds seen by the various MuSun detector systems.

B.1 Muon Background

B.1.1 Step Shape

A detail model has been developed for the kicker step itself, consisting of a linear slope that blends into the flat background with an error function onset and an exponential tail. These are constrained to be tangent to the linear portion at the point where they meet, but the positions of those points are free parameters. Figure B.1 shows fits to the beginning and end steps of the kicker window using this model. The main section of the step takes about 40 ns, with the exponential tail remaining significant for another 40 ns afterwards. This is a much faster than the muon decay rate, and has no relation to the background slope effect.




Figure B.1: Fits to the beginning (left) and end (right) step of the kicker window, for μ SC background data in μ CL events. The blue curve shows a model of the step, with linear slope in the center connected to an error function onset and an exponential tail. The red line shows just the linear term, and the green line indicates the center time of the step.

Fitting the kicker step with this detailed model requires high statistics, but it should be fairly consistent so it should not be a problem to use fixed shape parameters when fitting individual datasets. This is not particularly useful, however, as there are usually much larger effects distorting the observed time distributions. For real muon entrances there are large RF oscillations as well as a significant exponential muon decay component, while for secondary particles the details of the kick will be slightly different. For most use cases it is best to use the instantaneous kick approximation and simply avoid fitting near the step itself, with analyses focusing on early times typically excluding a narrow time window around the step while still fitting to times before and after that window.

B.1.2 Long-Range Correlations

Equation B.1 is only valid within the event window from -25 to $+25$ μ s. It can be extended to work up to $+50$ μ s rather trivially, since the delay until the next event follows the same probability distribution as the previous event. The next event also has the same 800 ns delay

before the beam is actually diverted, so after the kicker end step t_{k2} the background becomes

$$B(t) = \begin{cases} B_0 & t_{k2} \leq t \leq t_{k2} + t_k \\ B_k + (B_0 - B_k)e^{-r(t-t_{k2}-t_k)} & t > t_{k2} + t_k \end{cases} . \quad (\text{B.2})$$




Figure B.2: Probability distributions for previous three kick times before the muon entrance, as well as for the sum, using 40 ns binning.

When using a longer time window one must consider more than just the adjacent event. Figure B.2 shows a model of the probability distributions for the previous three kick times, obtained by convolving the original kick time distribution with itself. Adding these together produces an oscillating function that takes about $100 \mu\text{s}$ to settle down to a value of about half the nominal rate. The observed beam background would then be obtained by convolving this kick time distribution with the background distribution seen in the kicker window.

B.1.3 Clock Events

The backgrounds observed in μCL events are normally assumed to be identical to those in events triggered by real muon entrances, and this assumption is fairly accurate within the

kicker window. However, the event type does make a difference outside of the kicker window where the background shape is determined by the spacing between events. For real muon entrances the beam rate parameter should be approximately the sum of the muon entrance rate and the clock trigger rate, assuming the muons and the clock are uncorrelated. For clock events the adjacent clock triggers are always separated by the fixed clock period, and only the real muon entrance rate contributes. This means that outside the kick window the background has a steeper slope for real events than for clock events, and care must be taken when studying this area or using it for normalization.




Figure B.3: Simulated background for real muon events and clock events, with a 10 kHz clock frequency. This high frequency leaves only $100 \mu\text{s}$ between clock triggers, which is too short a time for the correlation with the muon entrance times to decay.

Simply adding the clock rate to the beam rate to obtain an effective combined rate is valid if the clock frequency is low enough that any correlations between the clock and the muon entrance times are smeared out before the next clock trigger occurs. As shown in figure B.3,

if the clock frequency is too high then this assumption breaks down, and a fairly complicated background structure is produced. This starts to be a problem for clock frequencies higher than about 5 kHz, so the 2.5 kHz and 1 kHz rates used in R2014 and R2015 should cause no issues.

B.1.4 Kicker Re-fire Offset

One final subtlety affecting the behavior of the kicker is that, although the kicker cannot be triggered again while it is already active, the delay for in-flight muons clearing the beam pipe is not properly accounted for. As a result the kicker is able to re-fire approximately 400 ns before the beam rate actually recovers at the end of the kicker window. Muon entrances occurring during this period will therefore see the kicked beam rate for up to 400 ns at the beginning of the event, before it briefly returns to the full rate before being kicked again at 800 ns.

This is a very rare occurrence for real muons, since the kicker reduces the beam rate by a factor of roughly 100. Once again the clock events behave differently, with clock triggers being just as likely to occur during this window as at any other time. As a result the clock background differs slightly from the real muon events at very early times, and is only really valid after 400 ns. This is not really a problem since it still accurately captures the normal kicker step, and since the background should be nearly flat before the kick it is also fairly simple to extrapolate backward and remove any artifacts. The effect also allows the previous muon entrance to occur at slightly later times when studying clock events.

B.2 Secondary Particles

In most cases the muon background itself is not particularly important, since the pileup protection cuts should eliminate most events with background muons. Instead the concern is mainly secondary backgrounds emitted by muon decays and other interactions upstream in the beam pipe. These are what produce the majority of the beam backgrounds observed in the electron detectors, as well as the neutron and gamma ray backgrounds seen in the

neutron detectors. The time dependence of these signals roughly matches the unseen muon distribution, but with some key differences.

B.2.1 Step Shape

The separator should remove most of the electrons from the beam regardless of the state of the kicker. In practice the effectiveness of the separator is limited by the issue with the upstream beamline magnet mentioned in section 4.2, but this also means that the kicker should produce relatively little change in the electron background as well. The kicker also has no effect on neutrons or gamma rays produced before that point. As such, the time-dependent component of these background signals is dominated by muons that decay between the kicker and the detectors, and in particular muons stopping in the collimators.

The observed backgrounds are therefore delayed slightly relative to the original beam muon, and the original shape of the kicker step must be convolved with a fast exponential decay curve. The entrance collimator is lead and the upstream collimators are copper, with lifetimes of roughly 75 and 165 ns respectively. Although these are quite fast they are still sufficient to wash out most details of the original step, and so the convolution is omitted in favor of a simpler model where the kick is simply replaced with an exponential decay with rate λ :

$$B_e(t) = \begin{cases} B_0 & 0 \leq t \leq t_k \\ B_k + (B_0 - B_k)e^{-\lambda(t-t_k)} & t > t_k \end{cases} . \quad (\text{B.3})$$

This is less of a concern for the gamma ray background, where the prompt atomic capture peak discussed in chapter 11 dominates over the exponential decay component.

The kick time t_k is also no longer the same as it was for the muon background. Recall that this was the sum of signal travel time and the time of flight of the particles, the latter of which varies significantly for different backgrounds. The beam muons travel at roughly 1/3 the speed of light, but the gammas and electrons travel at nearly the speed of light while neutrons have different speeds depending on their energy. The exact time offset depends on

the position where the secondary particles are created, but for the upstream collimator the electrons arrive about 100 ns before the muons.

B.2.2 Transient Spike

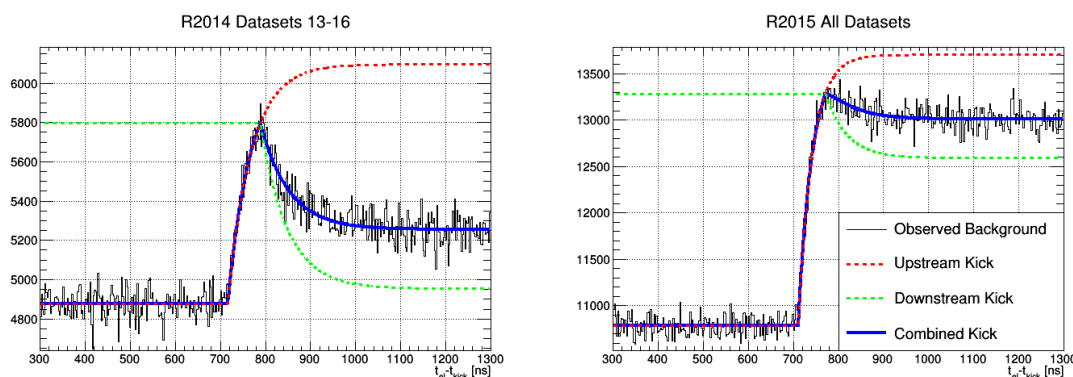


Figure B.4: Fine-binned views of the kicker step as seen by the eSC, for the R2014 and R2015 μ CL data. The signal is fit with two separate kicker steps of opposite polarity, producing a large transient spike for R2014 in particular.

In DSM4 and later the electron background actually increases when kicker fires. This occurs because the diverted beam stopping in the upstream collimator and emitting decay electrons, and the upward step starts close to 700 ns as expected from the time of flight difference. However, there is also a second source of beam electrons from muons stopping in the final entrance collimator. When the beam is diverted this signal is eliminated, resulting in a second downward step closer to the normal 800 ns kick time since the muons travel most of the distance. Adding these two signals together creates a transient spike at 800 ns, as shown in figure B.4.

The spike can be fairly large, but depends on the beam tune since the downward step is the result of off-axis muons that miss the entrance window. For DSM1-3 the upstream collimator was misaligned or entirely removed, so the initial upward step is eliminated and

only a downward step is observed. These early datasets have not been studied in as much detail since the clock was not enabled for most of them.

