

©Copyright 2019

Rachel Ryan

MuSun: A Precision Measurement of Nuclear Muon Capture in Deuterium with a Cryogenic Time Projection Chamber

Rachel Ryan

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

Doctor of Philosophy

University of Washington

2019

Reading Committee:

Peter Kammel, Chair

David Hertzog

Jason Detwiler

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

MuSun: A Precision Measurement of Nuclear Muon Capture in Deuterium with a Cryogenic Time Projection Chamber

Rachel Ryan

Chair of the Supervisory Committee:
Professor Peter Kammel
Physics

The MuSun experiment aims to measure the muon capture rate on deuterium to 1.5%, allowing for the extraction of a low energy constant required to quantify the axial two body current in the framework of modern effective field theories. Once obtained, this constant calibrates the two nucleon sector, enabling calculations of related astrophysical reactions of interest, including proton-proton fusion, an important input for stellar models, and the neutrino-deuterium breakup reactions required by the Sudbury Neutrino Experiment to quantify the neutrino flux from the sun. MuSun utilizes the lifetime technique – extracting the capture rate via a comparison of the negatively charged muon disappearance rate in deuterium to the free muon lifetime– in combination with a novel cryogenic time projection chamber. The full statistics necessary to reach the precision goal have been acquired in two separate data collection periods. This thesis presents the account of the first of these production datasets, from detector development to the final analysis. The construction, commissioning, and performance of the time projection chamber is described, followed by details of the 2014 data collection period. Aspects of software development and several in depth analyses are described. Based on this high-statistics analysis, limits are set on the effect of all sources of systematic error and several correction strategies are developed.

TABLE OF CONTENTS

	Page
Chapter 1: Introduction	4
Chapter 2: Motivation and Theory	9
2.1 Weak interaction and the Fermi Constant	9
2.2 Single Nucleon Sector and the Pseudoscalar Coupling	12
2.3 Few Nucleon Interactions	15
2.3.1 QCD Symmetries	16
2.3.2 Chiral Perturbation Theory	17
2.3.3 Pionless Effective Field Theory	22
2.3.4 Astrophysical Significance	23
2.4 Summary	25
Chapter 3: Design Solutions for Experimental Challenges	26
3.1 Precision	26
3.2 Muon Kinetics	28
3.2.1 Atomic and Nuclear Capture	28
3.2.2 Molecular formation and Muon Catalyzed Fusion	30
3.2.3 Optimizing running conditions	32
3.3 Muon Capture On Other Nuclei	33
3.3.1 Capture on Target Materials	34
3.3.2 Gas Purity	36
Chapter 4: Experimental Design	39
4.1 Blinding	40
4.2 Beamline	40

4.3	Muon Entrance Detectors	42
4.4	Muon Stop Determination	45
4.4.1	Time Projection Chamber (TPC)	45
4.5	Decay Electron Detectors	46
4.5.1	Electron Proportional Chamber	46
4.5.2	Electron Scintillator Hodoscope	47
4.6	Neutron Detectors	48
Chapter 5:	Time Projection Chamber Upgrades	49
5.1	Principles of Operation	49
5.2	Motivation for upgrades	51
5.2.1	Stable Grid and Cathode Voltages	51
5.2.2	Wall Capture Materials	53
5.2.3	Electronic Resolution	53
5.3	New TPC Construction	54
5.4	Electronic Readout	58
5.5	Commissioning	60
5.6	Understanding Detector Response	61
5.6.1	E0 vs E1	63
5.6.2	E9 Energy	65
5.6.3	MuSun Monte Carlo	65
5.6.4	Effect of PT fusions on Muon Tracking	66
5.6.5	Effect of n-d Scatter Events	70
Chapter 6:	R2014 Dataset Overview	72
6.1	Data Collection Summary	72
6.2	Datasets	73
Chapter 7:	Analysis Overview	76
7.1	Software Overview	76
7.2	Muon Entrance	77
7.3	Muon Stop Determination with the TPC	81

7.3.1	Pulse Processing in the TPC	83
7.3.2	Muon Tracking in the TPC	90
7.4	Electron Definition	94
7.4.1	Electron Tracking with the Multi-Wire Proportional Chambers	94
7.4.2	Simplified Scintillator-Only Electron Definition	98
7.5	Trend Histograms and Quality Cuts	107
7.5.1	Initial Quality Cuts	107
7.5.2	Error Conditions	108
7.6	Lifetime Histograms and Fitting	112
7.6.1	Beam RF Oscillations	114
Chapter 8:	Systematics	119
8.1	Muon-On-Request scheme	119
8.1.1	Entrance Efficiency	120
8.1.2	Electron Background	120
8.2	Muon Track Interference	120
8.2.1	Fusion Interference	121
8.2.2	Electron Interference	123
8.3	Muon Capture on Other Nuclei	129
8.3.1	Isotopic Purity	129
8.3.2	Chemical Purity	131
Chapter 9:	Muon On Request and Accidental Structures	138
9.1	Muon Pileup	142
9.1.1	Time Dependent Inefficiency from muSC Deadtime	142
9.1.2	Time Independent Entrance Detector Efficiencies	143
9.2	Electron Background	154
9.2.1	Observing a Time Dependent Beam Electron Background	155
9.2.2	Electron Multiplicity Technique	159
9.2.3	Quantifying the MORE Induced Background	161
9.2.4	muCL Background Subtraction	163
9.2.5	Quantifying other Backgrounds	168

9.2.6	Characterizing the Beam Electron Background	172
9.3	Background Model	181
9.4	Monte Carlo Background	184
9.5	Conclusions	188
Chapter 10:	Muon Catalyzed Fusion Interference Correction	190
10.1	Fusion Formalism	190
10.1.1	Characterization of the fusion correction function	192
10.1.2	Lifetime correction strategies overview	193
10.2	Application of Fusion Fraction Method to Data	199
10.2.1	Determining $\tilde{\epsilon}$	199
10.2.2	Determining ϵ	204
10.2.3	Determining κ	208
10.3	Applying correction to Data	209
10.3.1	Nominal Stop Distribution	209
10.3.2	Balanced Stop Distribution	210
10.4	Uncertainty Estimates and Conclusions	213
10.5	Application of Fusion Fraction Method to Monte Carlo	214
10.5.1	Determining $\tilde{\epsilon}$	215
10.5.2	Determining ϵ	215
10.5.3	Determining κ	216
10.5.4	Monte Carlo Correction	219
Chapter 11:	Consistency Checks	221
11.1	Dataset Consistency	221
11.2	Fit Variations	223
11.2.1	Lifetime Vs Fit Start Time	223
11.2.2	Lifetime Vs Fit Stop Time	224
11.3	Lifetime vs TPC Stop Position	225
11.3.1	Lifetime vs Stop X	225
11.3.2	Lifetime vs Stop Y	227
11.3.3	Lifetime vs Stop Z	228

11.4 Lifetime vs Electron Studies	229
11.4.1 Lifetime vs eSC Segment	231
11.4.2 Lifetime vs electron Z Position	231
Chapter 12: Conclusions	234
12.1 Quality Cuts and Data Consistency	234
12.2 Systematics	235
12.3 Outlook	238
12.4 Summary	239
Bibliography	241
Appendix A: Kinetic Solutions	248
A.1 Quartet Population	249
A.2 Doublet Population	250
A.3 μ^3He Population	251
A.4 Combined equation	251
A.5 Solutions for MuSun Conditions	252
Appendix B: Resolving the kinetic equations with impurity captures	254
Appendix C: Resolving the kinetic equations with protium impurity	257
Appendix D: Quadrapole Current Error and Effect on Beam Tuning	259
Appendix E: Error Inflation for Electron Double Counting	261
E.1 Sources of doubly counted electron events	261
E.2 Error correction for doubly counted electrons	261
E.3 Determination of double counted fraction	262
E.4 Error correction for double counted events	264
E.5 Correction for non-correlated events	265

ACKNOWLEDGMENTS

I cannot imagine making it through this without the support of so many wonderful collaborators, friends, and my family. I am grateful for the opportunity to develop not only research skills, but life-long friendships with the graduate students of the MuSun collaboration. None of the work in this thesis would have been possible without the foundations provided by Michael Murray, Frederik Wauters, Xiao Lou, Luis Ibanez, and many more. Extra thanks are necessary for Michael and Frederik for patiently answering all my questions in the early years. I couldn't have made it to the finish without the immense support of Ethan Muldoon and Dan Salvat. Thanks Dan and Ethan for being such a creative and fun team to work with!

During the periods of data collection, this experiment afforded me the opportunity to develop as a researcher under many knowledgeable collaborators from around the world. Claude Petitjean, thank you for being such a warm and welcoming host and teaching me the intricacies of beam tuning. To my PNPI colleagues— Alexei Vorobyov, Alexander Vasilyev, Peter Kravtsov, Marat Vznuzdaev, Nikolai Voropaev, Polina Kravchenko, Alexander Nadochy, and Kuzma Ivshin— I cannot thank you enough for all I have learned from you and I hope to see you all again in the future! Thanks to Malte Hildebrandt, Bernhard Lauss, and Florian Barchetti for creating such a lovely work experience at PSI. I never would have understood how to keep this experiment running without the guidance of Fred Gray, Rob Carey, Tim Gorringer, and Wes Gohn.

The CENPA staff and faculty have created a unique environment conducive to helping young graduate students flourish. The construction of the TPC would not have been possible

without the help of the staff in the machine shop or the detailed CAD drawings developed by John Amsbaugh. Tim van Wechel and David Peterson, thanks for all your electronic expertise over the years. All of the faculty at CENPA contribute to a community driven learning environment, but in particular I would like to thank David Hertzog. Dave, your guidance and encouragement has meant a tremendous amount to me over the years. Duncan Prindle, thanks for the years of advice on everything from software to the best hiking spots. It would take me another ten pages to really thank my advisor, Peter Kammel. Peter, I can't thank you enough for everything you've taught me – from nuclear physics to back-country skiing. I never would have made it through this tough analysis without your kindness and support. It has truly been an honor to work beside such a passionate, dedicated, and brilliant physicist.

Completing this analysis has not been an easy task, and I'm grateful to have had the support of so many amazing friends and family. Jenny Jensen, thanks for taking the time to explore the world with me, and all the food it has to offer. John Fuini, you were right all along; I could do it. Thanks for keeping me going, friend. Kyle Hadley, I couldn't imagine a better partner to keep me sane during the rough times. Laura Bodine, Brett Wolf, Michael Kruse and Stefan Janiszewski your love and support has been invaluable over the years. To my wonderful landlords Stan and Mary Jeppesen, your encouragement is always appreciated, both academically and in the garden. Thanks for providing me with such a lovely place to call home. To Eva Kammel and Ivona Prindle, you have been strong female role models, and I can't thank you enough for that. And last but certainly not least, I'd like to thank my family for sticking by me all these years. Dad, you fostered my scientific curiosities from an early age, and your endless support is the reason I made it this far. Mom, thanks for teaching me to be unabashedly myself, no matter what. To my brother Jason, even on my worst days I still felt like you were looking up to me and I hope you know how much I needed that.

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).

Chapter 1

INTRODUCTION

The visible matter of our universe consists of stable particles— electrons, and strongly bound states consisting of up and down quarks, interacting via their respective force carriers. However, beyond the typical energy scales of everyday life lie two more generations of these fundamental particles, carbon copies in every way except one: mass. Why these generations of each particle type exist is still a mystery. Of the three generations of leptons, the muon has key characteristics which make it particularly useful for probing the electroweak and strong forces. As shown in table 1.1, its properties are measured with outstanding precision, and stringent limits on lepton flavor violating processes have been established. It is sensitive to high mass scales, yet light enough that it may be readily produced, and its $2.2 \mu s$ lifetime lends enough time for precision measurements. Additionally, it only interacts via electroweak forces at tree level, allowing for clean calculations and interpretations.

Of particular use is the muon's ability to undergo atomic capture, replacing the valence electron to form an atom. In Gaussian units, the Bohr radius is given by

$$a_0 = \frac{\hbar^2}{me^2} \quad (1.1)$$

where e is the electric charge and m is the lepton mass. Because the muon is 200 times heavier than the electron, the radii of muonic atoms are 200 times smaller than their electronic counterparts. The probability of finding the orbiting lepton at the nucleus is quantified by the solution of the hydrogenic wavefunction evaluated at zero

$$|\psi(0)|^2 \propto \left(\frac{Z}{a_0}\right)^3 \quad (1.2)$$

Property	Symbol	Value	Precision
Mass	m_μ	105.6583745(24) MeV	22.7 ppb
Mean Lifetime	τ_μ	2.1969811(22) μs	1 ppm
Magnetic Moment Anomaly	$a_\mu = \frac{g-2}{2}$	116592089(63) $\times 10^{-11}$	540 ppb
Branching Ratio	Value	LFV modes	90% CL
$\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$	$\approx 100\%$	$\mu^- \rightarrow e^- \gamma$	$< 4.2 \times 10^{-13}$
$\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu \gamma$	$6.0(5) \times 10^{-8}$	$\mu^- \rightarrow e^- e^+ e^-$	$< 1.0 \times 10^{-12}$
$\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu e^+ e^-$	$3.4(4)^{-5}$	$\mu^- \rightarrow e^-$ conversion	$< 7.0 \times 10^{-13}$

Table 1.1: Selected properties, decay modes, and lepton flavor violating (LFV) limits of the muon, taken from the particle data group [45].

where Z is the atomic number, such that the overlap for the muon is 10^7 times larger than that of the electron. This was most recently exploited in muon spectroscopy measurements to determine the proton radius. Because of this large enhancement of the overlap, nucleon and nuclear size effects are intensified compared to electron spectroscopy, allowing for a 0.1% measurement of the proton radius [18]. Surprisingly, the results from the muon spectroscopy measurement differ from the previous electronic measurements by nearly 7 standard deviations. This exciting problem spurred ample new interpretations and experiments which are currently in progress, but remains unsolved to this day [40, 41].

The same enhancement factor also allows for the weak interaction conversion of the proton into a neutron,

$$\mu^- + p \rightarrow n + \nu_\mu \tag{1.3}$$

a process called muon capture. The muon thus provides a probe of low energy nuclear physics and QCD, a regime which has been notoriously difficult to constrain with first principle calculations. Modern effective field theories have allowed calculations of nucleon and even

few body nuclear interactions in this range using Lagrangians constructed based on the underlying symmetries of QCD. In order to reduce the theory to relevant degrees of freedom, the high energy physics is integrated out, resulting in low energy constants (LECs). There is currently significant effort to obtain these LECs via lattice calculations and precision experiments. In the following, a program of muon capture measurements aiming to test chiral symmetry predictions and determine an important LEC is discussed.

The first experiment in the muon program measured the muon capture rate on the proton. The MuCap collaboration determined this rate to the 1% level, allowing for extraction of the pseudoscalar coupling constant, g_p , which agrees well with theoretical predictions [49]. The next step in this program is the MuSun experiment, which aims to precisely measure the muon capture rate on deuterium,



The theoretical two nucleon sector currently contains one LEC which is not well known, referred to as either L_{1A} or d_R , which quantifies the coupling of the axial two body current. Once the LEC is extracted from this measurement, other relevant reactions of astrophysical significance can be calculated to the same degree of precision, including solar proton-proton fusion and the charged and neutral deuterium breakup reactions. Furthermore, as this LEC is currently obtained from the more complex three body system, its determination in the cleaner two nucleon system serves as a consistency check of the theory. Although it is beyond the scope of this thesis, it is worth mentioning that recently a first pioneering calculation of L_{1A} was achieved on the lattice [44].

Previous attempts to measure this capture rate have proved inconclusive, as shown in table 1.2. All attempts have yielded large uncertainties, and the most precise of these measurements deviates from the theoretical prediction by nearly three standard deviations. Several experimental challenges complicate a clean determination of the muon capture rate. The complex muon kinetics within deuterium gas yield capture from two different hyper-

Experiment	Target	T(K)	Detection	$(\Lambda_d(s^{-1}))$	Ref.
Wang 1965	liquid H_2/D_2	20	neutron	365 ± 96	[48]
Bertin 1973	gaseous H_2/D_2	300	neutron	–	[13]
Bardin 1986	liquid D_2	20	electron	470 ± 29	[9]
Cargnelli 1989	<i>gaseous</i> D_2	45	neutron	409 ± 40	[17]

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 upon capture.

fine states, as well as muon-catalyzed fusion processes. As the V-A structure of the weak interaction suppresses capture from the quartet state, the population of μd atoms in each hyperfine state must be well known for correct interpretation of the results. The $dd\mu$ fusion reaction produces a 2.45 MeV neutron background, complicating the already difficult task of measuring the large energy range of neutrons created in the three body final state.

Initial experiments used a primarily hydrogen target, with small concentrations of deuterium to suppress $dd\mu$ formation. However, this results in $pd\mu$ formation and a buildup of μ^3He , so the later experiments switched to pure deuterium. In 1983, a temperature dependence of the $dd\mu$ formation rate from each hyperfine state was discovered [30], which implied that the 1973 measurement needed to be increased by a factor of three. The 1986 experiment utilized the ‘lifetime method’, extracting the capture rate from the small deviation of the observed μ^- disappearance rate relative to the μ^+ decay rate. The low stop fraction within the low density target of the 1989 experiment resulted in backgrounds which dominated the physical signal.

The MuSun experiment aims to measure the muon capture rate on deuterium to 1.5% precision using an ultra-pure gaseous deuterium time projection chamber (TPC) in combi-

nation with the lifetime method. An active target allows for 3D stop reconstruction of each muon, bestowing the ability to reject events which leave the target gas and suppress backgrounds which significantly contributed to the earlier measurements. Target conditions have been carefully selected to suppress $dd\mu$ formation effects and cleanly interpret the hyperfine state populations. The lifetime method is utilized to eliminate complications of measuring capture neutrons directly. The full statistics necessary to achieve the precision goal have been collected in two separate beam periods, the first in 2014 (R2014) and the second in 2015 (R2015). This thesis presents several significant upgrades to the centerpiece of the experiment – the TPC – as well as the data collection and analysis of the R2014 dataset.

Chapter 2 details the evolution of modern nuclear theories via muon physics and describes the motivation for the measurement described in this thesis. Chapters 3 and 4 describe several challenges and the design considerations used to combat them in the experimental setup. Significant upgrades to the centerpiece of the experimental apparatus, the cryogenic TPC, are discussed in chapter 5. Chapter 6 summarizes the R2014 data collection period, and all remaining chapters of this thesis are details on the analysis. An overview of the software is given in chapter 7, and a summary of all systematic errors is presented in chapter 8. Two particular sources of systematic error are studied in detail in this thesis, relating to muon and electron backgrounds and muon-catalyzed fusion effects, in chapters 9 and 10. Several important consistency checks are considered in chapter 11 and a summary of the analysis is presented in chapter 12.

Chapter 2

MOTIVATION AND THEORY

The muon capture process studied in this thesis lies at the interface of two fundamental forces – the weakly interacting muon probes the nuclear bound QCD states. In this chapter, the weak interaction is introduced and illustrated via the purely leptonic process of muon decay. The inclusion of hadrons is then investigated via semileptonic processes. Symmetries of the QCD Lagrangian are explored in the context of developing low energy effective field theories which can describe nucleon and even few body interactions. The implications of calibrating this model independent, predictive theory for astrophysical reactions of interest are discussed in the final section.

2.1 Weak interaction and the Fermi Constant

A theoretical description of the weak interaction was first proposed by Fermi in close analogy with quantum electrodynamics to explain the neutron beta decay process

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

His initial formulation consisted of a current-current interaction described by

$$\mathcal{L} = G_F \bar{\Psi}_p \gamma_\alpha \Psi_n \bar{\Psi}_e \gamma^\alpha \Psi_{\nu_e} \quad (2.2)$$

where Ψ represent field operators. Although somewhat successful, this description did not explain the observation of Gamow-Teller transitions, in which the total change in angular momentum between final and initial states of a reaction takes values of $\Delta J = 0, \pm 1$. In 1956 the possibility of parity violation was proposed by Lee and Yang, indicating that the current

of equation 2.2 may actually include several other Lorentz invariant combinations:

$$\begin{aligned}
\bar{\Psi}\Psi & \quad \text{Scalar, (S)} \\
\bar{\Psi}\gamma^5\Psi & \quad \text{Pseudoscalar, (P)} \\
\bar{\Psi}\gamma^\alpha\Psi & \quad \text{Vector (V)} \\
\bar{\Psi}\gamma^\alpha\gamma^5\Psi & \quad \text{Axial Vector (A)} \\
\bar{\Psi}\sigma_{\alpha\beta}\Psi & \quad \text{Tensor (T)}.
\end{aligned}
\tag{2.3}$$

A cascade of illuminating experiments followed, in which the V-A structure of the weak interaction was deciphered [20]. In 1961, Glasgow proposed a gauge theory which unifies the electromagnetic and weak forces. The weak interaction is thus not weak due to a reduced coupling, but because it is mediated by massive vector bosons, with a propagator given by

$$-\frac{g_{\mu\nu} - (1/m_W^2)q_\mu q_\nu}{(q^2 - m_W^2)}
\tag{2.4}$$

where g is the weak coupling which can be related to the Fermi constant above by comparing at low momenta, $q \ll m_W$,

$$\frac{G_F}{\sqrt{2}} = \frac{g^2}{8m_W^2}.
\tag{2.5}$$

However, this theory was only gauge invariant for massless vector bosons. The inclusion of massive vector bosons was remedied independently by Weinberg and Salam several years later, via a mechanism of spontaneous symmetry breaking, in which the underlying theory of the Lagrangian is broken by the ground state. This standard model of electroweak interactions is completed with the Higgs mechanism, which results in massive gauge bosons and a massive scalar particle, the Higgs particle, which was finally observed at the LHC in 2012.

As a purely leptonic process, the measurement of the muon decay rate, illustrated in figure 2.1(a) is ideal for characterizing the strength of the weak interaction. In the low energy limit, where $q^2 \ll M^2$, the Lagrangian can be written as an effective current-current interaction as shown in figure 2.1(b) with a strength determined by the Fermi constant, G_F

$$\mathcal{L} = \frac{G_F}{\sqrt{2}} L^\alpha L_\alpha^\dagger.
\tag{2.6}$$

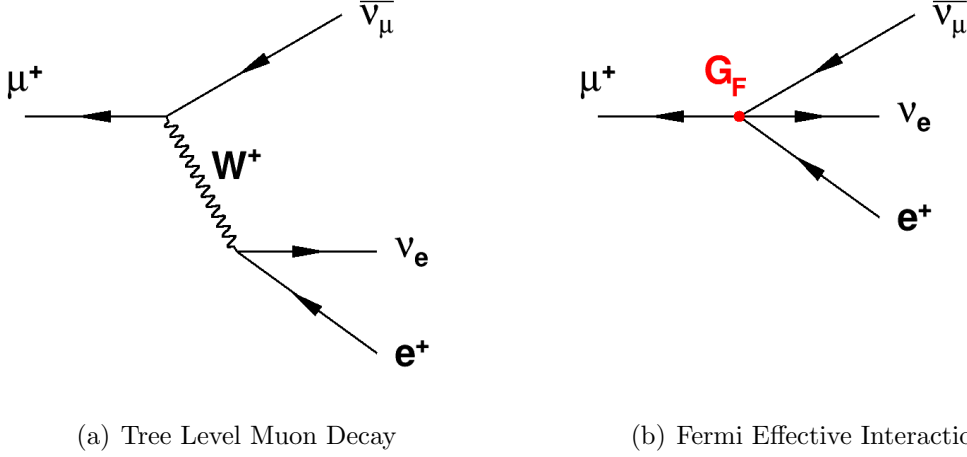


Figure 2.1: (a) Feynman diagram representing the tree level decay of the positively charged muon. As the momentum transfer is significantly smaller than the mass of the W boson, the decay can be reduced to a four point Fermi interaction, as shown in (b).

The leptonic current is given by

$$L^\alpha = \bar{\Psi}_e \gamma^\alpha (1 - \gamma^5) \Psi_{\nu_e} + \bar{\Psi}_\mu \gamma^\alpha (1 - \gamma^5) \Psi_{\nu_\mu} \quad (2.7)$$

where Ψ_i represent the field operators of the four leptons of the incoming and outgoing states. In order to calculate the decay rate, the invariant amplitude must first be determined. This is given by

$$M = \frac{G_F}{\sqrt{2}} \bar{u}_{\nu_\mu} \gamma_\alpha (1 - \gamma^5) u_\mu \bar{u}_e \gamma^\alpha (1 - \gamma^5) u_{\nu_e} \quad (2.8)$$

where the field operators have acted on the initial and final states, giving rise to an expression written in terms of Dirac spinors, u_i . The application of Fermi's golden rule yields for the differential transition probability

$$d\Gamma = |M|^2 (2\pi)^2 \delta^4(p_e + p_{\nu_e} + p_{\nu_\mu} - p_\mu) \prod_{k=\mu, e, \nu_\mu, \nu_e} \frac{m_k}{E_k} \prod_{f=e, \nu_\mu, \nu_e} \frac{d^3 p_f}{(2\pi)^3} \quad (2.9)$$

By integrating over the outgoing particle momenta, the muon lifetime from the tree level

diagram is obtained, and related to the Fermi constant via

$$\frac{1}{\tau_\mu} = \frac{G_F^2 m_\mu^5}{192\pi^3}. \quad (2.10)$$

All corrections to this result are due to radiative corrections, many of which were calculated in 1999 by van Ritbergen and Stuart, resulting in a sub-ppm relation between G_F and the muon lifetime [47]. In 2013, the MuLan experiment measured the positive muon lifetime to the ppm level to be $\tau_\mu = 2196980.3(2.2) ps$, allowing for an extraction of the Fermi coupling constant to 0.5 ppm [32].

Aside from being an important test of the standard model, both muon capture experiments discussed later in this chapter rely on the muon lifetime to high precision in order to obtain capture rates via measurements of the small deviation of the muon decay rate due to the presence of an additional disappearance term.

2.2 *Single Nucleon Sector and the Pseudoscalar Coupling*

Muon capture processes on light nuclei offer a unique realm in which observables can be both calculated by modern theories based on the underlying symmetries of QCD and measured experimentally. Muon capture on the proton, shown in figure 2.2, offers insight into the theoretical development of the weak current in the single nucleon sector. In this case, the




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

interaction is characterized by the following Lagrangian

$$\mathcal{L} = \frac{G_F V_{ud}}{\sqrt{2}} (J^\alpha L_\alpha + h.c) \quad (2.11)$$

where V_{ud} is the CKM matrix element which quantifies quark flavor mixing, L_α is the leptonic current described in the previous section, and J^α is the hadronic weak current, written in terms of the field operators as

$$\bar{\Psi}_d \gamma^\alpha (1 - \gamma^5) \Psi_u. \quad (2.12)$$

In order to calculate the transition probability, the invariant amplitude must be obtained. The leptonic contribution to the invariant amplitude is the same as in muon decay

$$\langle \bar{\nu}_\mu | L_\alpha | \mu \rangle = \bar{u}_\nu \gamma^\alpha (1 - \gamma^5) u_\mu. \quad (2.13)$$

The expression for the hadronic current contribution could be determined just as easily if the initial and final states contained free quarks. However, the initial and final states instead contain the bound quark states of the proton and neutron,

$$\langle n | J^\alpha | p \rangle = \langle n | \bar{\Psi}_d \gamma^\alpha (1 - \gamma^5) \Psi_u | p \rangle \quad (2.14)$$

and as it is challenging to calculate how the quark field operators act on a bound state of quarks. Instead, the matrix elements can be written in terms of form factors, including the most general form of all Lorentz invariant terms which may contribute,

$$\begin{aligned} \langle n | V^\alpha | p \rangle &= \bar{u}_n \left[F_V(q^2) \gamma^\alpha + \frac{i F_M(q^2)}{2m_N} \sigma^{\alpha\beta} q_\beta + \frac{F_S(q^2)}{m_\mu} q^\alpha \right] u_p \\ \langle n | A^\alpha | p \rangle &= \bar{u}_n \left[F_A(q^2) \gamma^\alpha \gamma^5 + \frac{i F_T(q^2)}{2m_N} \sigma^{\alpha\beta} q_\beta \gamma^5 + \frac{F_P(q^2)}{m_\mu} q^\alpha \gamma^5 \right] u_p \end{aligned} \quad (2.15)$$

where $q^\alpha = p_n^\alpha - p_p^\alpha$ is the four momentum transfer and $F_i(q^2)$ are the weak form factors encapsulating the strong interactions effects within the bound nucleon states [28]. The scalar term transforms differently under G-parity – defined by a rotation in isospin space followed by charge conjugation – than the other terms in the vector current and thus F_S is expected to be

nearly zero. Likewise, F_T is expected to be extremely small because it transforms differently under G-Parity than the axial and pseudoscalar terms. The conserved vector current (CVC) hypothesis of the standard model states that the vector hadronic current above, and its complex conjugate, form an isospin triplet with the isovector electromagnetic current [20], and that all three currents are conserved. The electroweak vector and magnetic form factors, F_V and F_M are then easily determined in terms of the well known electromagnetic form factors. The axial form factor, F_A has been determined via measurement of the neutron decay rate. This leaves one unknown coupling constant, the pseudoscalar F_P . Initially, values of F_P were calculated based on the partially conserved axial current hypothesis (PCAC), which relates the divergence of the axial current to the pion field,

$$\partial_\mu A^\mu(x) \propto \phi(x). \quad (2.16)$$

Expansion of the divergence of the axial current around a momentum transfer of zero results in the following first order term

$$F_P(q^2) = \frac{2m_\mu m_N}{m_\pi^2 - q^2} F_A(0). \quad (2.17)$$

A precision measurement of F_P was required to assert the validity of this assumption, and strengthen the understanding of the weak and strong interactions. In 2013, F_p was obtained via a 1% measurement of the muon capture rate on the proton by the MuCap collaboration [6]. Additionally, with the development of chiral perturbation theory, discussed in a later section, the estimate in equation 2.17 can be systematically improved [12] to give a precise QCD prediction

$$F_P^{chiral}(q^2) = \frac{2m_\mu g_{\pi NN} F_\pi}{m_\pi^2 - q^2} - \frac{1}{3} F_A(0) m_\mu m_N r_A^2. \quad (2.18)$$

Even if the current uncertainty in the nucleon axial radius is taken into account, the measured value of F_P from the MuCap experiment confirms the theoretical prediction to 8%, serving as an important confirmation of this chiral symmetry framework [29].

2.3 Few Nucleon Interactions

In extending to more complex nuclei, exchange currents between nucleons must be taken into account in addition to the one-body parameterization developed in the previous section. In order to determine the uncertainty in a calculation, there must exist some way to order such terms based on their contribution. Typically, this done by writing the theory as a combination of the free Hamiltonian and a perturbative interaction term,

$$H = H_{free} + H_{int} \tag{2.19}$$

and expanding in terms of some small parameter. In the case of QCD in the high energy regime, this can be done effectively in terms of the small running coupling constant, α_s . Increasingly higher terms contribute less, and due to the well defined power counting scheme, the error associated in truncating at each order is quantified. However, this does not work in the regime of few body nuclear processes, due to the increasingly large α_s at low energies.

The non-perturbative nature of QCD at low energies has stimulated the development of modern effective field theories (EFTs), in which an effective Lagrangian is created by including all terms permitted by the underlying symmetries of the theory. In this case, the theory is mathematically equivalent to the full theory, but may be expanded in terms of some small parameter which allows for a power counting scheme. This results in a model-independent, predictive theory, in which the truncation errors of each calculation are quantifiable. In order to limit the terms necessary in the expansion, a large separation of energy scales is exploited. The physics beyond the relevant scale is integrated out, and exists within the remaining theory in the form of low energy constants (LECs) leaving only the relevant degrees of freedom. There has been significant progress in calculating these constants from first principles on the lattice, e.g. the nucleon axial coupling, g_A has been determined to the percent level [19]. A complimentary program of experiments aims to determine these LECs empirically. Two EFTs are discussed in the context of muon capture, one which uses the chiral symmetry

breaking scale and a second which uses the pion mass as relevant energy scale.

2.3.1 QCD Symmetries

There are two important symmetries which contribute to the hierarchy of particles observed in nature; isospin and chiral symmetries. Consider the simplified two quark-flavor case, with a doublet field given by

$$\Psi = \begin{pmatrix} u \\ d \end{pmatrix}. \quad (2.20)$$

A rotation in isospin space is given by

$$\Psi \rightarrow \Psi' = \Omega\Psi \quad (2.21)$$

where Ω is as a unitary 2×2 matrix. This generates three vector currents, and the divergence of the charged current is given by

$$\partial_\alpha V^\alpha = \partial_\alpha(\bar{u}\gamma^\alpha d) = i(m_u - m_d)\bar{u}d \quad (2.22)$$

indicating the current is conserved in the limit where the quark masses are equal, as is nearly the case in QCD. This was exploited in section 2.2 in order to determine the vector and magnetic form factors for single nucleon capture.

The same transformation can be applied to the left and right handed components of the doublet field,

$$\begin{aligned} \Psi_L &\rightarrow \Psi'_L = \Omega_L\Psi_L \\ \Psi_R &\rightarrow \Psi'_R = \Omega_R\Psi_R. \end{aligned} \quad (2.23)$$

This once again generates three currents, and the divergence of the charged axial current is given by

$$\partial_\alpha A^\alpha = \partial_\alpha(\bar{u}\gamma^\alpha\gamma_5 d) = i(m_u + m_d)\bar{u}\gamma_5 d \quad (2.24)$$

indicating the axial current is only conserved in the case where both quark masses are zero. The partial conservation of the axial current was used to approximate the pseudoscalar coupling in equation 2.17. Initially the divergence was equated with the pion field because it happened to exhibit the same quantum numbers. As it turns out, this is no accident. Chiral symmetry is a spontaneously broken symmetry of QCD – the Lagrangian possesses a symmetry which is broken by the vacuum state. As formulated by Nambu and Goldstone in 1960 [39, 27], a spontaneously broken symmetry leads to the existence of massless Goldstone bosons, which in the case of QCD gives rise to the three lightest mesons, the pions. Due to the finite quark masses, this chiral symmetry is broken explicitly as well by a small amount, and the pions, albeit light, are not massless. Exploiting this approximate symmetry, the QCD Hamiltonian can be written

$$H = H_0 + H_{sb} \tag{2.25}$$

where H_{sb} contains the symmetry breaking mass terms of the u and d quarks, while H_0 contains all terms which respect the chiral symmetry. As the quark masses are relatively small, $m_u = 2.2_{-0.4}^{+0.6} \text{ MeV}$ and $m_d = 4.7_{-0.4}^{+0.5} \text{ MeV}$ [45], this allows for a perturbative expansion. This is the basis of Chiral Perturbation Theory (χ PT).

2.3.2 Chiral Perturbation Theory

In χ PT, an effective Lagrangian is constructed which retains the approximate chiral symmetry of QCD, and is organized into powers of the momentum over the chiral symmetry breaking scale p/Λ_χ . All physics above $\Lambda_\chi \approx 1 \text{ GeV}$ is integrated out and contained within low energy constants, leaving nucleons and pions as the relevant degrees of freedom. Only one complication remains in this case. Terms arise proportional to the nucleon energy, which if treated relativistically is of the scale of the nucleon mass $M_N \approx \Lambda_\chi$. In order to evade these terms, the baryons are considered heavy, static sources with small momentum transfer between pions [33]. This is referred to as heavy baryon chiral perturbation theory. Within

such a theory, the coupling of external currents to the nuclear system can be calculated in a self consistent manner. The focus of this thesis is the coupling of the weak current to the two nucleon system, described by the first order diagrams of figure 2.3. The two nucleon sector contains a single new LEC, d_R , which quantifies a four point nucleon interaction. Currently, the LEC d_R has not been measured cleanly in a two nucleon system.




Figure 2.3: Feynman diagrams which contribute to the two nucleon axial vector coupling at leading order. The solid lines represent nucleons, while the dashed line represents a pion. The contact interaction is characterized by a single LEC, d_R .

If calculated up to next-to-next-to leading order (N^2LO), the couplings can be related to the force terms in the three-nucleon system, as shown in figure 2.4. All contributing diagrams for one and two nucleon interactions up to N^3LO can be seen in figure 2.5 [33]. The LECs c_1 , c_2 , and c_3 enter at next-to-leading order (NLO) in the two pion exchange diagrams, while c_D and c_E first enter at N^2LO . The relation between d_R and c_D is given by

$$d_R = -\frac{1}{4} \frac{m_N}{g_a \Lambda_\chi} c_D + \frac{m_N}{3} (c_3 + 2c_4) + \frac{1}{6} \quad (2.26)$$

indicating there are three LECs which must be determined in order to extract d_R from the three nucleon system. The four LECs, c_i , which characterize πNN vertices at NLO are determined by fitting to data from elastic pion-nucleon scattering experiments[16, 23]. The constants c_D and c_E are most precisely determined using the binding energy of A=3 nuclei




Figure 2.4: Feynman diagrams which contribute to the 3N forces at N^2LO . The solid lines represent nucleons, while the dashed line represents a pion. The LEC c_D can be related to d_R .

and the tritium half life [26]. The extraction of d_R from the three nucleon system relies on the validity of the theory as well as the ability to calculate the N^3LO contributions to a high degree of precision. As it is the only LEC on the two nucleon system, and only relies on calculations at leading order, it can be determined much more cleanly from a measurement with only two nucleons. A precise measurement of muon capture on deuterium thus offers a clean determination of d_R as well as a check on the current development into the three nucleon sector.