A transient spike is also seen in the neutron detectors, with the gamma rays arriving first and producing the upward step while the neutrons generate the later downward step. The details are slightly different since the gamma signal is coincident with the original muon, so the spike is less symmetrical with a fast onset and relatively flat top before a similar slow exponential tail. In this case the PSD cuts can help to separate the two signals, but it is difficult to entirely remove the gamma backgrounds from the neutron signal so a small peak is unavoidable.

B.3 TPC Tracks




Figure B.5: Drift time distribution for muon tracks in the TPC. Times between 1 and 16 μs are physically plausible, while the tails are the result of Y tracker errors.

Finally, the background seen in the TPC is a special case because the muon can stop anywhere in the TPC and will only be detected after an unknown drift time. The drift time distribution is easily obtained by looking at tracks associated with real muon entrances, and

is shown in figure B.5. The time distribution of pileup muon tracks in the TPC is then produced by convolving the standard background shape with the drift time. Figure B.6 shows the results of this convolution applied to each combination of seen or unseen muons and normal or kicked beam, where seen and unseen still refer specifically to detection by the μ PC.




Figure B.6: Background time distributions after convolution with the TPC drift time distribution, split into seen or unseen muons and normal or kicked beam states. The time distribution of tracks from the previous muon entrance is also shown for comparison. The normalization of each signal is arbitrary.

The normalization of the curves in figure B.6 is arbitrary, since each signal has an unknown detection efficiency. The time distribution of the seen muons is very similar to the original drift time distribution, since the seen muons are approximated by a short pulse near the entrance time. The unseen muons also have a similar dip on the right after the maximum drift time, but are relatively flat on the left side of the plot since the unseen muons can occur before the entrance time.

These distributions are used for the pileup protection efficiency analysis mentioned in section 10.3.2, since the entrance detectors and the TPC are both sensitive to background

muons and provide an independent check of the detection efficiency. This procedure works by performing a fit to the μCL data, keeping the shape of each component fixed but allowing their amplitudes to vary. The analysis starts with no pileup protection applied to calibrate the relative sizes of each signal, as well as μSC -only pileup protection that eliminates the seen muon background but has no effect on the unseen background by definition. With the detection efficiencies calibrated the fits may then be repeated with increasingly strict pileup protection cuts, reducing the background by 5 orders of magnitude for R2015.

There is some concern that the muons seen in the TPC may not be entirely representative of the entire muon beam. About 1/3 of all muon entrances have no associated TPC track, presumably scattering in the entrance window or stopping in the flange of the deuterium vessel. Muons that are significantly off axis may be both less likely to be detected by the entrance detectors and more likely to stop before reaching the TPC, invalidating the assumption of independent detection efficiencies. Fortunately the final collimator in the entrance detector stack should block most such muons and the μPC distributions typically show that the beam is well localized inside the detector area, so this should not be a serious problem.

B.4 Conclusion

The kicker-induced background time dependence is well understood, and a very detailed model has been developed to explain all of the observed effects. Most of the complications added to the basic background model are only relevant outside of the kicker window or close to the actual kicker step itself. These details are important for various systematics studies looking at the background and pileup, but have very little effect on the standard lifetime analysis which uses a 40 ns time binning and starts the fits at 1 μs . The only way the kicker could affect the final result would be due to the exponential tail following the kicker step for beam electrons, which has largely disappeared before the fit start time and is estimated to cause less than a 0.25 s^{-1} rate shift.

Appendix C

COMPUTER SYSTEMS

The MuSun software framework consists of several analysis stages, discussed in chapters 7 and 8, with many separate analysis modules included in each stage. Running this analysis on the large number of data files produced by the experiment requires substantial data storage and computational resources. The collaboration makes use of the Ranch data archive and the Stampede2 supercomputing cluster, both of which are operated by the Texas Advanced Computing Center (TACC). Some smaller scale analyses have also been adapted to run on a local computer cluster at CENPA. This appendix provides a detailed explanation of how MuSun utilizes these computer resources.

C.0.1 Data Storage

Data is stored on the Ranch long-term data archive system. This system is based on the Quantum StorNext file system, and combines a 20 petabyte (PB, 10^{15} bytes) front-end disk with long-term tape storage for 50 PB of total capacity. Data files are initially stored on the disk and remain there while in active use, but are automatically transferred to the tape if they have not been accessed recently. However, Ranch is not meant to be a redundant system, which is why an additional copy of the raw MuSun data is backed up in the PSI archive.

The MuSun experiment currently has a 700 terabyte (TB, 10^{12} bytes) storage allocation on Ranch. Storage of the raw data files consumes 180 TB of this, with 90 TB of production data, 20 TB of data from R2016, and another 70 TB from previous runs. The event tree files require about 50% more space than the raw data, since individual files are roughly 60% larger but some files which fail quality checks are removed. The MTA output takes about

half the space of the original data files, including both the lifetime trees and the various skim files. In total, a single full analysis pass of the production data generates approximately 180 TB of output files, and with additional outputs from analysis of the other runs and previous passes of the production data these make up the bulk of the allocation. Finally, Monte Carlo simulation data is also stored on Ranch although in this case the intermediate data files produced by each analysis stage are not always retained.

Ranch is optimized for retrieval of very large files, in the range of 300 GB to 4 TB each, and retrieval of smaller files may be over an order of magnitude slower. This range is much higher than the typical MuSun file size of 1.5 to 3 GB, so the UNIX tar (tape archive) command is used to bundle many individual files into .tar files of approximately 500 GB each. Typically a tar file would be compressed, similar to the common zip format, but in this case it is unnecessary as both the raw MIDAS files and the output ROOT files already store data in a compressed format. After these tar files are created they are transferred to Ranch using the Globus file transfer service.

If the data is required for later analysis passes then it must be transferred off of Ranch, again using Globus, and the individual data files must be extracted from the combined tar files. The process of retrieving data from the archive and preparing the files for use in an analysis pass is known as staging. To make the staging process easier the output files from an analysis pass are always grouped by type and by dataset, so each tar file contains a logical grouping of data. However, it is impractical to extract individual files from the archive, so a selection of data files from throughout production is normally kept in local storage for testing purposes.

C.0.2 Job Submission

The analysis itself occurs on the Stampede2 supercomputer cluster. Stampede2 consists of several groups of compute nodes with different architectures, MuSun specifically utilizes the 1,736 "SKX" nodes based on Intel Skylake. These nodes use a pair of 24-core Intel Xeon Platinum 8160 processors for a total of 48 cores per node, or 96 hardware threads with

hyper-threading enabled. They also have 192 GB of DDR4 memory and 144 GB of local SSD storage.

MuSun does not implement any complicated multi-threading scheme, instead opting to simply run many separate instances of the analysis simultaneously on different input files. With nearly 50,000 individual production data files, the MuSun analysis is a so-called embarrassingly parallel workload which scales nearly perfectly with increased processing power. The exception to this is the limited memory capacity available, which can rapidly fill up with many instances of the analysis each creating duplicate copies of every output histogram. With 4 GB of memory per core this is not much of a concern for the SKX nodes, but in the past it was challenging to optimize the memory usage for more constrained systems. Hyper-threading was also found not to provide much benefit for the MuSun software, so normally one copy of the analysis is run per core giving 48 instances per node.

To access Stampede2 one must first connect to one of the login nodes using SSH (the Secure SHell protocol). These nodes are shared between all users and are not meant for significant data processing. Instead, they act as a prep area where users may compile code, perform file management, and submit jobs to run on the main compute nodes. Jobs running on the compute nodes will have access to any files or programs prepared on the login nodes.

Stampede2 jobs are controlled by the Slurm workload manager, and TACC also provides a "launcher" utility for submitting jobs including multiple serial applications. The MuSun team has developed a set of job submission scripts which interface with these systems and automate the process of submitting common analysis jobs. These submission scripts use a local SQLite database to maintain a record of the analysis progress of each data file. Analysis passes are each assigned a number, and the database can then select all input files remaining to be processed in that pass and also select the output files from the pass to feed into the next analysis stage. There are also utilities for tasks such as merging output histogram files together.

Stampede2 users are granted allocations in the form of service units (SUs). The exact value of an SU varies depending on the node type being used, but for the SKX nodes it

is simply equivalent to one hour of processing on a single node. Analyzing one TB of raw data, MU uses about 20 SUs, MTA uses about 12, and the stage 3 scripts vary but use less than half an SU. For a full analysis pass of the production data, this becomes 2000 SUs for MU, 1000 SUs for MTA, and roughly 30 SUs for filling histograms. This demonstrates the importance of the staged analysis framework, for the price of a single full pass it is possible to run 3 passes of MTA or to generate roughly 100 different sets of histograms.

C.0.3 CENPA Cluster

In addition to the resources provided by TACC, MuSun also has access to a local computer cluster at CENPA. This cluster was substantially upgraded in 2018, and now consists of a total of 660 cores with 3 or 4 GB of ram per core depending on the node. That is a similar configuration to the Stampede2 SKX nodes, and while the cluster is much smaller there are also significantly fewer users competing for computer time. Although the processor speed is not exactly the same it would theoretically be possible to finish a full analysis pass of the production data with less than two weeks continuous running on the CENPA cluster.

Unfortunately, the CENPA cluster has major limitations with regard to storage capacity and data transfer throughput. There is about 50 TB of local MuSun hard-drive storage capacity, much less than the 270 TB that would be needed to store all the data files for a full production pass. Even if the storage was available the CENPA network has insufficient bandwidth to handle these large data files, significantly limiting the effective analysis speeds achievable. Stampede2 will therefore continue to be used for large-scale analysis passes, although the CENPA cluster is ideal for other applications such as processing skim files. The Monte Carlo simulation has also been adapted to run on the CENPA cluster, avoiding these storage limitations by keeping intermediate data files in temporary storage on the compute nodes and only actually saving the final lifetime trees.