Several recent theoretical predictions of the muon capture rate on deuterium from the doublet state, which use the value of d_R from the three nucleon system, are compared to past experimental measurements in figure 2.6. The first predictions of Marcucci in 2011 use χPT to obtain the axial current, but used models to obtain the nuclear wavefunctions [37]. The following predictions of both Marcucci and Adam in 2012 use the χPT framework to obtain both the currents and the nuclear wavefunctions [35, 2]. In the 2009 paper relating the two and three nucleon LECs, the factor of $-1/4$ in front of the first term of equation 2.26 was missing [26]. This error was propagated in all subsequent calculations, and was only recently corrected in an erratum by Marcucci in 2018 [36], as displayed in the final




Figure 2.5: Contributing diagrams for two and three nucleon interactions up to N^3LO . Image taken from [33], the state of the art calculation in 2011. For a recent update with higher orders included, see [34].

theoretical prediction in figure 2.6. As noted in the erratum, this correction has a small effect on the conclusions for the 2012 results. In the meantime, both the chiral NN force and their associated current has been expanded to higher order. In these updated calculations, extraction of c_D and c_E using the nd scattering length leads to an overestimate the tritium Gamow-Teller(GT) matrix element by a few percent. An updated calculation [11], in which the tritium GT matrix element is used in place of the nd scattering length, obtains values for c_D and c_E which are significantly different. Moreover, N4LO terms in the weak current

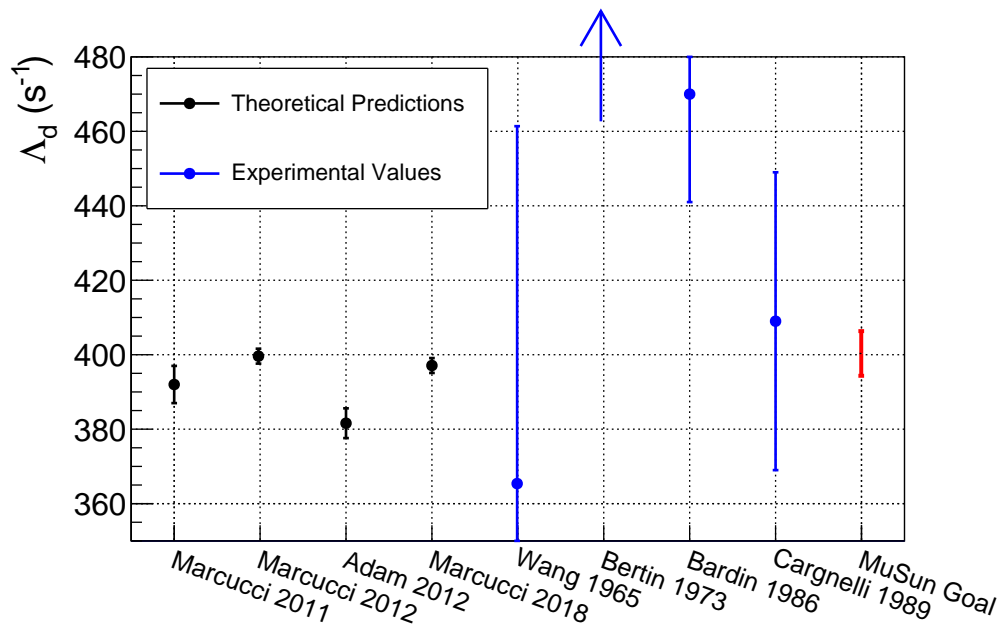


Figure 2.6: Theoretical predictions and past experimental measurements of the muon capture rate on deuterium. Theoretical predictions use the values of d_R from the three nucleon system. The results of the Bertin experiment are off the scale of this figure, and largely inconclusive due to undetermined hyperfine state populations.

have been calculated for the first time [11], and a large contribution calls into question the convergence of the theory. The impact of these recent findings on the extraction of d_R and calculations in the two nucleon sector is yet to be determined. These recent issues highlight the importance of a complimentary experimental extraction of d_R in the clean two nucleon system.

A precise enough experimental determination of the the muon capture rate, however, is yet to be provided. Figure 2.6 illustrates results from previous experiments which vary wildly, with large errors, and the most precise of the experimental values differs from the theoretical predictions by 2.9 sigma. The precision goal of the MuSun experiment, displayed as a red bar in the figure, aims to illuminate this inconclusive picture.

2.3.3 Pionless Effective Field Theory

In 1996, Kaplan, Savage, and Wise (KSW) suggested that the previous expansions in terms of Λ_χ suffered renormalization issues because of the non-perturbative nature of the NN force at low energy [31]. KSW instead proposed a pionless effective field theory (π EFT), in which the pion is integrated out as a high energy degree of freedom, leaving only nucleons as the relevant constituents of the theory. In this case, the theory gives rise to a well-defined, explicit power counting scheme in terms of the momentum over the pion mass, p/m_π . In the two nucleon sector, the axial vector coupling gives rise to a four point contact interaction, quantified by a single LEC, L_{1A} , as depicted in figure 2.7. Although this provides a concise




Figure 2.7: First order diagram contributing to NN axial vector coupling in pionless effective field theory.

theory where all interactions are contact interactions, it is only applicable in regimes where the energy is much lower than the pion mass of 140 MeV. The total 102.1 MeV of energy in the final state of μd capture is shared between the two neutrons and the neutrino, allowing for a large range of momentum transfer to the two body system, as described in figure 2.8. The region where π EFT applies contributes to roughly 90% of the total capture rate. The remainder of the calculation would need to be extended into higher energies via χ PT, or an experiment which only measures the relevant part of the phase space would need to be

completed.




Figure 2.8: Final state energy of the two neutrons and neutrino in $\mu^- + d \rightarrow n + n + \nu_\mu$. The energy (E_n) of each neutron is shown on the X and Y axes, and the diagonal lines designate constant neutrino momentum (p_ν) in units of MeV/c. Pionless effect field theory is applicable in the region where $p_\nu \geq 90 \text{ MeV}/c$. At lower neutrino momenta, meson exchange currents dominate over an impulse approximation.

2.3.4 Astrophysical Significance

Once the axial two body current is quantified via the determination of d_R , other two body systems can be calculated based on the input. Important examples of astrophysical interest include the solar proton-proton (pp) fusion reaction and the charged (CC) and neutral (NC) current reactions used by the SNO collaboration to determine the 8B solar neutrino flux, shown diagrammatically in figure 2.9. All three reactions have small enough cross sections that they cannot be determined experimentally at terrestrial conditions. As p-p fusion is the




Figure 2.9: Astrophysical reactions of interest characterized by the same LEC, d_R – (a) solar pp fusion and the (b) charged and (c) neutral current neutrino breakup reactions used by the SNO collaboration to calculate the 8B solar flux.

first step in a chain reaction which fuels hydrogen burning stars, its reaction rate is a critical component of stellar models. This rate is utilized to characterize the evolution of stars, and deviations of its value could impact age determination and neutrino fluxes of stellar clusters. The standard characterization of the pp reaction rate is the S-Factor – which is proportional to the cross section but eliminates energy dependencies arising from the Coulomb barrier, and recent calculations of the S-Factor using χPT obtain values about 3% higher compared to the literature values. The impact of the 3% change in S-Factor is found to result in a change of less than 1% in age determination and 8% in neutrino flux predictions [46].

The SNO experiment observed both charged and neutral current reactions from solar neutrinos within a large heavy water target [3]. The neutral current reaction, which is sensitive to all neutrino flavors, yields the total solar neutrino flux. A comparison of the charged current, which is only sensitive to electron neutrinos, to the neutral current gives direct evidence for neutrino oscillations. In order to extract the flux from the measured event rate, theoretical calculations of the cross sections are used. The standard SNO analysis uses the cross sections determined in the χPT framework based on the LEC obtained in the 3N system from tritium beta decay, and assumes a 1% uncertainty.

Both the p-p fusion and SNO cross sections referenced in the preceding discussion rely on the extraction of d_R from the three nucleon system. If a clean measurement of d_R in the two nucleon system confirms the validity of this extraction, the results of the previous solar models and the SNO measurements would remain unchanged. However, significant deviation in the value of d_R from the clean 2N measurement could have important implications.

2.4 Summary

From characterizing the the strength of the weak interaction to playing a significant role in the development of understanding low energy nuclear processes, muon physics has played an important role. Modern effective field theories based on the fundamental symmetries of QCD allow for model-independent predictions of few nucleon interactions, including the cross sections of several astrophysical reaction of interest. The strength of the axial two-body current, however, has not been cleanly determined in the two nucleon sector. Current predictions rely upon the determination of the LEC from the more complicated three nucleon system. The MuSun determination of the muon capture rate on deuterium will enable the extraction of this important LEC, reducing its uncertainty from 100% to 20% in the clean two nucleon system, and providing a consistency check on the chiral perturbation theory framework.

Chapter 3

DESIGN SOLUTIONS FOR EXPERIMENTAL CHALLENGES

The requirement for high precision, in addition to complications arising from muon atomic physics and muon-catalyzed fusion reactions in deuterium gas, create several challenges to the unambiguous extraction of the desired capture rate. These challenges are first mitigated with novel experimental design, followed by scrupulous analysis techniques. In this chapter, the most pressing obstacles are presented, along with the experimental techniques or conditions which have been carefully selected to suppress their effect on the final measurement.

3.1 Precision

The process of muon capture on deuterium is a rare process given by



offering only neutral particles in the final state. Although the rate could most directly be determined by measuring the neutrons, neutron detectors are difficult to calibrate, and the three-body final state leads to a continuous range of neutron energies. This method is further complicated by a background of 2.45 MeV neutrons produced by muon catalyzed fusion reactions. The large uncertainties of previous muon capture experiments were due in part to the difficulties arising from determination of the rate via direct measurement of the capture neutrons. A different method is therefore required to obtain the desired precision.

The MuSun experiment instead utilizes a technique dubbed the 'lifetime method', first used by the 1981 Saclay experiment [8], in which the difference between the measured negative and positive muon decay rates are compared to extract the capture rate. Both positively

and negatively charged muons will undergo the ordinary muon decay processes

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu \quad (3.2)$$

and

$$\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu. \quad (3.3)$$

According to the CPT theorem, the positive and negative free muon decay rates are identical¹, and are denoted here by λ_{μ^+} . When stopped in matter, the negatively charged muon can undergo an additional capture process at a rate Λ , whereas the positively charged muon cannot overcome the Coulomb potential barrier necessary to undergo capture. This additional disappearance channel for the negatively charged muon within a material manifests itself as a deviation from the free muon decay rate.

The MuSun experiment precisely measures the negative muon disappearance rate, λ_{μ^-} , in deuterium gas. This total rate includes both the nominal muon decay and the capture process, such that the observed time distribution is given by

$$\begin{aligned} N(t) &\propto e^{-(\lambda_{\mu^-})t} \\ &\propto e^{-(\lambda_{\mu^+} + \Lambda_d)t}. \end{aligned} \quad (3.4)$$

The free positive muon decay rate has been measured to the ppm level by the MuLan experiment [49], such that the capture rate can be extracted via

$$\Lambda_d = \lambda_{\mu^-} - \lambda_{\mu^+}. \quad (3.5)$$

This method eliminates the uncertainties from counting the capture neutrons. However, as the muon capture is a 10^{-3} deviation on the nominal muon decay rate, high statistics are

¹In the case of muonic atom formation, there is a small deviation in the decay rate due to relativistic effects for muon decay from an atomically bound state [21].

necessary to reach the precision goal of the MuSun experiment. The free muon decay rate is $455.1702 \times 10^3 \text{ s}^{-1}$, compared to the doublet muon capture rate of 400 s^{-1} , such that a 6 s^{-1} measurement of Λ_d requires the disappearance be measured to the 10^{-5} level. Assuming an even split of 4 s^{-1} statistical and 4 s^{-1} systematic uncertainty, $\sim 10^{10}$ decay events are required to reach this precision goal. This thesis presents the analysis of the first production run, consisting of 6×10^9 decay events.

3.2 Muon Kinetics

Several experimental challenges arise from the array of processes possible following the atomic capture of a muon within deuterium gas, as illustrated in Fig. 3.1. Once stopped in the target gas, the muon captures into one of two hyperfine states, the $F=3/2$ quartet state and $F=1/2$ doublet state with initial populations $q=2/3$ and $(1-q)$, respectively. Nuclear capture is possible from either hyperfine state, with rates denoted by Λ_q and Λ_d . Formation of $dd\mu$ molecules is also possible from either hyperfine state, with rates λ_q and λ_d . Once $dd\mu$ molecules are formed, a muon catalyzed fusion occurs in less than a nanosecond, producing either a ${}^3\text{He}$ and a neutron with branching ratio β or a proton and triton with $1 - \beta$. The quantity $\bar{\omega}$ denotes the probability that the muon will stick to the ${}^3\text{He}$ following a fusion, which can result in a nuclear capture process. Otherwise, the muon is recycled after the fusion. Each process is discussed further in the subsequent sections, and the most recent values for the rates and branching ratios can be found in table 3.1 [7]. In order to reach an unambiguous determination of the capture rate, the running conditions must be optimized to suppress the other processes which can interfere with the desired measurement.

3.2.1 Atomic and Nuclear Capture

A muon entering the target first slows and undergoes atomic capture on deuterium into an excited state. Through the Auger effect, radiative transitions, and Coulomb collisions the

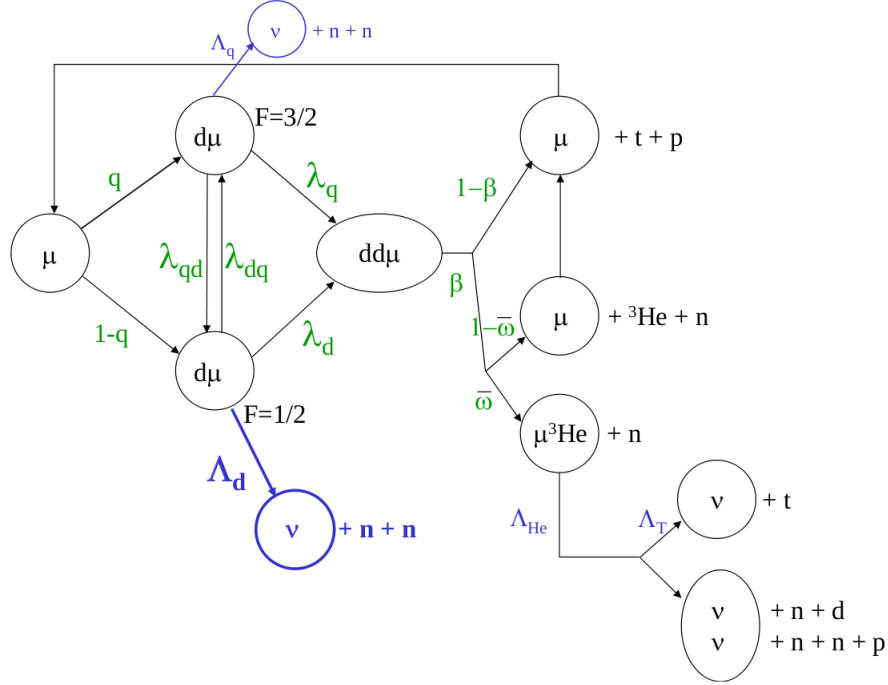


Figure 3.1: Simplified diagram of Muon Kinetics within deuterium gas.

muonic atom is quickly de-excited to the 1S ground state into one of two hyperfine states, the doublet ($F=1/2$) and the quartet ($F=3/2$), initially populated according to their statistical weights. Due to the helicity structure of the weak interaction, capture from the quartet state is suppressed, resulting in a capture rate roughly 40 times higher from the doublet state, Λ_d , compared to the quartet state, Λ_q . The hyperfine state upon capture must therefore be known well in order to correctly interpret the MuSun measurement.

The following hyperfine transition exists between the quartet and doublet state

$$\mu d(\uparrow\uparrow) + d \rightarrow d + \mu d(\downarrow\downarrow) + 0.049 \text{ eV}. \quad (3.6)$$

Being a collisional de-excitation, this hyperfine transition can be increased by selecting higher target densities. The reverse transition is suppressed by the energy difference of 49 meV

Process	Symbol	Value		Ref
		300k	30k	
μd Quartet Capture Rate	Λ_q	~ 10		
μd Doublet Capture Rate	Λ_d	~ 400		
3He Partial Capture Rate	Λ_{He}	1496.0(40)		[1]
3He total Capture Rate	Λ_T	2216(70)		
hf transition $q \rightarrow d$	λ_{qd}	35(5)	37.0(4)	[30]
hf transition $d \rightarrow q$	λ_{dq}	10.5	0	
Quartet $dd\mu$ formation	λ_q	3.75	3.98(5)	[7]
Doublet $dd\mu$ formation	λ_d	2.549(23)	0.053(3)	
Effective Fusion Fraction	β (unit-less)	0.590(6)	0.517(15)	
Sticking Probability	$\bar{\omega}$	0.1206(6)		

Table 3.1: Rates and branching ratios for Muon kinetics.

between the two hyperfine states. The faster depopulation of the quartet state leads to a cleaner interpretation of the measured decay rate.

3.2.2 Molecular formation and Muon Catalyzed Fusion

The process of muon catalyzed fusion has been studied extensively [7]. A $dd\mu$ molecule can be formed from either hyperfine state, primarily via a temperature dependent resonant molecular formation.



This resonant process relies on a matching of the kinetic energy to a loosely bound state in which the released energy is absorbed into vibrations and rotations of the formed molecule,

leading to a strong temperature dependence, as shown in Fig. 3.2.




Figure 3.2: Molecular formation rates from both hyperfine states (left) and hyperfine transition rates at various temperatures. Figure reproduced from [7].

As the muon mass is 200 times heavier than the electron, the radius of the $dd\mu$ molecule is reduced compared to the typical D_2 radius, bringing the two deuterium nuclei within 500 fm. The probability of overcoming the Coulomb potential barrier at this separation increases such that the muon quickly catalyzes a fusion reaction, producing either a proton and a triton, or a ${}^3\text{He}$ and a neutron. The muon is always recycled after a p+t fusion and able to re-form a μd atom in either hyperfine state, so there is no direct effect on the measured rate from this reaction. However, the observation of fusion products within the detector couple to various analysis cuts leading to a rate distortion, discussed further in chapter 10. In most cases, the muon is also recycled following a ${}^3\text{He}$ fusion, but, it is possible for the muon to stick to the ${}^3\text{He}$ and undergo nuclear capture. Although it can be suppressed by low target temperatures, this extra capture channel leads to a distortion of the MuSun measurement

and must be taken into consideration.

3.2.3 *Optimizing running conditions*

The MuSun target conditions are carefully selected to mitigate complications from the diverse muon kinetics outlined above. Standard running conditions are chosen with the following objectives:

1. Maximize population of doublet hyperfine state
2. Minimize population of quartet hyperfine state
3. Minimize population of μ^3He
4. Suppress muon-catalyzed fusion products

To satisfy condition (1) and (2) above, the collisional de-excitation between the quartet and doublet hyperfine states is accelerated by choosing a high deuterium target density. Conditions (3) and (4) are achieved by selecting lower target temperatures. This reduces the molecular formation rate, directly decreasing the μ^3He population while simultaneously reducing muon recycling back into the quartet state. Additionally, at low temperatures there is a large difference in the molecular formation rates from each hyperfine state, as illustrated in Fig. 3.2. Measurement of the fusion neutron time spectrum thus yields information on the relative populations of the hyperfine states; a critical aspect to correctly interpret the observed decay rate.

The kinetics displayed in figure 3.1 can be translated into a system of coupled differential equations, as demonstrated in appendix A. From the solution, one can determine the relative populations of the quartet and doublet hyperfine states, as well as the μ^3He population, for various temperatures and densities. Relative populations for various experimental conditions can be seen in Fig. 3.3. The densities are expressed as a fraction of the density of liquid hydrogen, $N_{LH_2} = 4.25 \times 10^{22}$ atoms/cm³.

The MuSun target conditions of 31 K and 5.1 bar are selected to accommodate a density of 0.065 LH2. As shown in frame (d) of Fig.3.3 below, this configuration suppresses the quartet state to the 10^{-4} level within $5\mu s$, and the 3He population to the 10^{-3} level.

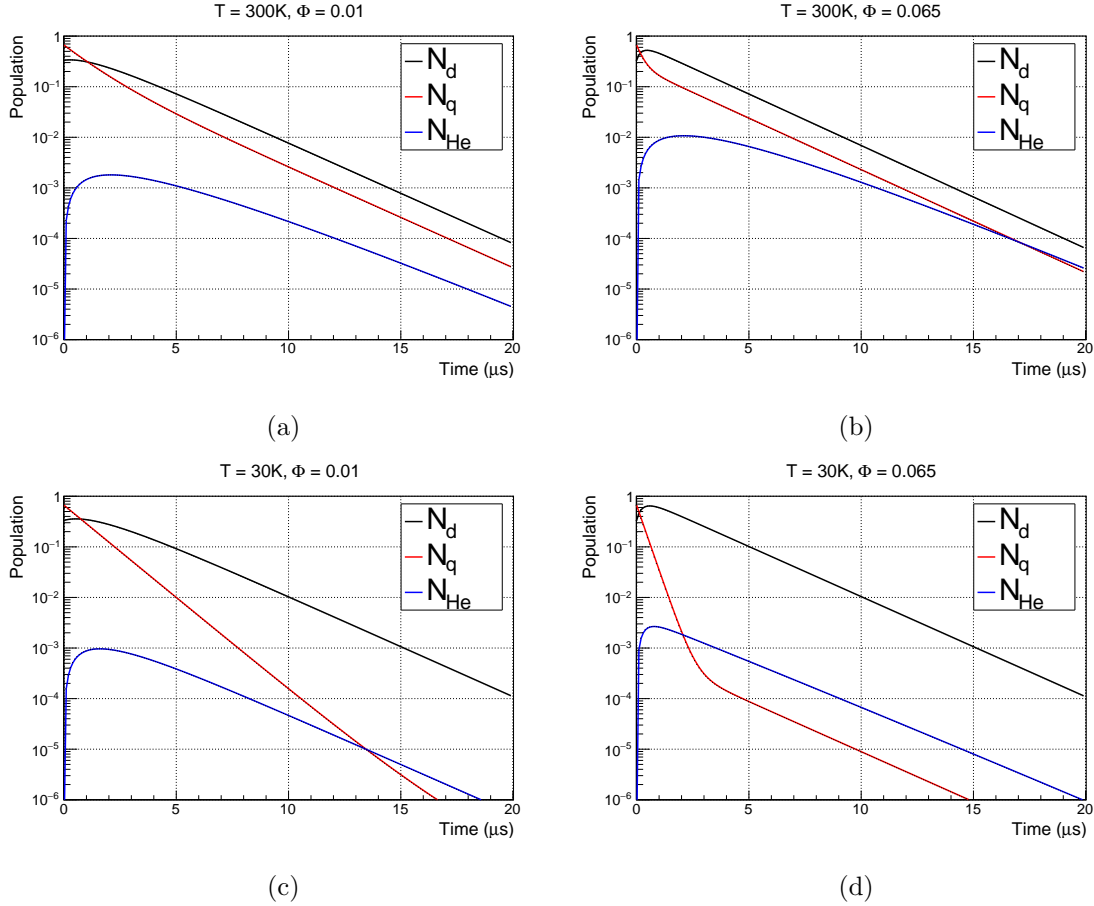


Figure 3.3: Populations of both hyperfine states and μ^3He for several variations of temperature and density.

3.3 Muon Capture On Other Nuclei

The dependence of the capture rate on a nucleus can be estimated by considering the overlap between the muon and the nucleus. The radial hydrogenic wavefunction takes the following

form

$$\psi(r) = 2\left(\frac{z}{a_0}\right)^{3/2}e^{-\left(\frac{z}{a_0}r\right)} \quad (3.8)$$

where z is the atomic number, r is the radial distance, and the Bohr radius is defined as

$$a_0 = \frac{\hbar^2}{(e^2/4\pi\epsilon_0)m}. \quad (3.9)$$

The probability of finding the muon at the nucleus is thus proportional to the cube of the atomic number

$$|\psi(r)|^2 \propto \frac{z^3}{a_0^3}. \quad (3.10)$$

Multiplying by the number of protons in the nucleus, we find the rate of muon capture scales roughly as z^4 . As the lifetime method measures a deviation on the free muon decay rate, capture on any nuclei other than deuterium will distort the observed time distribution and lead to an additional deviation. Two cases must be considered; direct capture and transfer to secondary nuclei.

For some fraction, α , of direct captures on nuclei other than deuterium, the observed electron distribution is given by

$$e(t) = N\lambda e^{-\lambda t}(1 + \alpha e^{-\Lambda_z t}) \quad (3.11)$$

where λ represents the free muon decay rate plus the muon capture rate on deuterium, and Λ_z is the capture rate on the secondary nucleus. Muons which are stopped in materials surrounding the active target are contributing sources of direct capture on other nuclei.

To consider the case of muons which transfer to a secondary nuclei, then undergo a capture process, the set of differential equations governing the kinetics must be solved taking the transfer rate to the secondary nuclei into account. This is shown in appendix B.

3.3.1 Capture on Target Materials

Muons which stop outside of the gas volume can be captured directly on the materials surrounding the target, contributing to a deviation of the observed rate illustrated in equation

3.11. To estimate the shift in the disappearance rate caused by a fraction of stops in a given material, pseudo-data was generated according to 3.11 and fit with a pure exponential. The fitted rate with a fraction of stops α is compared to the nominal decay rate with $\alpha=0$. The capture rates for a few materials comprising the detector and volume surrounding the detector, along with the fraction of stops, α_2 in each material required to cause a $2 s^{-1}$ shift in the measured decay time for two different fit start times, can be seen in table 3.2. A larger fraction of stops is permissible for higher Z materials with high capture rates, especially at later fit start times, as most muons have been captured by the time the fit starts.

Material	$\Lambda_Z (10^6/s)$	α_2 160 ns	α_2 1 μs
C	0.044	6×10^{-5}	6×10^{-5}
Fe	4.53	1×10^{-4}	5×10^{-3}
Ag	10.86	7×10^{-4}	6
W	11.92	1×10^{-3}	30

Table 3.2: Lifetime of muons stopped in various materials, and allowed stop fraction for a rate deviation of $\Delta\lambda = 2 s^{-1}$ for fit start times of 160 ns and 1 μs .

In order to essentially eliminate contributions from muons stops in wall materials, a cryogenic time projection chamber (TPC) is used to reconstruct the three dimensional stop position of each muon. Any event in which the muon is suspected to leave the gaseous volume is rejected in the final analysis. As the capture rate scales with z^4 , captures on materials with large atomic number occur quickly, allowing effects on the observed rate to be suppressed with delayed fit start times. As such, the TPC was completely redesigned in 2013 using only high Z materials. The operating principles and all upgrades to the TPC are discussed in detail in chapter 5.

3.3.2 Gas Purity

Isotopic Purity

Commercially obtained gas for the target will contain two hydrogen isotopes: 2H , the desired deuterium component, and 1H , referred to as protium. Muons in the presence of this isotopic admixture of protium and deuterium, can catalyze the following fusion reaction



The muons are ejected with 5.3 MeV, corresponding to a range of ~ 26 cm in our chosen gas density. These muons likely leave the active volume and can potentially be captured on wall materials, leading to a distortion of the measured rate. In order to mitigate its effect on the decay rate, the target gas must contain less than 100 ppm of protium. This requirement is discussed in more detail in section 8.3.1.

The isotopic purity of the gas is ensured by producing ultra-pure deuterium using a cryogenic distillation procedure, which utilizes the difference in saturation vapor pressures to separate hydrogen isotopes [4]. A simplified diagram of the distillation column is displayed in figure 3.4. A condenser at the top creates a reflux, which drains down along the column, while hydrogen in the liquid phase is vaporized via a re-boiler at the bottom and rises up the column. The column is filled with a spiral prismatic random packing designed to optimize surface area, increasing the interaction between the rising vapor and the draining reflux. The lighter component is discharged as vapor, while pure deuterium accumulates at the re-boiler and is extracted through the bottom of the column. Starting with a source gas of commercially available deuterium, which typically contains around 2000 ppm of protium, the ultra-pure deuterium extracted from the re-boiler contains less than 100 ppb of protium. A chromatography measurement of the target gas was completed several times throughout the 2014 data run, and a stable protium component of 16 ppm was measured.




Figure 3.4: Simplified diagram of the distillation column use for isotopic separation, reproduced from [4].

Chemical Purity

As the relatively small, neutral μd atoms can easily penetrate the electron cloud of heavier nuclei, and the difference in binding energies is large, muons will readily transfer to impurities such as nitrogen and oxygen within the gas. Literature values for the transfer rates from μd as well as capture rates on these impurities can be seen in table 3.3. Compared to the 400 s^{-1} capture rate on deuterium, even trace amounts of impurities impose a serious threat for distortion of the observed decay rate. The most likely impurities are constituents of air, namely water vapor, nitrogen, and oxygen, yet only the partial pressure of nitrogen is high enough at the MuSun running conditions to pose any potential systematic effect, and is the main concern. In R2013, a high statistics dataset was taken with a doped gas containing roughly 16 ppb of nitrogen impurity. From this dataset, it was determined that the observed rate shifts by 3 s^{-1} per ppb of nitrogen impurity. It is therefore imperative to retain and

Nucleus	λ_{dz} ($10^{10} s^{-1}$)	Λ_Z ($10^3 s^{-1}$)
N	14.5 ± 0.2	69
O	6.3 ± 0.5	102

Table 3.3: Muon transfer rates from deuterium to nitrogen and oxygen at thermal energies, appropriate for the MuSun conditions in which thermalization is much faster than transfer. Values taken from [22].

quantify the nitrogen impurity at the ppb level.

A cryogenic hydrogen ultra-high purification system (CHUPS) was developed for the MuCap experiment, and was used to continuously filter the MuSun target gas through zeolite filters held at liquid nitrogen temperature [25]. Two methods were used to monitor the nitrogen and oxygen contamination throughout the run. Gas samples were periodically extracted from the TPC for a chromatography measurement, with a calibrated system designed for the MuSun experiment capable of determining concentrations to the few ppb level [24]. In addition, a method utilizing the observed nuclear recoil energy deposited into the TPC from nitrogen capture events was developed as a in situ monitor of the impurity level. The results of the chromatography and in situ analysis are detailed in section 8.3.2.

Chapter 4

EXPERIMENTAL DESIGN

A cross section of the detector setup for the MuSun experiment can be seen in figure 4.1. Only events with a single muon in the detector system within a $\pm 25 \mu s$ window are selected. An upstream beamline element, referred to as the muon-on-request (MORE) and described in detail below, is used to increase the probability of this condition. As the muon proceeds downstream, a scintillator (μ SC) and a proportional chamber (μ PC) record the time and position of the muon entrance, respectively. The beam is tuned to stop muons within the cryogenic time projection chamber (TPC) filled with deuterium gas at 31 K and 6.5% LH_2 density. A three dimensional reconstruction of the muon stop position is obtained with the TPC to ensure the muon stops within the deuterium volume. Once stopped in the target, most muons undergo the decay process governed by equation 3.3, emitting an electron. Two cylindrical multiwire proportional chambers (ePCs) are used to reconstruct the electron tracks, and a hodoscope of 16 two-layer scintillators (eSCs) record the electron time. The time of the electron with respect to the muon entrance time is used to extract the muon decay rate.

Pileup protection in the entrance detectors is implemented to ensure that a single muon is in the detector upon decay. Cuts on the reconstructed muon stop position confirm that the muon stopped within the gas. Each detector system is discussed in more detail in the subsequent sections.




Figure 4.1: Cross sectional view of the MuSun detector setup.

4.1 *Blinding*

To prevent any bias in the MuSun analysis, the experimental data is blinded before collection. Blinding is achieved by de-tuning the master clocks by some unknown value, within 0.3% of the nominal clock frequency. The value of the blinding is recorded and checked each week of the experiment by an external member of the collaboration, and is unknown to all collaborators until the final analysis is completed.

4.2 *Beamline*

A continuous beam of 40 MeV/c muons with $\Delta P/P \simeq 3\%$ is provided by the Swiss muon source at the Paul Scherrer Institute. Starting with a hydrogen source, protons are stripped

and accelerated in a Cockcroft-Walton accelerator then injected into two ring stages. Protons emerging from the second stage at 590 MeV are then directed onto a carbon target, producing negative pions via



The pions then decay into a muon and a neutrino via



A series of quadrupole magnets and beam collimators are used to guide and focus the beam such that muons stop in the center of the detector system.

There are several beamline elements upstream of the main detector systems crucial to reaching the precision goals of the MuSun measurement, as outlined in figure 4.2. The

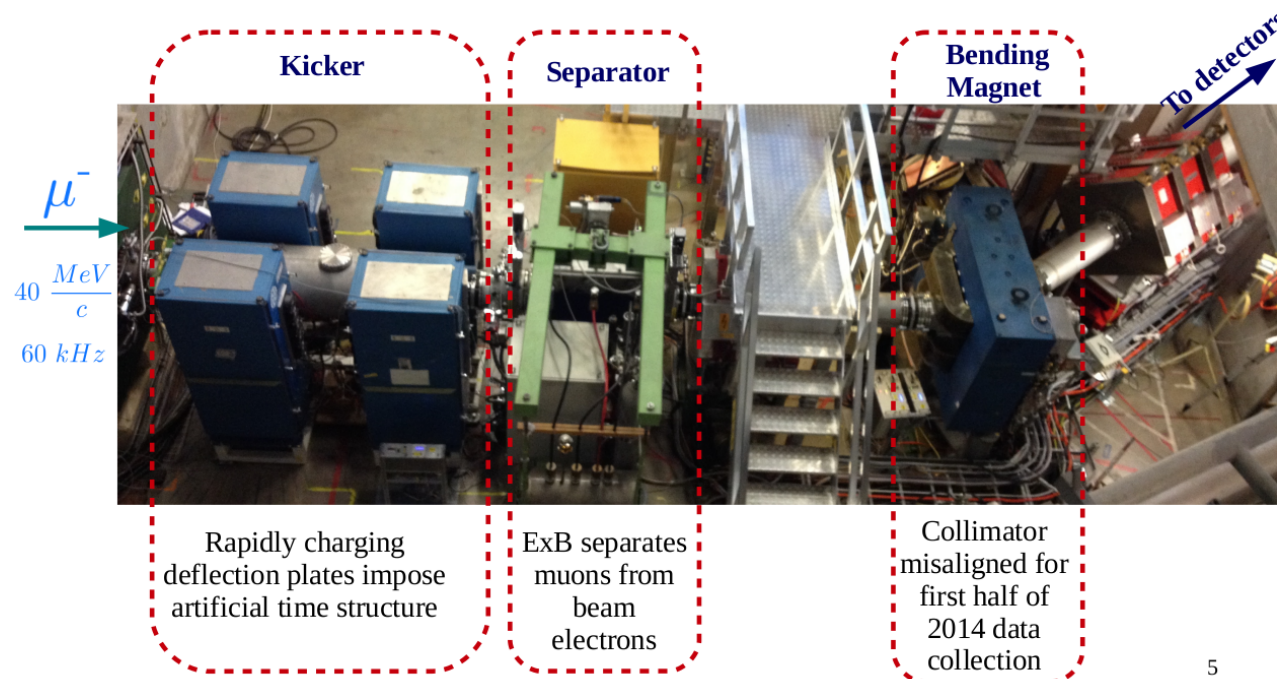


Figure 4.2: Upstream beamline elements of the MuSun Experiment.

first upstream element shown on the far left of figure 4.2 is the muon-on-request (MORE) system, also referred to as the kicker, a pair of rapidly charging plates which impose an artificial time structure to optimize muon stopping rate [10]. As the measurement technique requires that a single muon stop in the detector, it is advantageous to kick the beam and reduce the incoming muon rate once a muon has stopped in the target gas. A signal from the μ SC entrance detector is sent back to the kicker to trigger the charging of the plates to ± 12.5 kV, effectively diverting the muon beam by 22 mrad into an upstream beam collimator. Immediately following the kicker is the muon-electron separator, an ExB field intended to stop the beam electrons in the upstream collimator while allowing the muons to pass through unperturbed. The third beamline element depicted in figure 4.2 is the bending magnet, which contains the vertical slit onto which kicked muons and beam electrons are diverted. Extensive tuning was undertaken to optimize the clean muon stop rate while reducing contamination from beam electrons, and an error in the PSI database leading to an upstream quadrupole running at half of its nominal current, was found to be the culprit of many difficulties in this respect, unfortunately only after the completion of both MuSun production data collection periods. The effects on the beam envelope and tuning of this error are discussed in appendix D. In addition, roughly half of the R2014 data was taken in a configuration with a misaligned collimator, and its effects are further discussed in section 9.2.

4.3 Muon Entrance Detectors

A series of detectors, displayed in figure 4.3, are used to determine the time and quality of the muon entrance. A muon veto counter (μ SCa) is the first element proceeding the beam pipe, used to reject muon entrances off the beam axis. A thin scintillator, referred to as the μ SC, follows the veto detector, and is used as the primary timing signal for the muon entrance. Finally, a multiwire proportional chamber (μ PC) is used to determine the spatial location of each muon arrival.




Figure 4.3: Entrance detector elements of the MuSun Experiment.

Muon Veto (μSCa) The first element encountered by the muon upon exit of the beam pipe is a 5 mm thick scintillator with a 35 mm diameter hole in the center, used as a veto counter to exclude muons off the beam axis. This is coupled to a photomultiplier tube on one side, and the output is discriminated and sent to a TDC channel to record the time. A lead collimator with an identical 35 mm hole is mounted after the μSCa to further suppress off axis particles from entering the detector system. A subset of runs in the beginning of the R2014 dataset were collected with the μSCa misaligned. The effect of this misalignment on the data quality is discussed later.