Appendix D

ESC TIME ALIGNMENT

Section 7.4 describes the clock alignment issue for the eSC WFDs. The alignment procedure is somewhat complicated because the eSCs do not have a sync pulse that can be used as a timing reference. This appendix details the source of the issue and the correction procedure used to ensure good alignment between the eSC channels.

D.1 Problem




Figure D.1: Simplified diagram of an eSC segment, showing the PMT naming convention.

When the block start signal arrives, depending on the phase of the WFD's internal clock it may either latch immediately or on the following clock cycle. The WFD boards record groups of 4 samples at once, so a delay of one cycle of the internal clock results in a 4 tick (≈ 9 ns) offset in the output signal. Because the latching behavior depends on the details of how

the block start signal aligns with each module's clock, it is somewhat random whether each channel experiences a clock slip for any given block. The resulting four clock tick ambiguity (4cta) must be corrected on a block-by-block basis to ensure proper clock alignment.

Each eSC segment is composed of two layers of scintillator each read out by a PMTs on the upstream and downstream ends, as shown in figure D.1. Periods with consistent offsets are identified using trend plots, and these are summed to create time offset histograms such as the ones shown in figure D.2. Because these histograms are created before the clock ambiguity correction, they appear to have a double peak for most tubes. The inner upstream (IU) and outer upstream (OU) times are very tightly correlated, as are the inner downstream (ID) and outer downstream (OD) times. The time difference between upstream and downstream tubes has a spread of about 7 ns, primarily due to the time of flight delay for photons from the electron impact to reach the PMTs. Sub-nanosecond precision is required to extract the eSC Z position with any degree of accuracy, so it is important to get these offsets correct.

D.2 Intra-Segment Alignment

Alignment of the four PMT signals associated with a given eSC segment is achieved using a set of fixed relative offsets. Offset values are extracted by fitting the time differences histograms in figure D.2, using a function consisting of a pair of Gaussian peaks constrained to have identical widths and a separation of exactly 4 clock ticks. ID and OU are compared directly to IU and the time of the first peak is used. OD is instead compared to ID to take advantage of the sharper peak, and the second peak is used to avoid double counting a shift in ID. The global time offset is also checked by matching the IU tubes to the μ SC signal.

The fixed offsets should align the tubes, but the time difference can also be approximately ± 9 ns due to the 4cta effect. The clock ambiguity must be resolved separately for each block, by matching the OU, ID, and OD times against IU within a 50 ns time window. If the time difference is between 5.5 and 12 ns then an offset of 4 clock ticks is applied to that tube, otherwise there is no correction. By using these narrow time windows and a fixed shift, the clock ambiguity is corrected without otherwise interfering with the predefined calibrations.

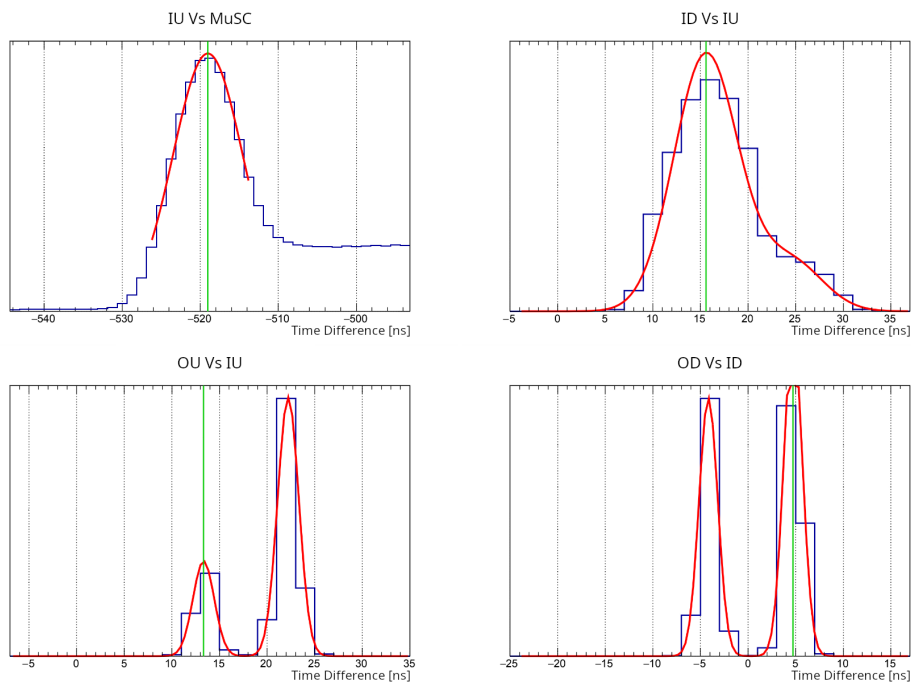


Figure D.2: Time offset calibration fits for eSC segment 3, files 93497-94765. Blue lines are input data, red curves are fits, and vertical green lines indicate the estimated time offsets. Comparing inner (I) to outer (O) detectors produces a sharp peak for both the upstream (U) and downstream (D) detector pairs, although the 4cta effect causes the peak to be doubled. When comparing upstream and downstream detectors on the same scintillator panel the time of flight delay becomes relevant leading to a broader peak and less separation from the secondary 4cta peak.

D.3 Inter-Segment Alignment

The original analog eSC signals were split and recorded by a set of TDCs in addition to the WFDs. Unfortunately the TDC times suffer from some nonlinearity and jitter, so they are not ideal as a reference. However, this is only really an issue for the alignment of signals from the same eSC segment, where improved timing resolution directly leads to a better Z position

estimate. For comparisons between different eSC segments the precise time alignment is less critical, as events merely need to satisfy coincidence windows on the order of 10 ns between eSC segments and 100 ns between the eSC and ePCs.




Figure D.3: Correlations between eSC WFD and TDC times for a single block. Tubes follow the order IU, ID, OU, and OD for each segment.

Global alignment is therefore achieved by applying a second correction step where the WFD signals are matched to the TDCs. This also uses a 50 ns coincidence window, but the average time offset is computed for all four PMTs in each segment. This offset is subtracted to align the times to the TDCs, without changing the relative offsets set previously. Figure D.3 shows the original time offsets between the WFD and TDC signals for each tube, and the offsets after applying the two correction steps. Note that after correction there is still some spread in the time differences, which is a result of the TDC nonlinearity and varies from block to block.

Appendix E

UNIQUE PULSER

As mentioned in section 6.3, a special analysis was included during the R2015 run to study a failure mode of the TPC WFDs identified in previous runs. This error causes the WFDs to start outputting corrupted data, most easily identified by a large negative value for the timestamp. The WFDs send their data as 96-byte packets consisting of 4 zero bytes, a 32-bit timestamp, and 88 8-bit samples. In corrupted data packets the timestamp appears shifted one byte to the right, with the least significant byte missing and the most significant filled with all 1s. For example, a packet with a time of 1,000,000 should appear as 0x000f4240 in hexadecimal, but this becomes 0xff000f42. These errors are somewhat rare and do not occur immediately, but when an error does occur the timestamps remain shifted for the remainder of the data file. Often the shift increases to two or more bytes later on, likely as a result of two separate shifts adding together.

These corrupted timestamps are easily identified and flagged, but the discovery of this error raised concerns that there might be other failure modes that were going undetected. In particular a shift to the left would instead result in timestamps 256 times larger than expected, but which could pass as legitimate values particularly if the original timestamp was small. This can be checked by looking for unusually late timestamps and does not seem to occur, but there could also be more subtle errors that would be less obvious.

To ensure that no unexpected errors were occurring, a "unique pulser" was introduced in R2015 that injected a distinctive triple-pulse signal into the TPC amplifier chain. This could then be identified by a dedicated analysis module and checked against the known pulse time 100 ms after the start of a block. Figure E.1 shows the unique pulser waveform and the times of the identified pulses, and the analysis found no unexpected issues with the pulse




Figure E.1: Unique waveform pulser (left) and times of identified unique pulses within a block. There are a few false positive detections, but otherwise the signal appears to align with the expected pulse time.

time reconstruction.

Another potentially interesting question is whether any usable data can be recovered from the corrupted packets. As shown in figure E.2 it appears that the packets do contain valid waveform data but it is shifted by the same number of bytes as the timestamp, with the missing data in this case coming from the end of the previous waveform. It seems likely that some form of electronics noise is interfering with the synchronization of the WFD internal clock domains, so the data is still present but the readout is misaligned within the packets. The waveforms thus appear to be recoverable, but the missing portion of the timestamp is entirely lost. Without accurate timestamps it becomes difficult to merge contiguous packets into longer waveforms or to properly locate them within the event, and it was determined that correcting these errors was not worthwhile. Instead events with WFD errors are typically discarded, representing a statistics loss of roughly 5 to 10 percent depending on the dataset.




Figure E.2: Selected waveforms from corrupted TPC WFD data packets. The waveforms in the second set are immediately after the corresponding waveform in the first set.