Muon scintillator (μSC) Directly downstream of the muon veto counter is a 500 μm thick plastic scintillator, 70 X 70 mm^2 , coupled to a second photomultiplier tube on one

side. Further referred to as the muon scintillator (μSC) this detector is primarily used to record the time of the muon entrance, although both muons and beam electrons are resolved. The output is split and discriminated with two thresholds, a HIGH threshold for the muons and a LOW threshold for the electrons. The HIGH signal is sent to two separate TDC channels for redundancy and is the time for the muon used in creating lifetime histograms. The LOW \overline{HIGH} signal is also sent to a TDC channel to record an approximate beam electron rate, used primarily in beam tuning.

Muon Proportional Chamber (μPC) A multiwire proportional chamber used to record the muon entrance position transverse to the beam axis is mounted behind the μSCa . The μPC consists of two $50\ \mu m$ Mylar windows, separated by 35 mm and filled with a mix of argon, ethane and Freon at 1 bar. The gas is comprised of 49.9% Argon as the main fill gas, 49.9% C_2H_6 to act as the quenching gas, and 0.2% Freon. Two sets of 24 anode wires, one running vertical the other horizontal, are strung at a pitch of 2 mm. $25\ \mu m$ Mylar foils on either side of each anode plane serve as cathodes, operated at 2500 V. Each of the 48 anode wires are read out separately and discriminated to determine an X and Y entrance position of the muons over a $48\ X\ 48\ mm^2$ area.

Muon Clock (muCL) In order to study several detector backgrounds, an in-situ muon clock was set up throughout the run, as seen in figure 4.4. A 2.5 kHz clock which simulated a muon entrance and triggered the muon kick was placed in an OR with the real muon signal from the muSC. Additionally, both the real muSC signals and the muCL signals were sent to separate TDC channels such that the muCL events can be tagged and sorted out of the nominal distributions. The utility in the muon clock signal is the ability to select events in which the upstream beamline elements are triggered without a real muon in the detector systems, allowing for detailed studies of electron backgrounds and muon pileup. The muon clock events are used to determine the entrance detector efficiency and to characterize the

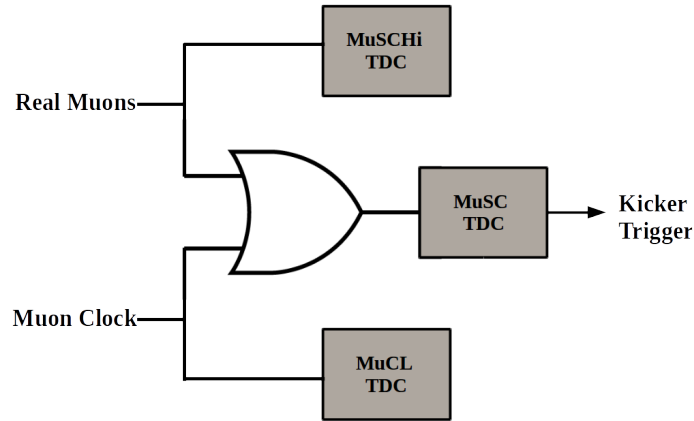


Figure 4.4: Set up of the in-situ muon clock signal

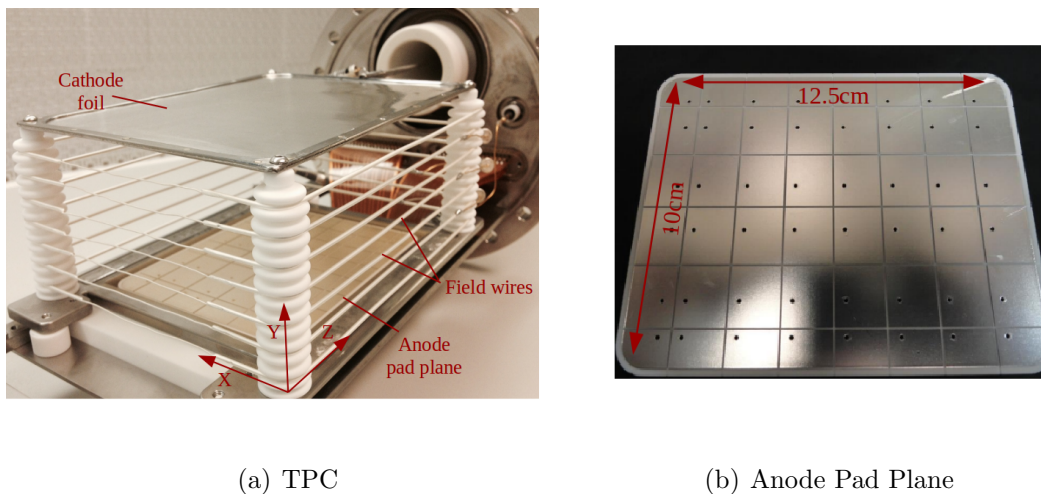
beam electron background in chapter 9.

4.4 Muon Stop Determination

4.4.1 Time Projection Chamber (TPC)

As discussed in section 3.3, it is imperative that only muons which stop within the gas volume be considered when extracting the muon lifetime, as muon captures in any materials surrounding the target can cause a distortion to the disappearance rate. A cryogenic time projection chamber (TPC) is used to reconstruct the three dimensional stop position of the muon. This chamber was significantly upgraded before the R2014 dataset. Motivations for the upgrades and a more detailed discussion of the chamber properties can be found in chapter 5. The TPC, located within the deuterium vessel, uses the target gas as the ionization gas and is operated without gas amplification. As shown in Fig. 4.5(a) a silver foil cathode held at -80 kV lies 73 mm above a silver coated anode pad plane. A series of

field wires wrap around the sides of the active volume to ensure uniformity of the field at the boundaries. A Frisch grid consisting of tungsten wires at a pitch of 0.4 mm are strung 1.4 mm above the anode plane at 2500 V. The 12.5 cm \times 10 cm pad plane, shown in more detail in Fig. 4.5(b), is divided into a 6 \times 8 array of 48 pads. Each pad contains a noise optimized readout chain, discussed further in section 5.4. The division of the pad plane allows for determination of the X and Z stop position of the muon, while the charge arrival time can be used to determine the stop position in Y.



(a) TPC

(b) Anode Pad Plane

Figure 4.5: (a) The MuSun cryogenic time projection table. The coordinate system, displayed in the bottom right, defines in Z in the beam direction and Y transverse to the beam and in pointing up in the vertical direction. X is transverse to the beam in the horizontal direction, forming a right handed coordinate system. (b) A top view of the anode pad plane.

4.5 Decay Electron Detectors

4.5.1 Electron Proportional Chamber

Two cylindrical multiwire proportional chambers, ePC1 and ePC2, shown as two shades of green in figure 4.1, surround the deuterium vessel and allow for a three dimensional

reconstruction of the decay electron track vector. Each chamber consists of a plane of anode wires strung along the beam axis sandwiched between two layers of cathode strips, each winding around the chamber at a $\approx 45^\circ$ angle with respect to the anode wires but in opposite directions to allow for precise determination of the electron location. The same argon ethane gas as utilized by the μPC flows continuously to both chambers as a fill gas. All 512 anode wires of ePC1 are held at a voltage of 2750 V, while the 1024 anode wires of ePC2 are held at a voltage of 2600 V. During the 2014 run, a wire broke, resulting in large amounts of current being drawn at the standard operating voltage. The section containing this wire was ramped down to 2300 V using a 100 M Ω shunt to ground. Both the anode and cathode planes are equipped with electronic readout chains, beginning with custom built cards mounted directly to the chambers which contain preamplifiers followed by a discriminator. Insulated twisted pair cables carry the discriminated signals from each card to a custom built data acquisition module, which records a time stamp for each hit.

4.5.2 *Electron Scintillator Hodoscope*

A cylindrical hodoscope consisting of two layers of 16 90x15 cm BC404 scintillators is used for the electron timing information. Shown in blue in figure 4.1, the 16 scintillator segments form a barrel 78 cm in diameter concentric with the cylindrical wire chambers. Each segment consists of two layers of 5 mm thick plastic scintillator panel. Four photomultiplier tubes coupled at each end of both layers allow for a fast timing determination of the decay electrons. The photomultiplier outputs are split and sent to a set of discriminators fed to TDC's, which record time stamps for each channel, as well as a set of custom waveform digitizers built at Boston University which record the pulse information of each PMT.

4.6 Neutron Detectors

Eight supplementary Bicron liquid-scintillator neutron detectors are arranged outside of several scintillator segments, as shown in figure 4.6, comprising a 1% solid angle coverage. Although not used for the MuSun measurement directly, the neutron signals are useful for constraining certain systematic errors.




Figure 4.6: Circled in red are four of the eight liquid-scintillator neutron detectors located symmetrically outside of the electron scintillator hodoscope.

Chapter 5

TIME PROJECTION CHAMBER UPGRADES

After several trial data collection periods, the MuSun Time Projection Chamber was completely redesigned and rebuilt. This chapter details the operating principles of a TPC in the context of the motivations for upgrades. The author aided in the construction of the new chamber, and spent significant time upgrading the electronic readout as well as characterizing the response of the upgraded detector after commissioning.

5.1 Principles of Operation

The MuSun time projection chamber (TPC) is an active target ionization chamber allowing for three dimensional reconstruction of the muon stop position. A simplified diagram showing a cross section of the TPC can be seen in figure 5.1. As muons come to a stop, the deuterium gas is ionized. The electrons from the ionization are collected on the anode pad plane, sectioned into 48 pads and read out by individual charge sensitive preamplifiers. The electron




Figure 5.1: Simplified cross sectional diagram of the time projection chamber.

and ion drift in opposite directions in the presence of the large electric field with a drift velocity, \vec{v}_d , obtained from solving the Langevin equation,

$$\vec{v}_d = \frac{e\vec{E}}{2m}\tau \quad (5.1)$$

where \vec{E} is the electric field, m is the particle mass, and τ is a collisional timescale dependent on the gas, proportional to the mean free path. The drift velocities in hydrogen and deuterium gas for a range of reduced electric fields can be seen in figure 5.2. At the operating temperatures of the MuSun TPC, the drift velocity of electrons is $5 \frac{mm}{\mu s}$.

As both the ion and the electron drift, an induced charge is observed on the anode pad plane. This would pose no issue as long as the charge were collected long enough for both the ion and the electron to reach the cathode and anode, respectively. However, because the mass is significantly higher for the ions compared to the electrons, the drift velocity is three orders of magnitude slower. As only the induced charge of the electron is seen with the short time scale of the preamplifiers, the energy of the signal depends on the Y position of the muon stop. This Y dependence is removed by implementing a Frisch grid which shields the the anode from any induced charge. An equation for the shielding inefficiency as a function of geometry and operating voltages is derived in [15] via conformal representation theory. This is given by

$$\sigma = \frac{d}{2\pi p} \log\left(\frac{d}{2\pi r}\right) \quad (5.2)$$

where d is the grid wire pitch, p is the distance from the grid to anode plane, and r is the radius of the grid wires. To ensure that no signal is lost on the grid, it must be fully transparent to the electron field lines. This is also derived in [15], and the following condition must be met in order to retain all electron charge

$$\frac{E_p}{E_Q} > \frac{1 + \rho}{1 - \rho} \quad (5.3)$$

where

$$\rho = \frac{2\pi r}{d} \quad (5.4)$$

and E_p and E_Q are the electric fields in the grid-anode and grid-cathode regions, respectively.

If a high enough field can be applied, wire chambers can be operated in a regime to provide significant amplification of the signal. At the high gas densities required to depopulate the quartet state, sufficiently high fields to obtain gas gain cannot be achieved, and the MuSun TPC is operated as an ion chamber. Nevertheless, the applied electric field is important in controlling particle drift velocity and mitigating recombination effects. The main concern is prompt recombination, where charges recombine and neutralize before the drift field is able to pull them apart. This arises primarily from the columnar formation of electron-ion pairs around a muon track. MuSun Monte Carlo simulations initially used a Birk's Law parameterization of past 4.78 MeV α particle data [7] within a gaseous hydrogen TPC to quantify the recombination coefficient as

$$R \propto \frac{1}{1 + kp^3/E} \quad (5.5)$$

where p is the pressure and E is the applied drift field. The second panel of figure 5.2 shows the observed pulse amplitude in arbitrary units for a 4.78 MeV alpha source versus the cathode field strength, E , for several gas pressures.

5.2 Motivation for upgrades

Several performance considerations instigated the upgrades completed for the R2014 TPC. The need for sufficiently high, stable grid and cathode voltages as well as the benefits of constructing the target with high Z materials and improving the energy resolution were key factors in the decision to upgrade the TPC.

5.2.1 Stable Grid and Cathode Voltages

Upon ionization, the electron ion pairs can recombine, leading to a reduction in the observed signal. As discussed above, this effect is minimized with sufficient cathode voltages. Additionally, if the grid high voltage is not high enough to satisfy equation 5.3, it is not fully

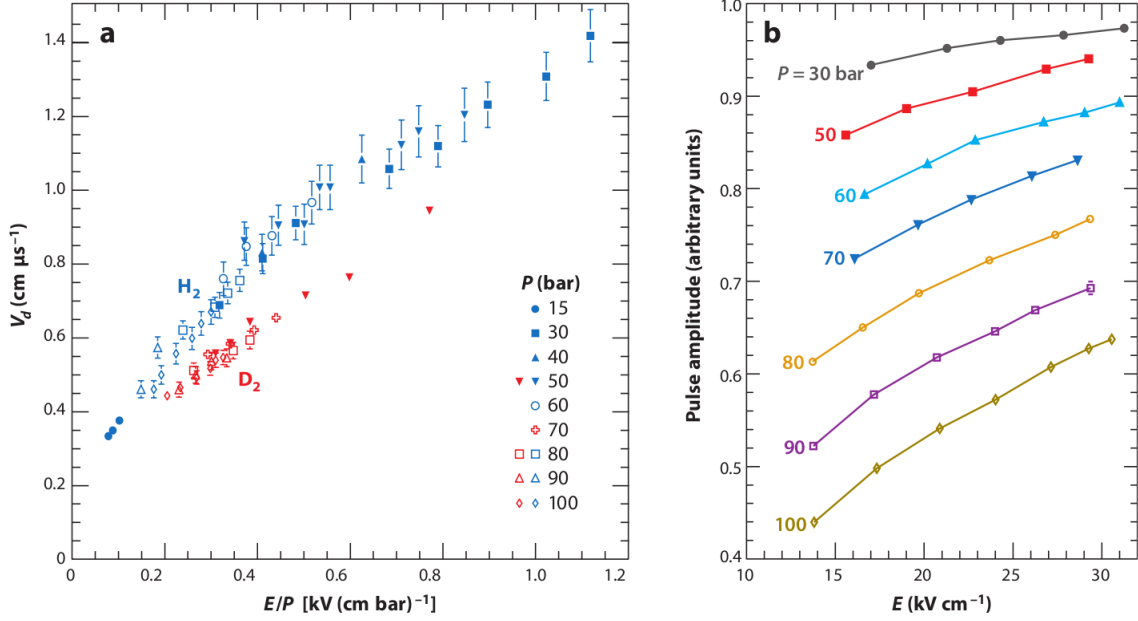



Figure 5.2: Measurements of a 4.78 MeV alpha source within gaseous hydrogen and deuterium TPC operated for muon catalyzed fusion experiments [7]. The left panel shows the drift velocity for a range of electric field to gas pressure values. The right panel shows the pulse amplitude over a range of applied drift field values.

transparent to the drift electrons and a reduction in the signal pulses is observed. It is thus critical that the MuSun TPC be operated at stable, adequate grid and cathode voltages. Retention of surface charges at the required potentials was not possible in previous constructions of the MuSun TPC, due in part to rough edges and grid wire attachment points which lead to sparking and discharges. One particular challenge arises from soldering the wires of the Frisch grid with a uniform pitch of 0.4 mm. The original MuSun TPC, shown in figure 5.3, constructed the frame by continuously wrapping the grid wire around the frame, soldering the sides, then cutting along the edges to remove the back side, leaving wires strung on a single side. Figure 5.5(c) shows the imperfections and rough surface left on the frame with this method, which lead to significant sparking and an inability to raise the grid voltage over

3 kV.



(a) 2012 TPC

(b) Frame Construction

Figure 5.3:

5.2.2 Wall Capture Materials

As discussed in section 3.3.1, any capture on nuclei other than deuterium leads to a significant deviation in the measured capture rate. As muon lifetimes in higher Z materials are relatively short, a delayed start time can mitigate effects from capture on wall materials. A significant push was made to move away from the stainless steel construction of the 2012 TPC, and redesign the detector using primarily tungsten and silver.

5.2.3 Electronic Resolution

From commissioning data taken in 2011, it was demonstrated that the gas impurity level can be monitored in-situ by directly measuring the deuterium recoil from nuclear capture within the TPC. The low energy tail of delayed ${}^3\text{He}$ fusion events is a significant background overlaying the 100-200 keV region of interest for capture recoil events. Events with a subsequent delayed electron indicate a fusion event, and can be used as a background subtraction

distribution. However, for a quantitative analysis of impurity captures, it was necessary to improve the electronic resolution of the TPC, to deepen the valley between the low energy fusion tail and the expected capture region.

5.3 New TPC Construction

A CAD drawing of the upgraded TPC can be seen in figure 5.4. The important design upgrades include:

- A composite tungsten-stainless steel grid frame
- A thin silver foil cathode
- Silver coated tungsten field wires
- Rounded Macor¹ support structures
- Tungsten grid wires

Tungsten is ideal for the grid frame, as any thermal expansion in the transverse direction must not change the tension on the tungsten grid wires. However, soldering the wires to a tungsten frame proved unsuccessful in 2013, when the grid wires began peeling away from the frame after several temperature cycles. The new design utilizes tungsten bars for the frame running parallel to the grid wires, and stainless steel bars for the frame running perpendicular to the wires, such that they may be efficiently soldered onto the frame. All components for the upgraded frame were machined at the CENPA and University of Washington machine shops, and the tungsten field wires were soldered along the stainless steel solder pads by the PSI detector group. An image of the tungsten-stainless steel connection point as well as the finished grid can be seen in figure 5.5. After all pieces were complete, the TPC was assembled by the UW group in a clean room at PSI.

¹Macor is a commercially available glass-ceramic material with low thermal expansion and little out-gassing.




Figure 5.4: CAD drawing of the upgraded design for the MuSun TPC.


The upgraded cathode, shown in figure 5.6(a), consists of a silver foil, held in place by a thin stainless steel frame with rounded edges to facilitate high voltage retention. A pin connects the cathode to the high voltage feedthrough. The cathode is mounted on four Macor posts which also received a new rounded design. Field wires, consisting of silver coated tungsten on the three sides of the TPC and silver wire at the muon entrance, are



(a) Frame Construction



(b) W-SS Connection




(c) Frisch Grid

Figure 5.5: (a) The updated grid frame consists of a composite stainless steel and tungsten construction. Tungsten bars running parallel to the Frisch grid wires prevent any loss or gain of tension under different thermal conditions. The stainless steel construction running perpendicular to the grid wires allows for soldering of the grid wires to the solder pad, denoted above. (b) Shows the joint between the tungsten and stainless steel bars. (c) A close up of the finished grid frame shows the grid wires soldered at a pitch of .4 mm.

held within the grooves of the Macor posts. High voltage resistors are soldered between each of the field wires, and placed behind the Macor posts as much as possible, to avoid muons

stopping within the medium Z material of the resistors. The field wire assembly can be seen in figure 5.6(b). The entire TPC was assembled in a clean room and sealed within the target




(a)

(b)

Figure 5.6: (a) The updated silver foil cathode, held by a rigid frame of stainless steel and connected to the HV feethrough. (b) Shows the series of HV resistors soldered along the silver coated tungsten field wires, held in place by the rounded Macor support bars which fix the cathode 71 mm above the grid frame.

volume to avoid collecting any dust or particulate that might interfere with the high voltage ramping. Pictures of the completed construction can be seen in figure 5.7.

Specifications for the finished TPC geometry, as well as standard operating voltages and resulting fields, can be seen in table 5.1. Using equation 5.2 with the parameters for the grid geometry listed, the inefficiency of the grid to shield induced charge is $\sigma=0.043$. The ratio of $E_p/E_Q=2.34$ is greater than the value of $(1+\rho)/(1-\rho)=2.29$, indicating that the conditions of equation 5.3 are satisfied and the grid should be fully transparent to electron charges at the standard operating voltages.



(a)

(b)

Figure 5.7:

Parameter	Symbol	Value
grid-anode distance	p	1.44 mm
grid-cathode distance	q	71.0 mm
grid pitch	d	0.40 mm
grid wire radius	r	0.025 mm
grid voltage	V_g	3.6 kV
cathode voltage	V_c	80 kV
grid-anode field	E_p	25.7 kV/cm
grid-cathode field	E_Q	11 kV/cm

Table 5.1: Specifications for the upgraded TPC, constructed for the R2014 data taking period.

5.4 Electronic Readout

Each of the 48 pads in the TPC are equipped with individual readout chains to allow for determination of the X and Z position of muon stops. Flat kapton cables carry the signal

through the TPC flange to the preamplifiers, followed by a shaping amplifier, the output of which is fed into custom built waveform digitizers. As discussed previously, the long tail from the ^3He fusion products which overlaps the signal region for in-situ measurement of impurity capture recoils motivated upgrades to the electronics in an attempt to improve resolution. Several upgrades were paramount in the noise reduction and resolution improvements.

As a first step, the charge sensitive preamplifiers were tested and modified to improve resolution via studies of the equivalent noise charge (ENC). The ENC curves allowed for determination of optimal component selections. Previously, the preamplifiers were mounted outside of the insulating vacuum, such that the flat, non-insulated kapton cables traveled several meters before reaching the preamplifiers. This long distance allowed for pickup, especially in the region near the high voltage feedthrough. However, coaxial cables increase the input capacitance of the preamplifiers, thereby increasing the ENC. In order to mitigate the noise effects without significantly increasing the input capacitance, the preamplifiers were redesigned to fit within the insulating volume, directly behind the TPC flange. Additionally, the preamplifiers were cooled by a dedicated liquid nitrogen line, further improving the resolution. More details on the development of the preamplifiers can be found in [42].

The best resolution is obtained by utilizing the full dynamic range of the digitizers. However, if the full dynamic range is used for muon pulses, the larger signals such as proton triton fusion events hit the rails of the digitizer and are clipped. In order to maximize the resolution without losing these high energy signals, a dual gain readout was implemented in 2014. The shaping amplifier was split into a low gain output, which could catch the larger signals, and a high gain output with ten times higher gain which recorded the muon signals with optimal resolution.

5.5 Commissioning

The new TPC was installed in 2014 for a commissioning period before production data taking. The target grid and cathode voltages were easily achieved, and remained stable for the entire commissioning period with minimal high voltage trips. An in-situ alpha source located in the corner of the MuSun TPC allowed for scans of signal energy versus the applied grid voltage, allowing for the determination of signal saturation as shown in figure 5.8. At the standard operating voltage of -3.6 kV, the grid is sufficiently transparent, and nearly no electron signal is lost.




Figure 5.8: Observed pulse amplitude for the in-situ alpha source at several values of the grid high voltage.

The electronics were calibrated using a small test circuit, which applied a square voltage pulse over a capacitor to inject a known amount of charge into each preamplifier channel. An electronic resolution of 10keV was measured; a three fold improvement compared to the previous readout chain.

5.6 Understanding Detector Response

Confident determination of the muon stop position is critical in ensuring that each event consists of a muon decay in deuterium, as decays in wall materials lead to a significant distortion of the measured lifetime. The TPC is also used to monitor the impurity level in-situ via measurement of the capture recoil spectra. Additionally, a fusion interference systematic will be addressed, which relies on the tracking properties within the TPC. As such, it is of great importance to understand the response of the chamber.

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

$$-\left\langle \frac{dE}{dx} \right\rangle = Kz^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 T_{max}}{I^2} - \beta^2 \right] \quad (5.6)$$

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, and T_{max} is the maximum energy transfer allowed in a single collision for a particle of mass M ,

$$T_{max} = \frac{2m_e c^2 \beta^2 \gamma^2}{1 + 2\gamma m_e/M + (m_e/M)^2} \quad (5.7)$$

and K is given by

$$K = 4\pi N_A r_e^2 m_e c^2 \quad (5.8)$$

where N_A is Avagadro's number and r_e is the classical electron radius. Equation 5.6 can be integrated to obtain a range-energy relation, which can be used to determine the maximum distance each particle of a given energy can travel within the target gas. The energy and range, obtained via SRIM² calculations, of muons and fusion product particles at the running conditions of the MuSun target gas can be seen in table 5.2.

Due to recombination effects, the total energy of some particles is not fully observed by the detector. As mentioned in section 5.1, a Birk's parameterization has been found to

²The Stopping Range of Ions in Matter software, available at www.srim.org.

characterize the recombination effects of alpha particles within a gaseous hydrogen TPC. For the MuSun data, this parameterization fails to reproduce the low energy tail of the ${}^3\text{He}$ peak. In order to correctly quantify the TPC recombination response, an angular dependence must be added such that the recombination coefficient is given by

$$R = \frac{1}{1 + \frac{k \frac{dE}{dx} \phi^2}{E_d \sin(\theta)}} \quad (5.9)$$

where ϕ is the target density, E_d is the drift field, θ is the angle of the track with respect to the drift field, and k is a constant extracted from fitting the ${}^3\text{He}$ peak data. The observed energies after recombination effects have been taken into account are included in table 5.2.

	$E(\text{MeV})$	$E_{obs}(\text{MeV})$	$R(\text{mm})$
μ	3.3	3.0	80.8
${}^3\text{He}$	0.82	0.35	0.292
$\mu{}^3\text{He}$	0.8	0.6	0.653
t	1.01	0.8	.935
p	3.02	2.7	13.2

Table 5.2: Energy, both expected and observed, and range of muon catalyzed fusion products in the TPC. The observed energy of the second column is lower due to recombination effects of the gas. All ranges are obtained via SRIM calculations in D_2 at the MuSun target density of 0.065 LH2.

As the energy loss scales roughly as $1/\beta^2$, the energy deposition peaks at low energies, just before the particle comes to rest. The energy loss as a function of the depth into the target, typically referred to as a Bragg curve, results in the majority of the energy from a muon stop deposited in the last pads of the track, as illustrated in figure 5.9. As the muon deposits this characteristic energy as it stops, and muon catalyzed fusion products




Figure 5.9: The energy deposition of an incoming muon increases as it slows to a stop, resulting in the largest amount of energy being observed near the end of the track. In order to better characterize specific event topologies, the energy in the most downstream pad (E0) and one pad upstream (E1) are considered.

deposit extra energy in the pads surrounding the stop depending on their range, it is useful to characterize events based on the presence and distribution of extra track energy. Three essential energies are defined in order to investigate specific event topologies:

- **E0** : The energy deposited on the most downstream pad
- **E1** : The energy deposited one pad upstream of the most downstream pad
- **E9** : The sum of the energy on the most downstream pad and the eight adjacent pads

The utility in each of these is discussed in the subsequent sections.

5.6.1 *E0 vs E1*

The range of the proton produced in a muon catalyzed fusion event at the target conditions is large enough that a portion of its energy may be deposited in the pads surrounding the stop. The range of the ^3He particles on the other hand is short enough that its energy is deposited into the same pad as the muon stop. As such, a useful tool in characterizing particular event geometries is histogramming the energy in the most downstream pad Z row (E0) vs the energy one pad row upstream of the stop (E1). The E0 vs E1 distribution for all events in the R2014 dataset can be seen in figure 5.10. The regions circled in figure 5.10




Figure 5.10: Energy deposited on the most downstream pad (E_0) versus energy deposited on one pad upstream (E_1).

are understood to result from the following event topologies:

1. Clean muon stops within the chamber. As the muon stops further into the last pad, E_0 increases as E_1 decreases, leading to the band shown.
2. Punch-through muons. Muons which pass out of the TPC leave a reduced amount of charge on the last and second to last pad.
3. Muons stops with a 3He . A muon catalyzed fusion event resulting in a 3He gains an additional 400-600 keV from the fusion products resulting in the shifted band shown.
4. Muon stops with a p-t on the stop pad. A Muon catalyzed fusion event resulting in a proton and a triton receives an additional 4MeV from the fusion products. If the energy from both fusion products is added to the stop pad, the E_0 is shifted up by 4 MeV.

5. Muon stops with a p-t, proton leaves the stop pad. With the running conditions of MuSun, the proton has enough range to travel to an adjacent pad. These events have the extra energy from the p-t fusion distributed either upstream or downstream of the muon stop, resulting in two bands.
6. Muon stops with a deuterium recoil. Neutrons produced from muon catalyzed fusion events can scatter off deuterium within the chamber, resulting in a recoil energy observed in this region.
7. Same as 6.

5.6.2 *E9 Energy*

A second metric used to identify specific event types in subsequent analyses is the sum of the energy on the stop pad and the eight pads adjacent to it, referred to as the E9 energy and depicted in 5.11(a). A distribution of the E9 energy obtained from all events of the R2014 data can be seen in 5.11(b). This metric is utilized to tag muon-catalyzed fusion events which result in a proton and a triton, described further in chapter 10. Although the proton has a long enough range to leave the stop pad, the E9 energy will include the charge deposition from the proton even if it travels into an adjacent pad.

5.6.3 *MuSun Monte Carlo*

A full GEANT4 Monte Carlo simulation has been developed by many members of the collaboration, and is used to create 2×10^9 muon events emulating the TPC geometry and beam conditions of the R2014 data. Sophisticated response code applies electronic noise, recombination, and drift effects via models developed to replicate aspects of the TPC response observed from data. The output files from the response stage are formatted to be utilized in the standard analysis chain, with the addition of truth information for all particle tracks within each of the MuSun detector systems. This provides a powerful tool in understanding

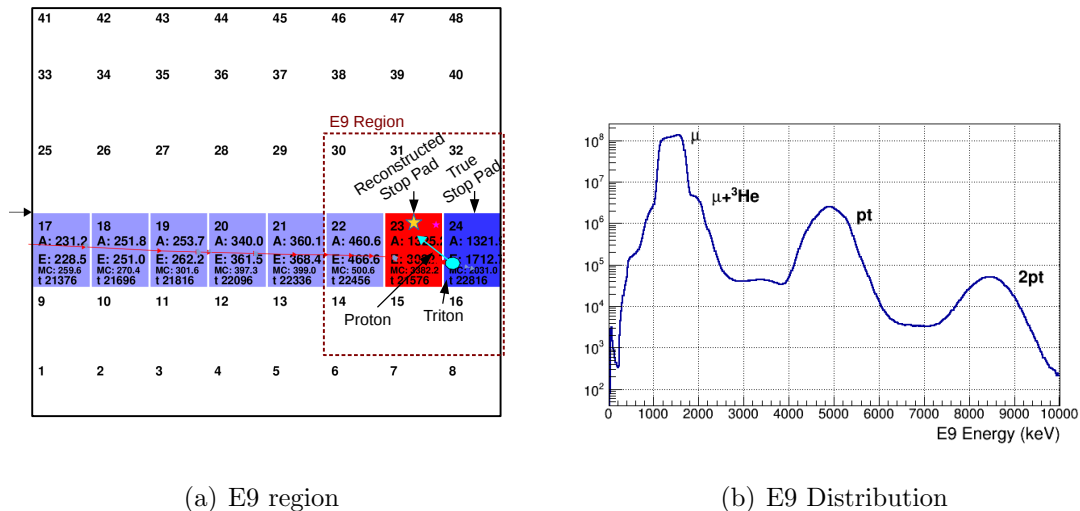


Figure 5.11: (a) Definition of E9 energy– the sum of the energy on the reconstructed muon stop pad (indicated by the yellow star), and the eight surrounding pads. In this event, the muon stops in pad 24 and a fusion event results in a triton, denoted by the light blue circle, and a proton which travels upstream, denoted by the light blue arrow. The E9 energy is the sum of the energy on the stop pad and the eight adjacent pads, such that all proton and triton energies are included even if the proton traverses to an upstream pad and the stop pad is mis-reconstructed. (b) The E9 distribution observed from the R2014 data.

the detector response of the TPC as well as the reconstruction abilities of the analysis.

5.6.4 Effect of PT fusions on Muon Tracking

In order to confirm the origins of regions detailed in figure 5.10, MC simulations were used by the author to select specific event topologies using the truth information of the TPC tracks. In particular, for the development of robust muon tracking algorithms it is useful to study the muon catalyzed fusion events resulting in a proton, which can travel outside of the stop pad. Three specific topologies are considered; events in which the proton stops in the same pad as the muon stop, one pad upstream, and one pad downstream.

PT Stop Pad = μ Stop Pad The E0 vs E1 distribution for muon catalyzed fusion events where the proton stops in the same pad as the muon, according to the MC truth information, can be seen in figure 5.12. An explanation of the events populating each of the labeled regions

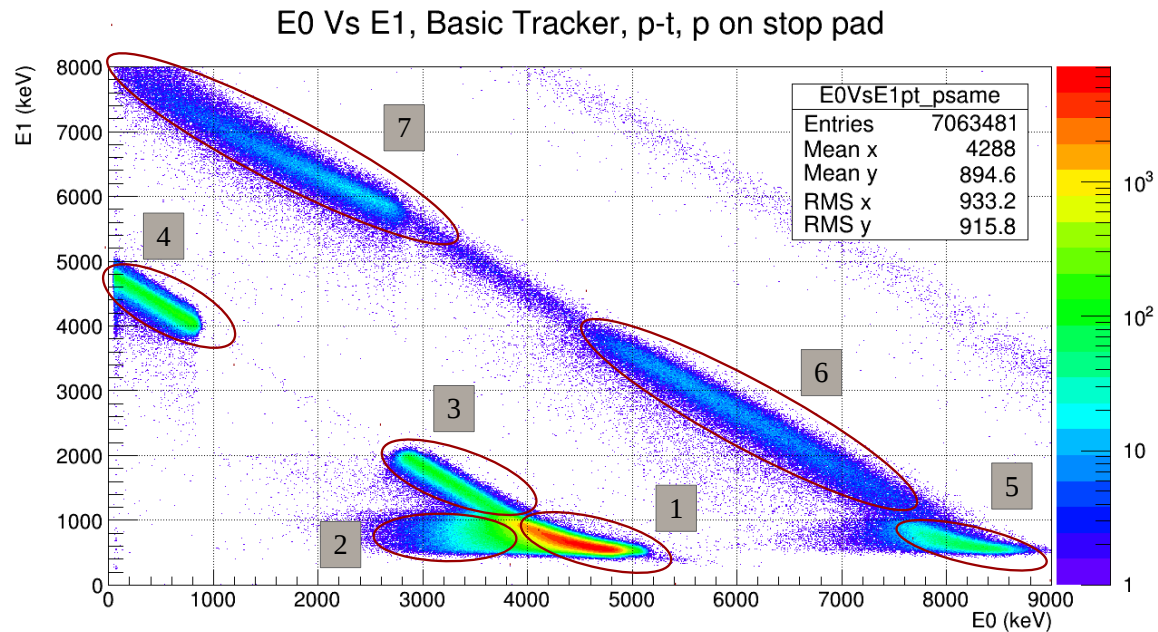


Figure 5.12: E0 vs E1 for MC events with a Proton-Triton fusion in which the true proton stop position is on the same pad as the true muon stop. See text for explanation of the labeled regions.

is as follows:

1. Both the proton and triton energies are deposited in the muon stop pad. This increases E0 by 4 MeV and does not change E1.
2. This is a smearing of band 1 due to recombination effects. The events with the smallest energy in this band contain a p-t emitted parallel to the drift field, such that the recombination is maximal.

3. The proton stops in the same pad as the muon, but the triton stops one pad upstream. This reduces E0 and increases E1 of the previous band by up to 1 MeV (the triton energy).
4. The proton stops in the same pad as the muon, but the triton moves one pad downstream. E0 then takes a maximum value of 1 MeV (the triton energy) and E1, containing the muon and proton, ranges from 4-5 MeV.
5. Band 1 for 2 pt events.
6. Band 3 for 2 pt events.
7. Band 4 for 2 pt events.

PT Stop Pad = μ Stop Pad + 1 The E0 vs E1 distribution for events where the proton travels downstream of the muon stop according to the MC truth info can be seen in figure 5.13. An explanation of each region depicted in the figure is as follows:

1. Proton travels one pad downstream of muon stop, depositing up to 3 MeV into the downstream pad. The cut around 3 MeV in E0 denotes the proton range.
2. Band 1 for 2 pt events.
3. Events where the Proton truth indicates that the proton moved downstream, but it was not far enough into the next pad to be above threshold, thus reflecting the most prominent band in the E0 vs E1 histogram from events where the proton stops in the same pad as the muon.
4. Band 3 for 2 pt events.
5. Events where one proton goes upstream and the second downstream, such that E0 is only up to the 3 MeV deposited by the first proton, and E1 contains the muon, triton, and other proton.




Figure 5.13: E0 vs E1 for events with a Proton-Triton fusion in which the proton travels one pad downstream of the muon stop pad.

PT Stop Pad = μ Stop Pad - 1 The E0 vs E1 distribution for events where the proton travels upstream of the muon stop according to the MC truth info can be seen in figure 5.14.

1. The muon stops and the proton travels upstream, depositing at most 3 MeV into E1. The sharp edge on this band at low E0/ high E1 denotes the range of the proton.
2. Band 1 for 2pt events.
3. One pt goes upstream, the other goes downstream. The E0 can then have a maximum energy of the proton (3 MeV) and E1 contains the muon stop and triton energy.