Appendix F

NEUTRON PSD CUT CALIBRATION

Section 8.3 discussed the interpolation procedure for neutron detector waveforms and the application of PSD cuts to distinguish neutron pulses from gamma rays or charged particles, but did not explain how these PSD cuts are generated. Ideally there would be a clear separation between the two signals in the PSD plots, and a cut could be drawn through the gap between them. In practice there is significant overlap between the two bands, particularly for R2015 where the digitizers operated with a reduced sampling rate for the majority of the run. This appendix will describe a more sophisticated procedure for generating cuts, and also discuss the R2015 digitizer problems in a bit more detail.

F.1 Cut Generation

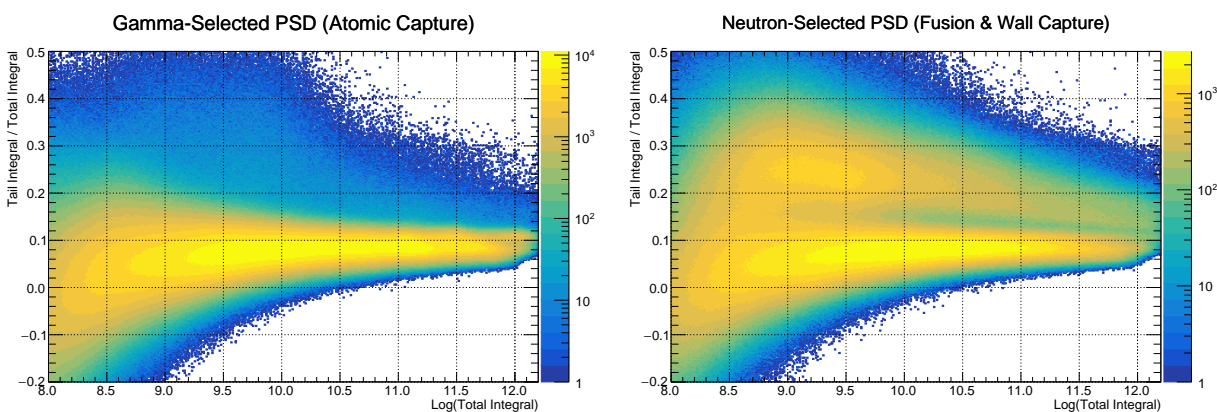


Figure F.1: Left: PSD plot selecting wall stops and using a prompt time cut to isolate the atomic capture peak. Right: Sum of PSD plots selecting ^3He fusion signals in the TPC and wall stops with a prompt veto to eliminate the atomic capture peak.

The various sources of neutron and gamma ray signals are well understood, and are described in chapter 11. Using this information, we can make cuts to select particular signals and change the numbers of neutrons and gamma rays in the PSD plots. A very pure gamma sample may be obtained by selecting wall stop muons with a coincident neutron detection to pick out the prompt atomic capture gammas. A neutron-enhanced sample is obtained by instead selecting the slightly delayed nuclear captures after a wall stop, and adding the signal from ^3He fusions detected in the TPC which gives a relatively pure sample of low-energy neutrons. Figure F.1 shows a comparison of the PSD plots for these two samples. It is also easy to select charged particle signals by looking for coincidences with the eSC segment in front of the neutron detector, these appear similar to the gamma rays but are slightly shifted towards higher tail fractions.

The next step is to compare the neutron and gamma-enhanced PSD plots and estimate the shapes that would be produced by pure neutron or gamma samples. This is essentially done by subtracting one plot from the other to cancel one of the bands, but because the energy spectrum is different for the two samples this must be performed individually for each vertical slice. If F_1 and F_2 refer to the gamma and neutron selected events respectively, we can write them in terms of the true distributions F_γ and F_n :

$$F_1 = N_{1\gamma}F_\gamma + N_{1n}F_n \quad (\text{F.1})$$

$$F_2 = N_{2\gamma}F_\gamma + N_{2n}F_n \quad (\text{F.2})$$

where the N parameters are the numbers of each kind of event in either set. The true distributions are then extracted by scaling and subtracting the two histograms:

$$F_\gamma = (F_1N_{2n} - F_2N_{1n})/(N_{1\gamma}N_{2n} - N_{2\gamma}N_{1n}) \quad (\text{F.3})$$

$$F_n = (F_2N_{1\gamma} - F_1N_{2\gamma})/(N_{1\gamma}N_{2n} - N_{2\gamma}N_{1n}) \quad (\text{F.4})$$

To estimate the N values the distributions are each fit with a pair of Gaussian peaks. This initial estimate is not particularly accurate, and often leaves a residual peak or over-subtracts and creates a dip. The results are refined by adjusting the estimated N values to minimize

any bumps or dips remaining after subtraction, leaving only a smooth tail from the selected signal. The gamma sample is already fairly pure to start with, so the subtraction makes little difference but it is difficult to determine the correct size of the neutron signal. The neutron-selected sample still has more gamma rays than actual neutrons, so it is relatively easy to estimate the size of the signal but subtracting two large peaks still leaves large fluctuations and statistical error bars. To remove these subtraction artifacts, the tails of the peaks are removed and the distributions are smoothed before they are analyzed further.

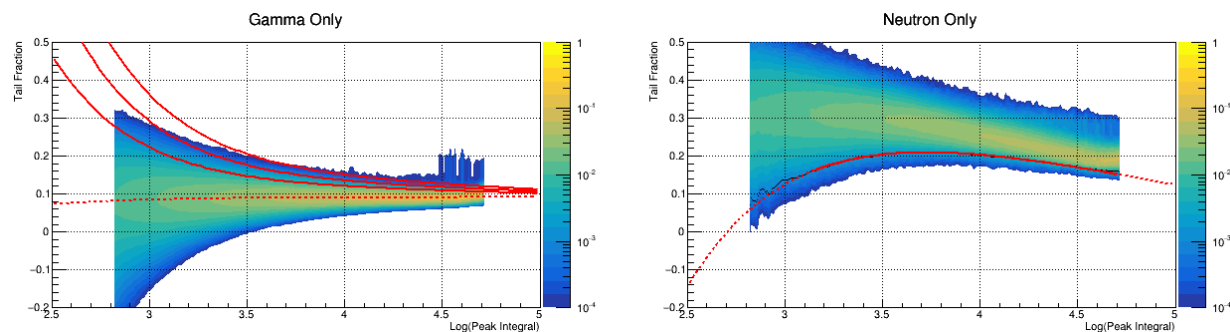


Figure F.2: Estimated PSD plots for pure gamma and neutron samples. The dashed curve on the left shows the peak of the distribution, with the red lines showing cuts at approximately 2, 3, and 4 sigma separations from the peak. The cut on the right is designed to select 99% of neutron signals, with no regard to gammas.

Figure F.2 shows the final PSD plots after subtraction, with each vertical slice now independently normalized. There are some artifacts, but the procedure appears to do a fairly good job of isolating the relevant bands. Because the relative contribution from these two bands can change drastically as shown in figure F.1, the selection efficiency and false positive rate of a given PSD cut is also inconsistent. We therefore define a series of increasingly strict cuts and give each pulse a quality rating based on the strictest cut it passed, with quality 0 signals failing to pass any cut. For a given analysis a higher or lower cut may then be used depending on how many neutrons and gammas are naturally selected by the other analysis cuts.

The majority of the PSD cuts are defined based on the gamma band, with the goal of eliminating a certain fraction of the gamma signals. Some attempts were made to derive these cuts directly using integrals of the estimated pure gamma distribution, but the subtraction artifacts at the top edge of the band made the strict cuts unstable using this method. Instead, for each vertical slice the peak is fit to a Gaussian distribution, and cuts with quality $q > 2$ are simply defined as being q sigma above the peak. This does not really capture the details of the shape, but produces much more consistent and reliable results. The exact quality numbers are rather arbitrary in any case, and the true cut efficiencies are instead determined by fitting the time distributions of the resulting signals as described in chapter 11.

The quality 1 cut is special, being designed to select 99% of the neutron band and ignoring the gamma band entirely. It is defined by taking the integrals of each vertical slice, and then fitting the cut values for each slice to a smooth curve. Ideally this cut would always be above the higher-quality cuts, as this shows the signals are well separated and means that any pulse that passes the first cut will also pass all other cuts and be labeled as a neutron with high confidence. Typically there is poor separation between the signals at low energy and better separation at high energy, but the position where the cuts cross each other and the maximum degree of separation vary between detectors.

F.2 Datasets

There are a few distinct periods that require different PSD cuts. The Bicron detectors used in R2014 typically had much better PSD separation than the upgraded DEMON detectors from R2015. Also, during the first R2015 maintenance period after DSM6 the neutron FADC board was replaced, with the new one misconfigured to run only 85 MHz rather than the intended 170 MHz. Furthermore, for some portions of R2015 the FADC lost the external clock signal entirely due to a faulty cable with poor grounding. As a result each channel defaulted to its own internal clock, which were not synchronized and operated much slower than intended with sampling rates between 33 and 39 MHz. These periods may be seen in the trend plot shown in figure F.3, as they cause the separation between the bands to drop




Figure F.3: Neutron PSD separation trend for detector NU6 in R2015. The figure of merit is defined as the peak separation divided by the sum of the peak widths, for a vertical slice through the center of the PSD plots. Aside from the μ^+ data, the other dips are caused by loss of the external clock signal in the digitizers.

significantly.

Note that while the figure of merit in the trend plot appears to increase after the FADC board was replaced, this is an early diagnostic plot created before the development of the improved waveform interpolation procedures described in section 8.3. This increase is therefore an artifact of the inaccurate pulse reconstruction, and the early data is in fact significantly better than the later data. The reasonable average values allowed the sampling rate issue to go undetected until the end of the run, although the large amount of scatter after the board change is an indication of the unreliability of the pulse processing. With the new pulse interpolation the majority of the data is still usable, albeit with poor PSD separation, but the datasets without the external clock signal are completely unusable.

Appendix G

TPC TRACKING ERRORS

Section 9.5 studied the performance of the TPC tracking algorithms using Monte Carlo truth information, focusing primarily on the effects of p+t fusion interference and the differences between fusion and non-fusion events. This appendix will describe the tracking errors unrelated to p+t fusions in some more detail, and explain the behavior of the upstream tracker which was only mentioned briefly in the main text. Finally, the effects of the TPC pulse-finding threshold on the P-Dir tracker will be discussed.

G.1 Non-Fusion Migrations

Figure G.1 shows the migration probabilities for events without p+t fusions when using the various trackers, relative to Monte Carlo truth. Recall that the goal is for tracks with interference effects to match the behavior of pure muon tracks, so if the true stop position is within the first 1% of a pad (0.15 mm) it is counted as part of the previous row to mimic the effect of the pulse-finding threshold. Without this correction all of the plots except the upstream tracker would show a large spike in upstream migrations for stops near each pad boundary. Even after the correction there is still a prominent short-ranged migration spike, but now it is roughly balanced with similar migrations in both directions.

The basic tracker is the easiest to interpret, since it always uses the most downstream pulse as the stop position. The observed upstream migration is therefore entirely due to the pulse-finding threshold effect, and evidently it is possible for a muon to stop relatively deep in the row without exceeding the threshold although the probability is very low. The fusion veto only removes p+t fusions, so ^3He fusions presumably contribute somewhat to the short-ranged migration peak, and the pileup electron migration effect discussed in chapter

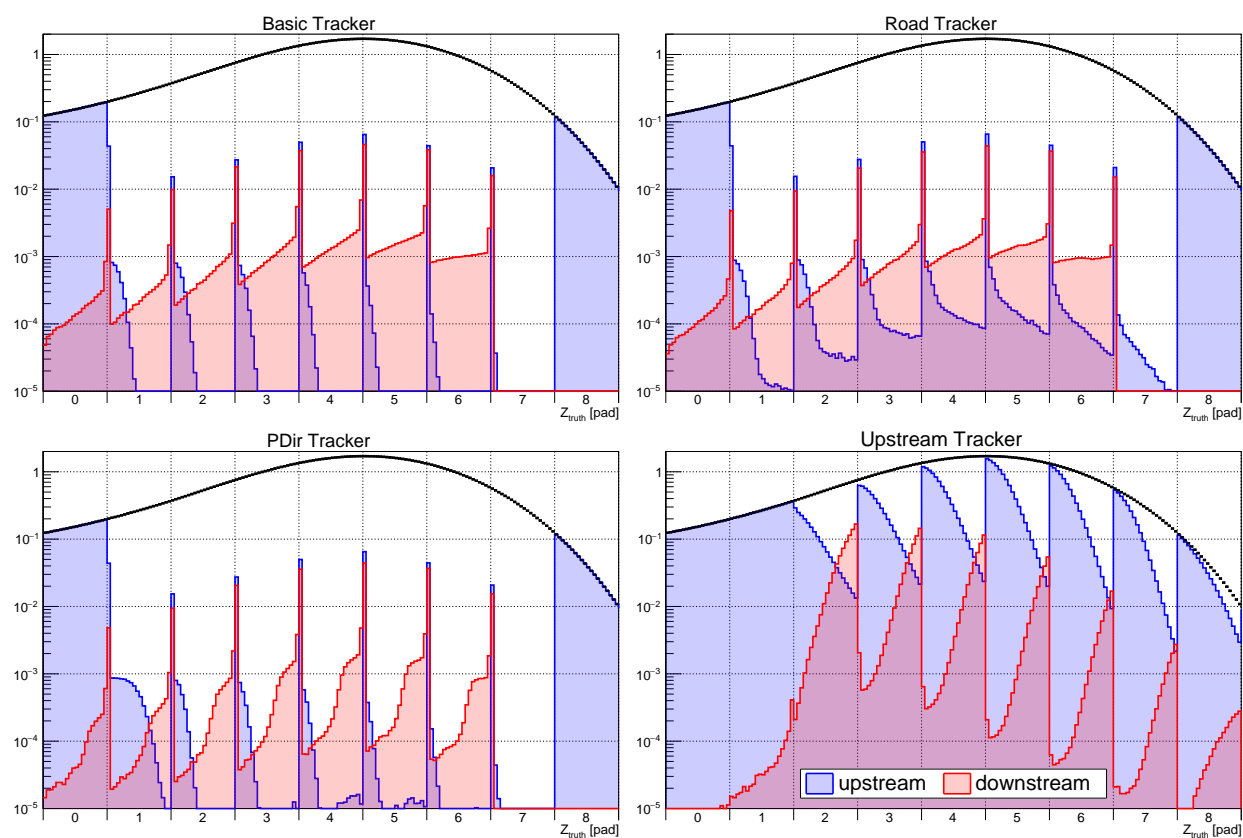


Figure G.1: Tracker migration for events without $p+t$ fusions, relative to the true stop Z position. The stop distribution for all events is shown by the black curve, while the number of events which migrate either upstream or downstream are drawn in blue and red, respectively.

14 also contributes to this. Finally, the long-ranged downstream migrations are the result of forward-going electrons adding energy to the next pad row, which is also described in chapter 14.

The downstream migration effect looks nearly identical for the road tracker, since the road energy threshold is completely irrelevant to downstream electron migration. The short-ranged peaks also look very similar, although there should be a small quantitative difference due to their different handling of ^3He migrations. Finally, the road tracker exhibits additional

upstream migrations caused by events where E_1 is high enough to exceed the road threshold, usually as a result of electron interference adding additional energy. Upstream migrations can occur even for stops at the far end of the pad, this tail is likely due to neutron scattering after $n+{}^3\text{He}$ fusion events.

With the P-Dir tracker the upstream migrations look similar to the basic tracker, showing that the more sophisticated cuts are less prone to false positives than the simple road energy threshold. The downstream migrations look similar to the previous two trackers for stops deep in the pad, but the migration probability drops sharply for stops in the first half of a pad row where the P-Dir electron cut is able to distinguish between migrations and real stops. For pad row 1 there is also a significant increase in upstream migration, in this case caused by tracks that are too short to properly utilize all of the P-Dir cuts depending on E_2 . The standard fiducial Z cut starts at row 2, so these errors are not relevant in most cases.




Figure G.2: Mapping of upstream tracker Z estimates to true stop Z positions.

Finally, the behavior of the upstream tracker is very different from the others, and is best visualized by drawing separate curves for each reconstructed Z bin as shown in figure G.2. This makes it clear that although the upstream tracker is relatively unaffected by

fusions, it is relatively imprecise in general and outputs a wide spread of reconstructed stop positions. The valid Z range for the upstream tracker is also shifted downstream relative to the other trackers, since it requires two upstream pads to perform the fit but can also report projected stop positions outside of the TPC. Finally, if the upstream fit fails for some reason it can prevent the track from being reconstructed, which may produce a subtle correlation to TPC interference effects similar to the crosstalk issue that caused problems for the first row analysis.

G.2 P-Dir Cuts

The Monte Carlo E_1 vs E_0 plots shown in section 9.4 did not include the stop Z correction described above, so while the cuts match the true stop position well it may be better to alter them to deliberately induce migrations similar to those expected for clean events. Figure G.3 shows the E_1 vs E_0 plots with and without the Z correction applied. The majority of events where a ^3He ion crosses the row boundary are now classified as downstream migrations, as even in the case of an upstream-going ^3He the low stop depth means that the final muon pulse would likely fail to pass the pulse-finding threshold and the event would be reconstructed one pad upstream even without the fusion. There is also a significant peak of downstream migrations at the left side of the muon stop band, caused by pileup electrons with a very small E_0 pulse. A similar pileup effect can even be seen at the top edge of the non-migrating p+t fusion band in the bottom center of the plots. The actual proton direction portion of the P-Dir cut is nearly unchanged, while the E_0 cutoffs at either end are now slightly offset. Little can be done about the pileup effect, but some minor adjustments to these E_0 thresholds and to the ^3He cuts may be implemented before the final planned analysis pass.




Figure G.3: E_1 vs E_0 plots for Monte Carlo data, split into downstream migrations (left) and non-migrating (right) events, and with the P-Dir energy cuts shown in red. The top plots use the true stop position as a reference, while in the bottom plots muon stops within the first 0.15 mm are shifted to the previous row to mimic the TPC pulse-finding threshold. This mainly affects the low- E_0 portion of the muon and ^3He bands, where events are shifted from the right plot to the left, and the E_0 thresholds at either end of the p+t cut, which are now slightly misaligned.

Appendix H

MUON BEAM OSCILLATIONS

The time-dependent background observed in chapter 12 raised some concerns about the stability of the accelerator beam rate. If the beam rate fluctuates then muons would be more likely to enter during periods of high beam rate, and would then naturally see a falling background as the beam rate returns towards its average value. These fluctuations would be distinct from the expected fast RF oscillations, instead occurring with a time scale on the order of at least tens of microseconds in order to produce an effect over the length of the event window. Beam rate fluctuations cannot be the sole source of the background slope as they would be averaged out in the μCL data, but it could be a secondary effect that adds to the background slope for real events.