These studies illustrate the difficulty in designing a robust muon tracking algorithm. If the most downstream pad is taken to be the muon stop position, all PT fusion events in which




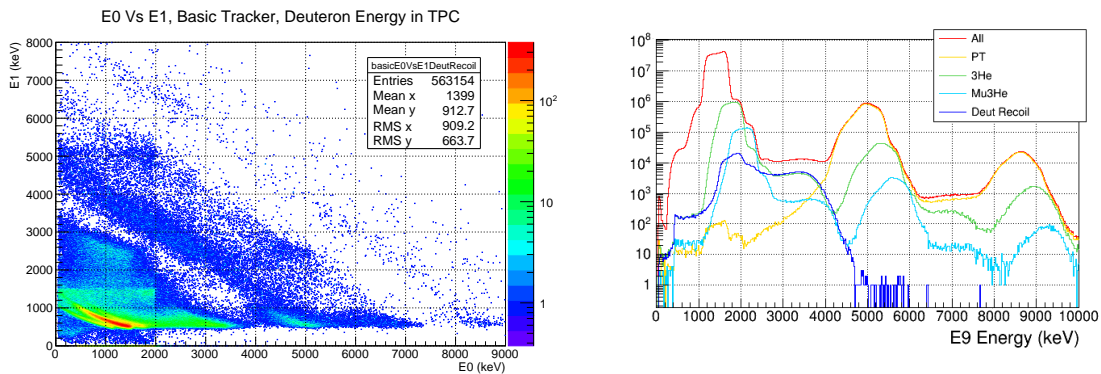
Figure 5.14: E0 vs E1 for events with a Proton-Triton fusion in which the proton travels one pad upstream of the muon stop pad.

the proton travels downstream will cause the muon stop position to be mis-reconstructed. Several tracking algorithms developed to reduce the mis-reconstructions based on the detector characterizations shown above are discussed in section 7.3.2.

5.6.5 Effect of n - d Scatter Events

For both the pt fusion event tagging required for the fusion correction and identification of nuclear recoil events in the impurity analysis, the E9 energy is used, but is contaminated by a background of deuteron recoil events. The fusion neutron created by ${}^3\text{He}$ fusion events can scatter off deuterium nuclei, depositing energy in the TPC of a few hundred keV. An E0 vs E1 distribution for events in which the MC truth confirms energy deposition from a deuteron is displayed in figure 5.15(a). This distribution confirms that the suspected regions,

labeled 6 and 7 in figure 5.10, are indeed due to n-d scatter events. The MC E9 distribution separated by event type using the MC truth information in figure 5.15(b) demonstrates the effect of these events in E9 space. The MC truth information was used to determine the extent of the background, as well as develop a background subtraction technique detailed in section 10.2.1.



(a) E0 vs E1, n-d scatter events

(b) MC E9 Distribution

Figure 5.15: (a) E0 vs E1 in which the MC truth information confirms an n-d scatter event within the TPC and (b) Monte Carlo E9 energy distribution, separated by event type.

Chapter 6

R2014 DATASET OVERVIEW

The R2014 collection period marked the first production data taking for the MuSun experiment, newly situated in the PIE1 beamline. Data was continuously collected for ten weeks, and the accumulated muon-electron events, passing basic quality cuts, from the fast online analysis over the course of data collection can be seen in figure 6.1.




Figure 6.1: Accumulated statistics for both muon polarities for the R2014 data collection period. The dashed bands show periods of accelerator maintenance.

6.1 Data Collection Summary

Several attempts were made in the first several weeks to reach an optimal beam tune, maximizing the muon rate while suppressing the background beam electrons. As such, the running conditions were variable over the course of the run. In addition, several larger changes were made throughout the run leading to possible changes in data quality. For the first week of data collection, the muon veto counter, the muSCa, was not connected. Once connected, the

veto counter was misaligned until a few weeks into the run. Although this has not been found to have large impact on data quality, it was taken into consideration when creating datasets to check for any inconsistency. For the first half of the R2014 dataset, an upstream beam collimator, located within the bending magnet and used to stop beam electrons following the separator, was misplaced, such that an optimal beam tune could not be achieved. For a small period the collimator was removed completely, then replaced and correctly aligned. Additionally, an incorrect current value for an upstream magnet lead to poor muon-electron separation resulting in a large beam electron background throughout the data collection period. This is discussed further in appendix D.

Ultimately, 6×10^9 good μ^- decay events were collected over the ten week period, resulting in half of the necessary statistics to reach the MuSun precision goal. Additionally, 1×10^9 μ^+ decay events were acquired for systematic studies. All proceeding chapters of this thesis focus on the analysis of this 40 TB dataset, including data organization and quality assurance, all performed by the author.

6.2 Datasets

The R2014 data was divided into 16 datasets based on running conditions. Each dataset with the number of muon entrances and a brief description of the conditions can be seen in table 6.1. The number of events listed for each dataset quatifies the number of good muon-electron pairs remaining after all selection cuts have been applied.

For the purpose of obtaining high statistics lifetime fits, a set of five larger datasets is created from the datasets outline above. These datasets are detailed in table 6.2. The first five datasets are the μ^- data, broken into five run periods. The last two datasets are the full statistics broken into two periods of the collimator configuration. A trend histogram of the good event rate, consisting of the number of muon-electron events which pass all standard cuts, is seen in figure 6.2 with colors denoting the different datasets.

Dataset	Events	Running Conditions
DS01	1.6×10^8	no muSCa
DS02	7.2×10^7	muSCa connected
DS03	7.7×10^7	variable file size
DS04	4.2×10^8	muSCa misaligned
DS05	2.3×10^8	muSCa misaligned
DS06	7.5×10^7	muSCa misaligned, muCL added
DS07	1.8×10^8	muSCa realigned
DS08	5.2×10^7	unstable beam
DS09	–	low beam
DS10	7.8×10^8	collimator removed
DS11	1.3×10^8	new collimator
DS12	1.9×10^8	muPC wire swapped
DS13	9.2×10^8	muCL turned on
DS14	5.5×10^8	muPC wire back
DS15	1.1×10^9	mu+
DS16	1.2×10^9	back to mu-

Table 6.1: R2014 sub-datasets for running on stampede, based on running conditions. The number of events listed are the good muon-electron pairs after all selection cuts have been made.

Dataset	Combination	Events	Conditions
DSM1	DS01-DS03,DS07-DS09	5.4×10^8	muSCa aligned
DSM2	DS04-DS06	7.2×10^8	muSCa misaligned
DSM3	DS10	7.8×10^8	no collimator
DSM4	DS11-DS13	1.2×10^9	new collimator
DSM5	DS14,DS16	1.8×10^9	new collimator
DSP1	DS15	1.1×10^9	mu plus

Table 6.2: R2014 sub-datasets used for comparison of high-statistics lifetime fits.




Figure 6.2: Number of good decay events per data block versus run number. The different colors denote the different datasets which have been created based on running conditions.

Chapter 7

ANALYSIS OVERVIEW

The MuSun data analysis framework, developed over years by many members of the collaboration, aims to process the raw data files to create event related objects, build muon decay events based on these objects, and create quality lifetime distributions which are fit to extract the observed decay rate. This chapter provides a general overview of the analysis and highlights several upgrades and developments contributed by the author, including

- updated muon entrance framework
- simplified electron definition development
- improvements to the muon tracking algorithms based on Monte Carlo studies
- trend histogram development
- improvements to error handling
- solutions to lifetime fitting issues.

The overview of the MuSun analysis is presented in section 7.1, followed by a discussion of processing the data from the entrance detectors, TPC, and electron detectors. Before the final analysis, several quality and error checks are used to remove any problematic data runs, as discussed in section 7.5. Finally, the lifetime fitting procedure is detailed in section 7.6.

7.1 Software Overview

The MuSun analysis chain consists of three main stages, an initial stage (MU) to make objects for each detector from the raw data, a second stage (MTA) which forms muon events containing all objects within a certain time range around each muon, and a final stage which

produces a tree of high level information used to create lifetime histograms with various cuts. A diagram of the analysis chain can be found in figure 7.1. The Stampede cluster at the Texas Advanced Computing Center (TACC) is used to analyze the 40TB of R2014 data, stored on the Ranch mass storage facility.

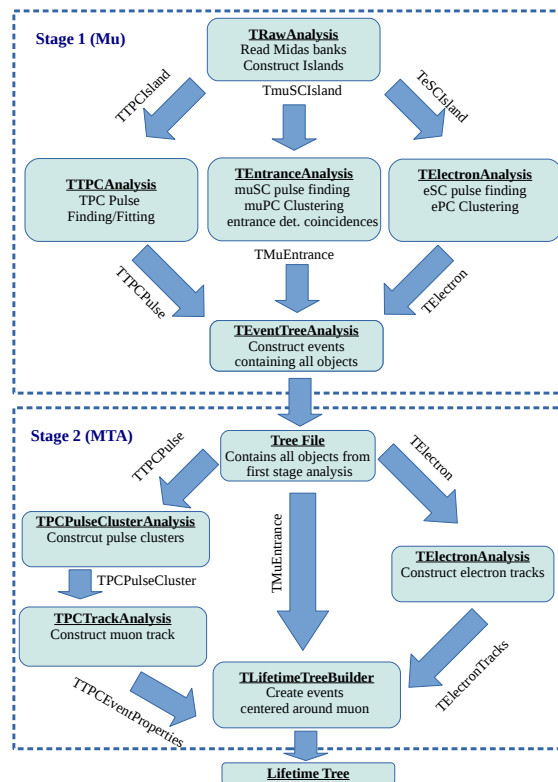


Figure 7.1: Diagram of the MuSun analysis chain.

7.2 Muon Entrance

The first stage of the analysis begins with the creation of muon entrance objects, comprised of objects from the three entrance detectors, the muSC, muSCA, and the muPC, all described in section 4.3. Each of these objects are created with the following information.

MuSC Objects As discussed in section 7.5.2, the timing information from the muon scintillator is discriminated and sent to two separate TDC channels. MuSC objects are created from the coincidence of both TDC signals. These objects contain the time from the TDC of the muSC output, which is the time primarily used for creation of entrance objects. In addition to seeding the creation on muon entrance objects, the MuSC objects are used for pileup protection.

Additionally, pulse information including time, amplitude, and energy of the fitted muSC pulses from the WFD are stored. The advantage of the muSC pulse information is the much lower threshold, illustrated in the comparison of the muSC WFD amplitude of all pulses to those in coincidence with a TDC signal of figure 7.2. The low amplitude peak is from electrons in the scintillator. The muSC WFD is not used to seed muon entrances, but is useful in pileup studies.




Figure 7.2: Comparison of the amplitude of all muSC WFD pulses and those in coincidence with a TDC signal. The TDC threshold is around 35 ADC channels.

MuSCa Objects MuSCa objects are created using the TDC time from the muon veto scintillator. The time, amplitude, and energy of the WFD pulse information from the veto

scintillator is also contained within the MuSCa objects. These objects are used to veto muon events off the beam axis.

MuPC Objects Objects are created by forming X and Y clusters from individual wire hits from the muon multiwire proportional chamber. X and Y clusters which are temporally and spatially coincident are then used to create MuPC objects, which consist of time and position information of X and Y clusters. These objects are used to determine the position of the muon entrance, as well as aid in the application of pileup protection.

MuEntrance Objects MuSC objects are used for the primary seed of MuEntrance objects. All MuSCa and MuPC objects in coincidence with each MuSC object are saved as references to the MuEntrance objects. Several coincidence and pileup flags are set within each MuEntrance object in order to make quality selections later in the analysis. All muon entrances in the final analysis must satisfy a set of “HasBestEntrance”(HBE) conditions. The entrance statistics after each additional cut can be seen for all five datasets in figure 7.3, and are explained below:

1. both muSC copies triggered and in coincidence (muSC Triggered)
2. artificial deadtime applied to the muSC (muSC DT)
3. muSC anti-coincident with muSCa veto (!muSCa Coinc)
4. bookending applied (muSC BE)
5. muon coincident with kicker signal (kick)
6. muPC X, Y, and XY coincidence required (muPC Coinc)
7. pileup protection on muPC XY clusters (muPC PP)
8. pileup protection on the muSC (muSC PP)

A standard deadtime of 29 ns is applied to the muSC and any event within 30 μ s of the beginning or end of the data block is rejected for book-ending. The muSCa is not used

for pileup protection, due to the 15% reduction in statistics resulting from the muSCA misalignment in datasets DS04-DS06. Lifetime scans over the pileup definition show no shift in rate in the removal of the muSCA from the pileup window, and no increase in efficiency of pileup protection is observed when the muSCA is included in pileup. Trend histograms displaying the effect of the increasingly strict cuts, as well as the pileup protection, can be found in figure 7.4.




Figure 7.3: Entrance statistics for each of the five R2014 sub-datasets.

As discussed in chapter 6, the beam collimator was misaligned for the first half of the data collection period, such that the electron background was much higher. Eventually the collimator was removed and replaced. As the thresholds are low enough for electrons to be detected by both the MuPC and the MuSC WFD, the effects from this electron background increase are observed in both the muPC and the muSC WFD pileup, as seen in figure 7.4(b). As such, only the muSC WFD Hi signals, with a threshold on the energy, are used in the

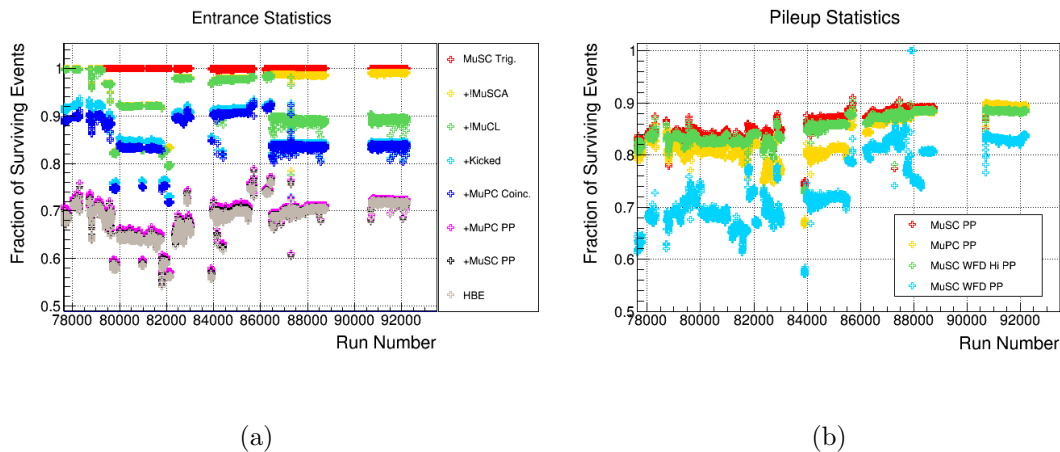


Figure 7.4: Trend plots showing the fraction of data block in each run file which survive the increasingly strict entrance cuts. (a) Illustrates the fraction of remaining events after each cut is applied, and (b) shows the pileup protection specifically.

pileup protection. Roughly 30% of all data was taken with the old collimator, 16% was taken with the collimator removed, and 54% was taken with the improved collimator. About 70% of all muon entrances satisfy every requirement for a HasBestEntrance muon, except for DSM2 where the muSCA was misaligned and only 60% pass all entrance cuts.

7.3 Muon Stop Determination with the TPC

The WFD waveforms from each of the 48 charge collection pads are processed in the first stage of the analysis. In order to quantify the energy deposition leading to each pulse, the baseline must be subtracted, and a calibration factor must be applied. Pulse clusters are formed using spatially and temporally coincident pulses. Finally muon tracks are formed from the pulse cluster objects, using several algorithms.

For the purpose of analysis, a standard coordinate system within the TPC is adopted, as denoted in figure 7.5. The Z direction is defined to be along the beam axis, the X direction is defined as the horizontal dimension parallel to the beam along the pad plane,

and the Y direction is the third dimension, pointing vertically upwards to form a right handed coordinate system. Frequently throughout the analysis it is convenient to sum the charge along X or Z rows, which are zero indexed and labeled in the figure. As a reminder, the quantities E0 and E1 are defined for the purpose of characterizing the TPC response in section 5.6. E0 is the sum of the energy deposited on the most downstream Z row, and E1 is the sum of the energy deposited one row upstream of the final row. In addition, the E9 energy is defined to be the sum of the energy deposited on the stop pad and the eight adjacent pads. In each case, the sum includes only the pulse included within the stop cluster, as discussed in section 7.3.2 below.




Figure 7.5: Coordinate system of the TPC.

7.3.1 Pulse Processing in the TPC

Dual-Gain processing

The final stage in the TPC electronic readout consists of a shaping amplifier fed into a custom built waveform digitizer. To optimize the resolution of muon signals in the TPC, it is best to amplify the 1 MeV muon signal to utilize the full 1 V dynamic range of the 8-bit digitizer, resulting in 4 mV for each of the 255 channels. However, if the amplification is set in this way any signals above 1 MeV, such as muon catalyzed fusion events, will hit the rails of the digitizer, resulting in clipped pulses. In order to maximize the resolution of the muon signals without sacrificing the large energy pulses, two amplifications were sent to two separate digitizers. The full dynamic range is utilized in a high gain (HG) digitizer, while a second low gain (LG) digitizer records the same signals with 10x less amplification to catch any larger signals without clipping. The first step in analysis of the TPC data merges the two gains in order to recover any clipped high gain pulses. For any waveform in the HG digitizer which hits the 255 ADC channel rail, the LG waveform which contains the time at which the HG waveform was clipped is retrieved, scaled by the gain, and used to replace any samples in the HG waveform which hit the 255 channel rail. An example of a waveform in which the pulse is clipped in the HG, as well as the final merged waveform can be seen in figure 7.6.

The gain ratio of each channel is determined in two ways: comparing the amplitudes of the pulses which trigger both gains, and matching the amplitudes of forced trigger events that have fewer statistics but a larger range of pulse amplitudes. For both histograms, a linear fit to the profile of the HG versus LG maximum sample is used to extract the relative gain calibration. An example of the two-dimensional histogram and the profile fit for one stable region with pad 20 can be seen in figure 7.7.

For each data file, histograms are created for each pad, using the triggered pulses in the overlap region to find the relative gain as the FT waveforms are too statistically limited to


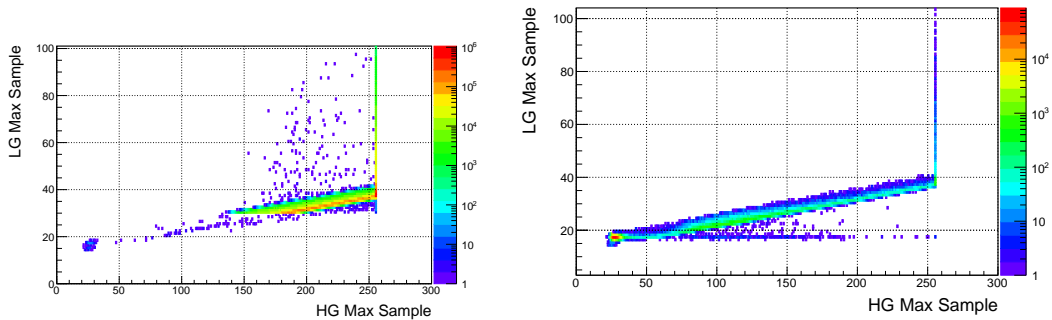



Figure 7.6: Example of a waveform in which a large energy event hit the rails of the high gain digitizer, shown in red. All samples of the HG waveform above 255 ADC channels are replaced by the scaled low gain waveform, resulting in the final merged waveform shown in blue.

determine the relative calibration for a single run. A histogram, created by extracting the information from each file, is used to display trends over the data collection period. A trend histogram with the regions of stability and a comparison of relative gains acquired from the two methods for each stable region can be seen in figure 7.8.