Suppose the beam follows a simple sinusoidal oscillation with amplitude a and frequency f , such that the number of muons as a function of time is given by

$$N(t) = N_0 [1 + a \cos(2\pi ft)]. \quad (\text{H.1})$$

Ignoring details such as time of flight differences, the beam electron time distribution would be identical apart from an additional scale factor b . The electron times are always measured relative to the muon entrance time, so the observed background is obtained by integrating the muon time over a full period:

$$B(t_e) = b \int_0^{1/f} dt_\mu N(t_\mu) N(t_\mu + t_e). \quad (\text{H.2})$$

Performing this integral yields a background oscillation with the same functional form, but reduced oscillation amplitude.

$$B(t) = B_0 \left[1 + \frac{1}{2} a^2 \cos(2\pi ft) \right] \quad B_0 \equiv \frac{b}{f} N_0^2. \quad (\text{H.3})$$

Returning to the basic slope model, the slope parameter dr is approximated by taking the relative background difference between the beginning (t_1) and end (t_2) of the event window:

$$dr = \frac{1}{t_2 - t_1} \frac{B(t_1) - B(t_2)}{B(t_2)} = \frac{a^2}{t_2 - t_1} \frac{\cos(2\pi ft_1) - \cos(2\pi ft_2)}{2 + a^2 \cos(2\pi ft_2)}. \quad (\text{H.4})$$

This equation would be useful if the oscillation amplitude were known, but in this case it is more interesting to solve for the amplitude instead:

$$a = \sqrt{\frac{2dr(t_2 - t_1)}{\cos(2\pi ft_1) - (1 + dr(t_2 - t_1)) \cos(2\pi ft_2)}}. \quad (\text{H.5})$$




Figure H.1: Beam oscillation amplitudes required to produce various observed background slopes.

Figure H.1 shows the required oscillation amplitudes to produce various slope values. A frequency of 20 kHz corresponds to a period of 50 μs , which would make the oscillation minimum precisely aligned with the end of the event window. With a time window starting at $t=0$ the curves would be symmetrical about this center point, but with a delayed start time they are biased towards lower frequencies. The time window has relatively little effect on the overall sensitivity though, so in any case it is clear that percent-level beam fluctuations are required to produce a meaningful background slope.

Such a large oscillation in the beam rate should be relatively easy to measure directly. For each 140 ms MIDAS data block a fast fourier transform is used to convert the muon entrance times into frequency space. The timing of separate blocks is not known precisely, so the phase information is discarded before adding the amplitudes from each block. This procedure results in a significant noise floor which appears to add to the real signals in quadrature.




Figure H.2: Observed beam oscillation frequency distributions. Production data shows significant structure due to the beam kicker and μCL , but the kickerless data from R2016 is clean apart from a small peak at 5 kHz. Backgrounds of 2% and 3.5% have been removed from the production and kickerless data, respectively, and simulated data is processed the same way to verify the procedure.

Figure H.2 shows the measured oscillation frequencies obtained by this analysis, with the constant background removed. A production data file is shown for reference, but exhibits substantial structure from known effects related to the kicker and the muon clock. For a significant portion of R2016 the kicker was not in working order, and analyzing this dataset allows the true beam rate to be monitored without artificial distortion by the MORE system.

The kickerless data does require a lower beam rate which in turn nearly doubles the noise floor, so a basic simulation was also used to verify that an input oscillation could be accurately identified.

The kickerless data has a nearly flat frequency distribution apart from a single peak at 5 kHz with an amplitude of approximately 0.5%. Inserting the observed amplitude and frequency into equation H returns a dr value of 2×10^{-7} , more than two orders of magnitude smaller than the actual estimated slope. Thus, while it appears the beam rate does oscillate slightly, the amplitude is much too small to affect the MuSun measurement.

Appendix I

NUCLEAR ISOMERS

Section 12.3.4 described a possible time-dependent contribution to the background from slow decays following wall stops. This would occur due to the production of specific excited nuclear isomers that are much longer lived than one would normally expect for such states, with half-lives potentially on the order of the 25 μs event window length. This appendix describes the effect in more detail and discusses the isotopes that might be created by muon stops in various parts of the detector. It also summarizes some of the tests performed in an attempt to definitively demonstrate the effect.

I.1 Long-Lived Isomers

As mentioned in section 2.2.3, muon capture on a high- Z nucleus leaves it in a highly excited state. Often the nucleus de-excites by emitting one or more neutrons, but this may still leave the resulting nucleus in an excited state. The remaining energy is typically lost via the emission of gamma rays. In most cases this gamma ray cascade will quickly return the nucleus to its ground state, but certain isotopes have metastable states with much longer lifetimes.

Excited states of a given isotope are known as nuclear isomers, and many isomers have well measured excitation energies and half-lives and are included in tables of isotopes. Thus, while the experiment should not produce any relevant radioactive isotopes it may produce several interesting excited isomers listed in table I.1. Measurements of neutron multiplicity as well as proton and alpha production probabilities for muon capture on various elements[50] are used to estimate the probabilities of producing each isotope, although the fraction that result in a specific nuclear isomer is harder to determine.

Source	Isotope	Probability [%]	Half-Life [μ s]	Energy [keV]
Aluminum	$^{26}_{11}Na$	0.7	4.35(16)	82.4(4)
Nickel	$^{58}_{27}Co$	23	10.4(3)	53.15(7)
Tungsten	$^{180}_{73}Ta$	11	∞	77.1(8)
			31.2(14)	1452.40(18)
			2.0(5)	3679.0(11)
			17(5)	4172.2(16)
	$^{181}_{73}Ta$	19	6.05(12)	6.238(20)
			18(1)	615.21(3)
25(2)			1485(3)	
210(20)			2230(3)	
Lead	$^{202}_{81}Tl$	3	572(7)	950.19(10)
	$^{203}_{81}Tl$	4	7.7(5)	3400(300)
	$^{204}_{81}Tl$	13	63(2)	1104.0(4)
			2.6(2)	2500(500)
	$^{205}_{81}Tl$	22	2.6(2)	3290.63(17)

Table I.1: Selected isotopes produced by muon capture, with half-lives and energies of potentially relevant nuclear isomers[41]. Transition probability estimates are based on neutron multiplicity measurements[50], and only reflect the total probability of producing the listed isotope with no regard to individual isomers. The 77 keV state of tantalum-180 is the only known observationally stable nuclear isomer, unlike the actual ground state which has a half-life of 8 hours.

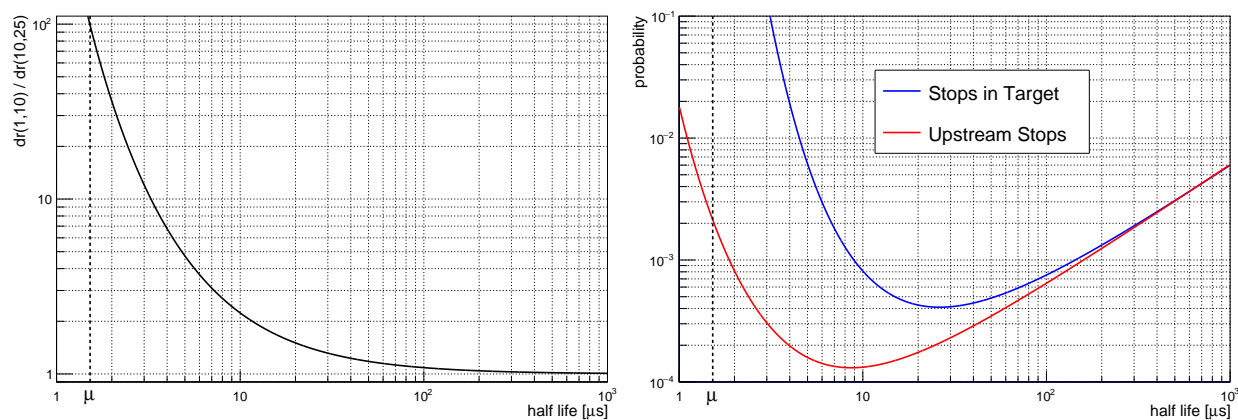


Figure I.1: Left: Ratio of dr values measured from 1 to 10 μs and from 10 to 25 μs . Right: Slow decay detection probability required to produce a dr value of 10^{-4} , measured starting from 10 μs . Upstream muon stops are not affected by pileup protection and thus have a larger impact for fast decays, but should be rare.

If the background slope is really an exponential decay then the half-life must be longer than the muon decay half-life of 1.6 μs . In fact it must be several times that to reproduce the observed lack of background curvature, as shown in figure I.1. With a half-life of 10 μs the slope early in the event is roughly double that later in the event, which may be consistent with R2014 but is still too much curvature for R2015.

Figure I.1 also shows the sensitivity of the slope measurement to an exponential decay signal as a function of half-life. The sensitivity is highest around a 10 μs half-life, and for half-lives longer than about 1 ms the required number of slow decay events starts to become unreasonably high. For faster decays the muon stop position is important, as upstream muon stops in the beam pipe or collimators are not affected by the pileup protection cuts and can occur as late as the kicker step 800 ns after the primary muon entrance. On the other hand muons that stop somewhere in the target after passing through the entrance detectors should have a much higher detection probability, but any entrances within the pileup window will be eliminated and the effect must instead come from the slow decay tail left by muon stops

in the previous event. The Tantalum isotopes produced by stops in tungsten are particularly interesting, as the existence of multiple excited isomers may allow a single capture event to initiate a chain of several slow decays.

Finally, these decays should emit an electron if they are to produce a background in the eSCs. Because the issue is related to muon stops in solid material, it is possible for a gamma ray emitted by the isomer transition to scatter and emit a secondary electron as it exits the material. Nuclear isomers may also decay via a process known as internal conversion, in which an orbital electron absorbs the extra nuclear energy and is ejected from the atom. The internal conversion probability can be quite large for heavy atoms and low transition energies, for instance a 6 keV transition in tungsten-181 decays this way 98% of the time[43] and the 6 keV isomer of tantalum-181 is presumably similar. This particular state has too little energy to trigger the eSCs, but internal conversion is still likely to be a significant contribution for the other states in table I.1.

I.2 Locations of Stopped Muons

Estimated probabilities of muons stopping in various parts of the experiment are shown in table I.2. These are grouped into three broad categories: Muons that stop upstream before reaching the entrance detectors, muons that pass the entrance detectors but stop before reaching the TPC, and muons that do stop in the TPC. In the latter case the stop position can be determined from the track, with the majority of muons stopping in the deuterium gas but also a significant contribution from wall stops at the Y boundaries. Very few muons stop upstream, although as shown in figure I.1 these events may have an out-sized effect particularly for faster decays since they are not affected by pileup protection. Finally, roughly a third of all incoming muons stop before reaching the TPC, so even a relatively small fraction of these producing slow decays could potentially explain a 10^{-4} level background slope effect.

The stops without a TPC track should be a mixture of muons directly impacting the front flange of the deuterium vessel and of muons that scatter in the entrance window and

Material	Location	Stops [%]
Upstream Stops		2
Steel	Beam Pipe	≈ 0
Lead	Entrance Collimator	< 2
Entrances Without TPC Tracks		34
Beryllium	Entrance Window	≈ 0
Steel	Entrance Flange, TPC Frame	16
Silver, Tungsten	TPC Frame	10
Deuterium	Gas Around TPC	8
Aluminum	Deuterium Vessel	< 1
Entrances With TPC Tracks		64
Silver	Cathode	1
Tungsten	Anode, Frisch Grid	4
Deuterium, Carbon	Punch-through	< 0.5
Deuterium	Outside Fiducial Volume	24
Deuterium	Fiducial Volume	34

Table I.2: Approximate probabilities of muons stopping in different materials.

then stop somewhere inside the target but outside the active volume of the TPC. Scattered muons may still stop in deuterium gas, but should also have a high probability of stopping in the steel and tungsten components of the TPC frame. Figure I.2 shows the front of the deuterium pressure vessel, the flange holding the entrance window is mounted in front of the aluminum pressure vessel so off-axis muons are likely to hit that rather than stopping in the aluminum itself. The flange is made of AISI 321 stainless steel made with approximately 17-19% chromium and 9-12% nickel, so these stops should have a roughly 2.5% chance of producing cobalt-58 which had a relevant isomer listed in table I.1.




Figure I.2: Front view of the aluminum pressure vessel containing the TPC, showing the beryllium entrance window welded to a stainless steel flange. The copper color comes from a thin layer of material deposited during the welding process, and should have no significant effect.

The approximate stop fractions for the entrances without a TPC track are estimated by fitting the observed lifetime distribution for these events, including a muon lifetime component for deuterium stops and additional medium and fast decay components for stops in steel and in high-Z materials. Silver and tungsten are difficult to distinguish with this method,

but assuming stops in the TPC frame have a similar Y distribution to the stops in the TPC itself we can estimate a total probability of roughly 12% for stopping in tungsten and 3% for silver. This analysis has only been performed for the combined R2015 data, and the exact probabilities would vary depending on the beam tune for a specific dataset. An increased probability of stopping in steel may explain the higher slope and slight curvature observed in the R2014 clock data, which are roughly consistent with the $10 \mu\text{s}$ half-life of the expected cobalt-58 isomer. It may be interesting to repeat the fit procedure for individual datasets to check whether the differences background slope correlate with changes in the estimated stop fractions in different materials.

I.3 Testing




Figure I.3: Fitted decay rates for double electron events using either standard HBE cuts or also applying μSCa pileup protection. The error bars on the $\mu\text{SCa PP}$ points indicate the uncertainty in the difference from the standard HBE fit.

Some attempts were made to definitively prove the existence of this slow decay back-

ground, but the small size of the effect makes it difficult to isolate in real data. There was more luck with the beam background test data from R2016, where muons stopped in the lead plate produced a clear slow exponential decay signal as shown in section 12.3.4. These events should be fairly similar to stops in the lead collimator attached to the μ SCa, and by adding the μ SCa to the pileup protection we can eliminate most of these events and thereby estimate their effect. Figure I.3 shows the change in lifetime fit results from applying this cut, and while it does produce about a 10% change in the fitted rate it is clearly not the dominant effect.

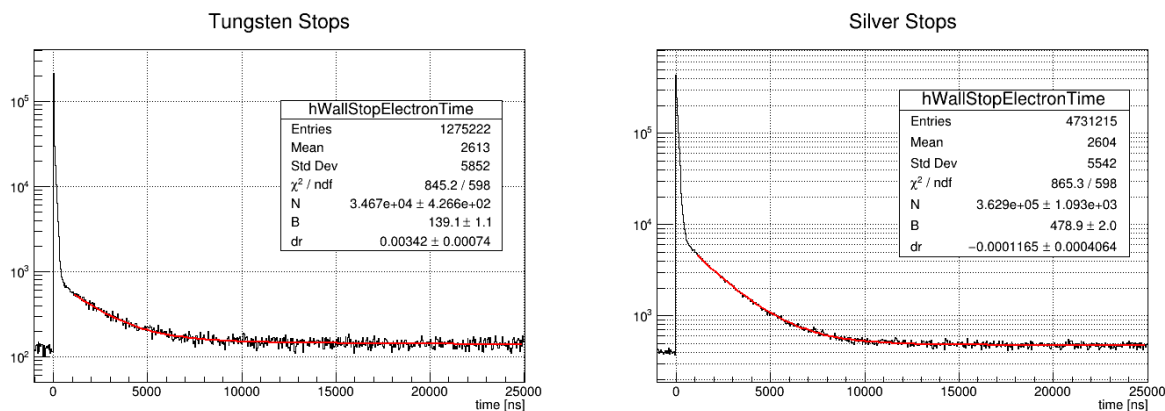


Figure I.4: Background slope estimates using wall stop calibration data. The slower exponential component comes from muon stops in scintillator itself, and are assumed to decay according to the muon lifetime in carbon. The important fit result is the slope parameter dr , which is consistent with 0 for silver but finds a significant positive value for tungsten.

Stops in tungsten appear to be a promising source of slow decays, and the R2016 wall stop calibrations included a dataset with a tungsten target. This was intended for measuring the response of the neutron detectors to the large prompt capture signal so the statistics are relatively low, but a linear fit to the background does find a significant slope as shown in figure I.4. Applying this analysis to the data for a silver target in the same setup shows

no significant slope, indicating that this is a real effect rather than some artifact of the apparatus or fitting procedure.

The estimated 3.4×10^{-3} dr value for the calibration data should be reduced to approximately 2×10^{-4} for real data, taking into account the $\sim 12\%$ muon stop fraction in tungsten and a factor of 2 reduction to account for the delay of roughly one half-life between muon entrances. This is still about twice the observed slope for R2015, leaving room for additional differences in detection efficiency and beam background rate between the calibration data and the real experimental conditions. As mentioned above the larger slope in R2014 could be produced by changes in the stopping distribution, and additional testing of the wall stop fractions for different datasets may be able to demonstrate this.

1.4 Conclusion

The MuSun experiment includes several materials that could produce metastable nuclear isomers as a result of muon capture, with half-lives on the order of 10s of microseconds which are suitable for producing a small background slope across the length of the event window. In particular there appears to be a large fraction of muons which stop in either stainless steel or tungsten, the latter of which can produce many relevant isomers of tantalum while the former can create an isomer of cobalt due as a result of captures on the nickel atoms included in the steel. The decays of these isomers should have a significant probability of emitting a detectable electron background, either due to scattering of gamma rays in the wall materials or by directly emitting electrons via internal conversion. The exact efficiencies for both creating and detecting the decays of specific nuclear isomers is unknown, but some evidence of this effect has been seen in R2016 test data where the muon beam stopped directly in solid target materials. Stops in tungsten appear sufficient to explain the observed background slope in R2015, although these measurements are not definitive and additional testing would be required to fully understand the R2014 behavior.

Appendix J

FUSION TAG

Section 13.2.1 discussed the cuts used to identify fusion events in the TPC, which are used to study how fusion-related tracking errors can distort the muon lifetime. This appendix details some modifications to extend the fusion so it works near the edges of the TPC, where fusion products may escape the active volume. The fusion tags will inevitably be less effective in this region, but with careful tuning it is possible to maintain a reasonable balance of detection efficiency and accuracy.

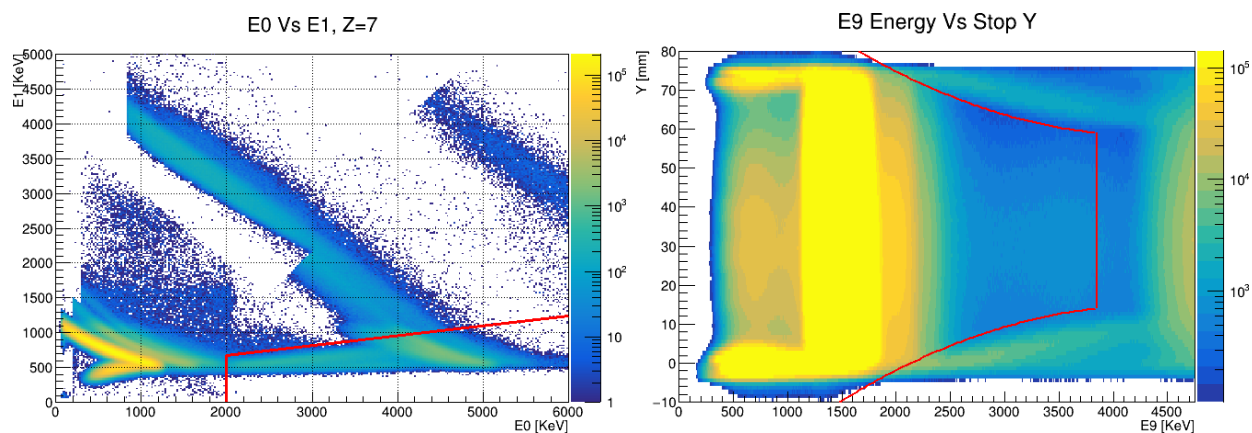


Figure J.1: Left: E_0 Vs E_1 cut to identify fusion punch-throughs on the last Z row, indicated by the bottom-right region bordered in red. Left: modified E_9 threshold which attempts to capture the degraded p+t energy at extreme Y values while still avoiding normal muon stops and ^3He fusions. All events to the right of the cuts are identified as fusions.

As usual the primary concern is the proton from p+t fusion events, with a range of roughly 13 mm. If these escape the active volume then it reduces the observed fusion energy,

potentially causing the event to fall below the E_9 threshold. Punch-through events in the final Z row may be identified by adding a 2D E_0 vs E_1 cut similar to those discussed in chapter 9. There is less information available in Y , but the energy losses may be accounted for by reducing the E_9 threshold near the boundary as shown in figure J.1. The X resolution is insufficient to determine whether a proton can reach the edge so little can be done about this direction, although a reduced threshold for the edge pads may also be a good idea.

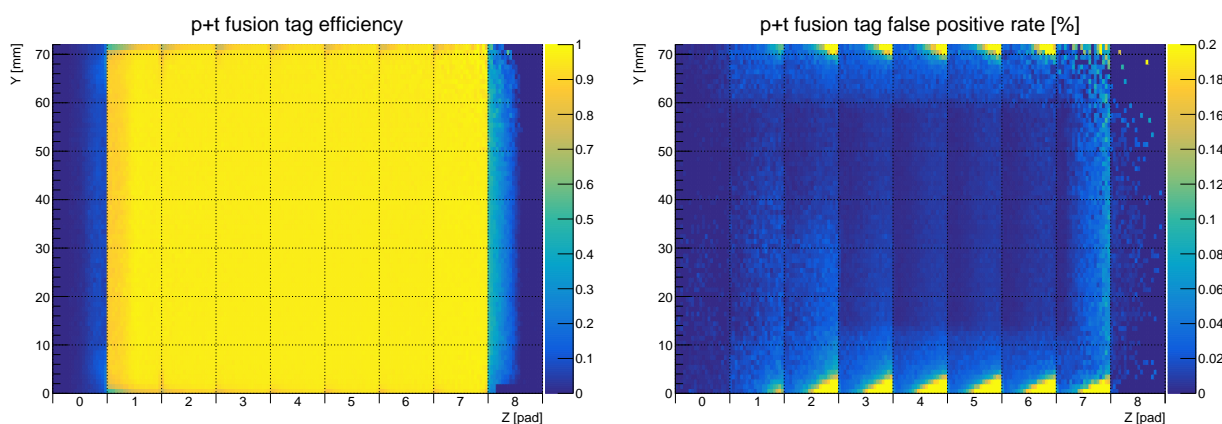


Figure J.2: E_9 p+t fusion tag efficiency (left) and false positive rate (right).

This modified version of the fusion tag has high efficiency for most stops within the TPC. However, it does have some issues with false positives close to the edges in Y where the threshold becomes quite low, as shown in figure J.2. It also is not particularly useful for muon stops that are actually outside of the active volume, although the punch-through cut does still identify some events stopping just past the end of the TPC.

Appendix K

BACKGROUND DIP

Chapter 13 assessed the impact of fusion events on the muon lifetime measurement, and discussed the determination of the early time deviation, D , as a way to quantify the amount of fusion migration. The biggest challenge when attempting to accurately determine D is that because it only uses the lifetime histograms those spectra must be free of any other systematics or backgrounds. At this point the behavior of the kicker-induced background step and the downward slope of the background are quite well understood, and relevant systematic effects such as wall stops are also well controlled. However, a new observation made during the early time deviation analysis was the presence of a small dip in the background just before the muon entrance, as shown in figure K.1.




Figure K.1: Background dip observed just before the muon entrance using standard HasBestEntrance (HBE) cuts. The dip disappears with different μ PC pileup protection variations, being replaced with a peak instead.

The background dip was eventually identified as an interaction between the pileup protection system and a flaw in the μ PC clustering algorithm. Electrons do not deposit enough energy in the μ PC to trigger both the X and Y wires reliably, but often trigger a single layer without the other. To avoid unnecessary losses due to pileup protecting on these electrons, the pileup protection cuts typically require an X-Y coincidence instead of single clusters.

These coincidences are formed from all pairs of X and Y clusters within the 190 ns coincidence window, allowing a single X or Y cluster to be used in multiple X-Y pairs if there is more than one matching cluster. In most cases this double-counting has no effect, as the pileup protection merely checks for the presence or absence of any μ PC pairs and makes no distinction between different quantities. The exception is near $t=0$, where there must be a single X-Y pair corresponding to the muon entrance itself. An unpaired X or Y cluster produced by a beam electron may thus be promoted to a full X-Y coincidence by grouping with one of the clusters from the entrance muon, even though that cluster will have already formed a separate pair with the real matching muon cluster. Figure K.1 shows that the dip disappears when μ PC pileup protection is either disabled or applied strictly even to single unpaired clusters, as in either case the pileup protection is then consistent regardless of whether this effect occurs.

Before the dip was fully understood an attempt was made to incorporate it into the background model, but after the cause was identified this was abandoned for several reasons. Firstly, the background increases again after the initial dip due to a separate effect where muon and electron clusters on the same plane overlap and merge together, rendering pileup protection ineffective. The shape after the muon entrance is therefore difficult to model, and most likely consists of a brief excess from these overlapping clusters followed by another small dip from clusters which are within the coincidence window but do not overlap. Secondly, the ability to remove the dip with the strict pileup protection cut provides an easy way to see the impact this effect has on the results, which turns out to be negligible. Finally, a software update preventing this double-counting has also been implemented and will be included in the final analysis pass, though no major differences are expected.

Appendix L

DOWNSTREAM ELECTRON MIGRATION MODEL

Section 14.5 described the modeling of downstream electron migration events, and an estimate of their effect on the muon disappearance rate. That analysis used a two-parameter model to approximate the migration probability as a function of stop depth, which is described in more detail in this appendix.

For a migration to occur the electron pulse must be clustered with the muon stop, giving a coincidence window of ± 1 pad in X and approximately $\pm 2 \mu\text{s}$ in time. Using the drift velocity of $5 \text{ mm}/\mu\text{s}$ and the pad width of 18 mm , this essentially means the electron signal will be accepted if it occurs within a rectangular volume 54 mm wide, 20 mm tall, and also 16 mm long since the electron signal can occur at any depth within the target pad. The exact solid angle will vary with the muon stop depth, but also with the X position within the pad as well as the Y offset from the decay time.

To simplify the problem the X and Y effects will be ignored, and the acceptance window will instead be approximated as a simple disc with radius W located at the downstream pad boundary. This is similar to the p+t fusion migration model described earlier, except in this case the range is effectively unlimited and instead the width of the intersection area is kept fixed. This results in a migration probability of

$$P_{mig}(d) = \frac{\alpha}{2} \left(1 - \frac{1 - d/l}{\sqrt{W^2 + (1 - d/l)^2}} \right), \quad (\text{L.1})$$

where d is the stop depth, l is the pad length, and α represents the probability of the electron depositing enough energy to exceed the pulse-finding threshold.

By ignoring the Y dependence this simplified migration probability has eliminated the vertical asymmetry which was the primary feature of the electron migrations. To restore

this behavior, it is assumed that all events have approximately the same time dependence as is observed for the tagged low- E_1 forward migration events. The detection efficiency is also assumed to be the same for all events in each pad row, although it varies significantly between rows and in principle should be calculated for each pad individually.




Figure L.1: Comparison between electron migration E_1 distributions predicted with various values of W , and the actual E_1 distribution for events with E_2 below 200 keV.

The W parameter is calibrated by matching the shape of the predicted E_1 distribution to the observed low-energy events, as shown in figure L.1. A value of 12 mm appears to reproduce the observed distributions well, and is similar to the 10 mm radius one would predict from the $2 \mu\text{s}$ clustering window. An error of ± 1 mm in W is still reasonably consistent with the observed E_1 distribution, giving an uncertainty of about 10% in the height of the peak.