It is important to note that the majority of the clipped pulses are muon catalyzed fusion events. Thus, any loss in these waveforms, or inability to properly reconstruct the energy of these pulses can lead to fusion interference effects, discussed in detail in chapter 10. The above method relies on the ability to find the corresponding LG waveform, and accurately determine the pedestal and gain ratios of each channel to accurately merge the waveforms. Trend histograms reveal that a matching LG waveform is missing at the 10^{-4} level for all



(a) Pad 20 HG vs LG for all triggered pulses (b) Pad 20 HG vs LG for forced trigger pulses



(c) Profile fit for all triggered pulses (d) Profile fit for forced trigger pulses

Figure 7.7: The maximum high gain versus low gain sample for (a) waveforms in the region of overlap and (b) forced trigger waveforms. A TProfile is created for each (c,d) and a linear fit is used to extract the relative gain for each pad.

channels during most segments of the data collection period, as shown for Pad 11 in figure 7.9. This is sufficiently low that no adverse effects are expected. However, for some periods this fraction is much higher. For the clipped waveforms where no coincident low gain island is found, a template fit created from typical waveforms is used to recover the missing pulse information. As the majority of overloaded pulses are due to fusion products, a precise resolution is not necessary; enough of the pulse energy must be retained to tag the event as a fusion for the interference correction. As such, the template fit is sufficient in recovering




Figure 7.8: Run by run relative HG/LG calibration for pad 20. The points in black show the relative gain using the triggered pulses, as the FT waveforms are too statistically limited to determine the relative calibration for a single run. The dashed vertical lines denote the stable periods used to group the runs for high statistics extractions of the relative gains. The red and blue horizontal lines signify the calibration obtained using the triggered and forced triggered fits to obtain the calibration for each stable period, respectively.

the information lost by the clipped waveforms with no matching LG pulse.

Baseline Determination

As pulses within the TPC have an undershoot, each digitizer channel has a DC offset applied to allow for negative pulse samples. This channel by channel pedestal must be subtracted in order to accurately determine the amplitude and integral of each pulse. Nominally, the digitizers only begin recording once a channel goes above a set threshold. In order to characterize the baseline of each WFD channel, a forced trigger (FT) chain was implemented, which causes the WFD to record for $25 \mu s$ following the muon entrance for every 1000th




Figure 7.9: Shown in the first panel are the fraction of all pad 11 waveforms which were overloaded throughout the course of the data collection period. The second panel shows the fraction of the overloaded waveforms which did not contain a time coincident low gain channel.

muon. As any muon signals should be gone after the $16\ \mu s$ drift time of the chamber, the remaining $9\ \mu s$ of each waveform contains a signal free region, from which the baseline can be extracted. For each run file, a Gaussian fit to the sample distribution in this region is used to extract run by run pedestals for each pad. This information is stored in the run database, which is accessed at the beginning of the stage one analysis to apply the correct pedestal to each pad. A trend histogram of the pedestal value for each pad over the course of the data collection period can be seen in figure 7.10.

Pulse Processing

A time over threshold (TOT) algorithm is used to find and characterize pulses within the digitized TPC waveforms. The method, developed by Michael Murray and explained in greater detail in his thesis [38], searches for stretches of pedestal subtracted samples above

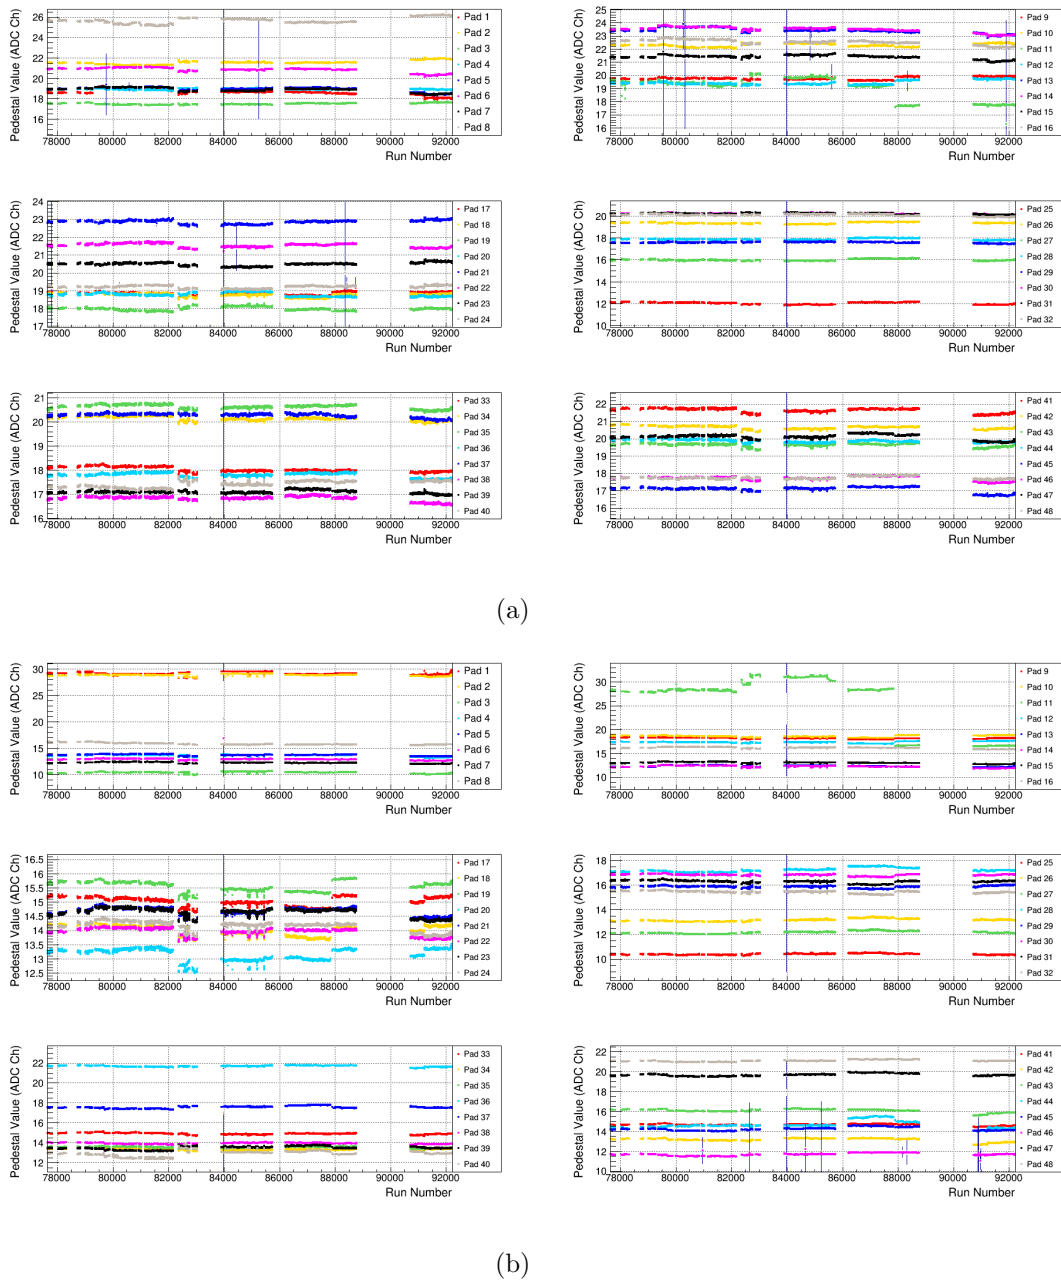


Figure 7.10: Trend plots showing the run by run pedestal values obtained for each pad for (a) the high gain and (b) the low gain WFD channels.

a threshold of 10 ADC counts. The longest stretch in each waveform is used to create a TOT pulse object if the amplitude of the pulse is above a pedestal subtracted threshold of 17 ADC channels, corresponding to a pulse amplitude of about 70 keV. The TOT pulse objects contain the amplitude and center time of the pulse, determined by the value and location of the maximum sample, in addition to the pulse energy obtained from a raw integral over threshold. Additionally, the start and stop time of the pulse, determined by the points at which the waveform goes above and below threshold, respectively, are saved to the pulse object.

Calibration

As mentioned above, any changes in efficiency over the TPC readout pads can impede the ability to correct for muon-catalyzed fusion. In addition, the impurity monitoring relies on searching for small signals following a muon stop. As such, it is of prime importance that an accurate pad by pad relative calibration is achieved. Several physical signals allow for the determination of a relative calibration, including the muon stop, proton triton fusion events, and delayed ${}^3\text{He}$ pulses. As the impurity monitoring requires the best resolution, the delayed ${}^3\text{He}$ signals are used to obtain a relative pad by pad calibration for all pads, except the first pad row, where muon tracks are not created because a track requires at least two pads. For the first Z row, the Bragg energy deposition is utilized, and the energy of tracks for which the second pad upstream lies on the first pad row is used to obtain a pad by pad calibration. As the muon energy deposition is modeled by the Bethe-Bloch formulation of equation 5.6, there is a maximum energy deposition in the last pad for each muon cluster, depending on how far into the final pad the muon stops. The calibration can be checked by looking at the distribution of energy in each pad for events in which the pad is the most downstream pad within the cluster. This is shown for each pad in figure 7.11. The pads displayed in red and pink are on the edges of the anode plane and are susceptible to punch-through events, which

do not deposit the full muon energy and distort the observed energy distribution. However, the high energy edge demonstrates the consistency of the calibration for all pads.

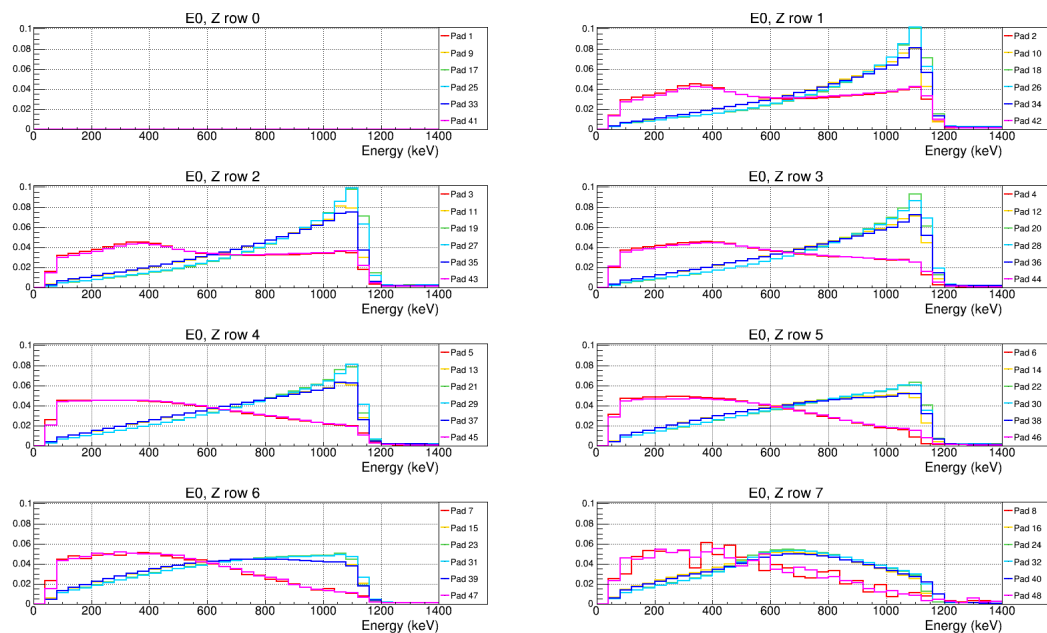


Figure 7.11: Energy distribution for each Z row in which the most upstream pad within a cluster was the pad shown. This is used to check the consistency of the energy calibration of all pads, as the high energy edge should coincide for all distributions. The pads displayed in red and pink lie on the X edges of the anode plane and are susceptible to punch-through muons, accounting for the secondary peak at low energies.

7.3.2 Muon Tracking in the TPC

A full 3D reconstruction of the muon stop position is required to ensure that each accepted event stops within the TPC. In order to determine muon stops and their location within the target, pulse clusters are first formed from temporally coincident pulses. Muon tracks which determine the exact stop location are then determined from the pulse clusters.

Pulse Clustering

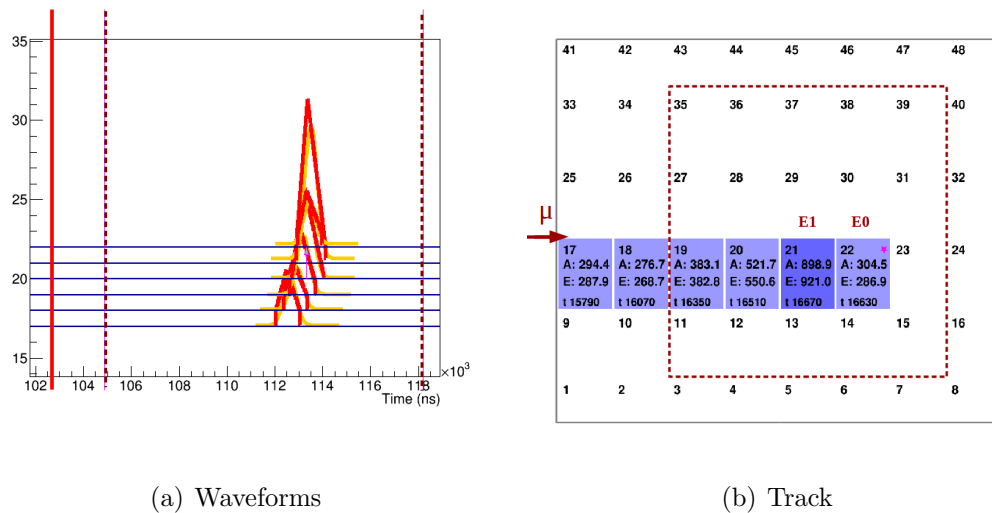
In the second analysis stage, pulse clusters are formed from TOT pulse objects which are spatially and temporally coincident. A pulse is included within a cluster if it satisfies the following conditions with respect to any of the other pulses in the cluster:

- $\Delta X \leq 1$
- $\Delta Z \leq 2$
- $\Delta t \leq 1000 \text{ ns}$

where ΔX and ΔZ are physical separation of pulses in pads and Δt is the time difference between either the beginning or end of adjacent pulses, whichever is smaller. A typical muon stop cluster from the R2014 data can be seen in figure 7.12. Figure 7.12(a) shows the time over threshold characterization of a typical pulse cluster in time (the Y dimension) for each pad. Each of the pulses satisfy the requirements listed above. The pulses for the 6 adjacent pads are displayed spatially in X and Z in 7.12(b), with an increase in amplitude and energy denoted on each pad until the muon comes to a stop in pad 22.

Stop Determination

A fiducial volume cut is made to ensure that the muon stopped within the target; any stop determined to be outside of the red dashed region in figure 7.12 is rejected. As such, correct determination of the muon stop position in the presence of muon-catalyzed fusion products is crucial. Events which are artificially accepted or rejected due to mis-reconstructions from fusion products lead to a distortion of the lifetime measurement, discussed in detail in chapter 10. Several tracking algorithms, which determine the muon stop position from the TPC pulse cluster objects displayed in 7.12, have been developed by Michael Murray, and extensive studies can be found in his thesis[38]. Upgrades have been made to the algorithms based on the detector characterization and Monte Carlo studies presented in section 5.6. However, as



(a) Waveforms

(b) Track

Figure 7.12: Typical event display showing (a) TPC waveforms and (b) X and Y pulse distribution for a muon stop in the TPC. The solid red line indicates the muon entrance time. Dashed red lines denote the fiducial volume region.

the tracking algorithms are already discussed at length in [38], only the general framework of three tracking algorithms used to analyze the R2014 dataset are presented here.

Basic Muon Tracking Algorithm The basic muon tracking algorithm takes the most downstream pad in Z as the muon stop position for all events. As discussed in section 5.6, this results in mis-reconstruction of the muon stop for all muon catalyzed fusion events in which the proton travels downstream. The Y position is determined by the latest charge arrival time, and the X stop position is taken to be the X pad on the last Z row with the highest energy. The utility of the basic tracker is its transparency. The general behavior and migration of this tracker is well defined and easy to interpret. As such, this tracker is used to study muon catalyzed fusion effects and aid in the development in more advanced tracking algorithms discussed below.

Road Muon Tracking Algorithm The tracker used primarily for the purpose of the R2014 analysis is the road tracking algorithm, which evades use of information near the stop pad to reduce effects from muon catalyzed fusion interference. The general concept is to distinguish the “road” consisting of all pads leading to the stop, from stop pads which may contain fusion products. As the muon deposits its energy characteristically, as discussed in section 5.6, there is a specific range of energies deposited on the most downstream pad, and each pad upstream, depending on the stop depth into the last pad. The range of energies for E0 and E1, also defined in section 5.6, are displayed in figure 7.13. As there is a maximum of 1100 keV in E1 for a clean muon stop, the road tracker sets a threshold at this value and determines all preceding pulses to be the “road cluster”. The first pad row in Z to go over the threshold is determined to be the stop pad. The road cluster pads are projected out to the stop in Z to determine the Y stop position. The X stop position is taken to be the X pad on the most downstream road cluster Z row with the highest energy. As upstream traveling protons can cause the track to go over the road threshold, the migration from fusion products is primarily upstream with this tracking algorithm, in converse to the basic tracker. The stop distributions in X , Y , and Z for the road tracker for each dataset can be seen in figure 7.14.

Proton Direction Tracking Algorithm The most advanced tracking algorithm developed attempts to distinguish between event topologies in which the proton is emitted upstream or downstream relative to the muon stop position. A cut based on the findings in section 5.6 is used to select events with a downstream going proton. The E0 vs E1 distribution along with the cut used to select the downstream events can be seen for both data and Monte Carlo in figure 7.15. For the events tagged as downstream, the stop Z position is taken to be the second to last Z row. Otherwise, the last Z row is used. The road threshold defined above is once again used to determine the road cluster, defined to be independent of fusion products, and is used to project the Y stop position as before. The X stop position is




Figure 7.13: Energy distribution for the most downstream pad (E0) and one pad upstream (E1), for the R2014 dataset.

again taken from the highest pulse on the last pad in the road cluster. Although this tracking algorithm is the most robust in terms of mitigating fusion interference, migrations may occur in both the upstream and downstream directions, making the task of quantifying its effectiveness extremely difficult. As such, the methods used to develop a fusion interference correction are primarily developed with the road tracking algorithm described above, and are only later applied to the proton direction tracking algorithm.

7.4 *Electron Definition*

7.4.1 *Electron Tracking with the Multi-Wire Proportional Chambers*

Two cylindrical electron multi-wire proportional chambers (ePCs) located concentrically around the TPC and inside of the electron scintillator hodoscope can be used to determine electron track information. If a cut is made on the distance of closest approach of the electron track to the muon stop position – an impact parameter cut – electron backgrounds

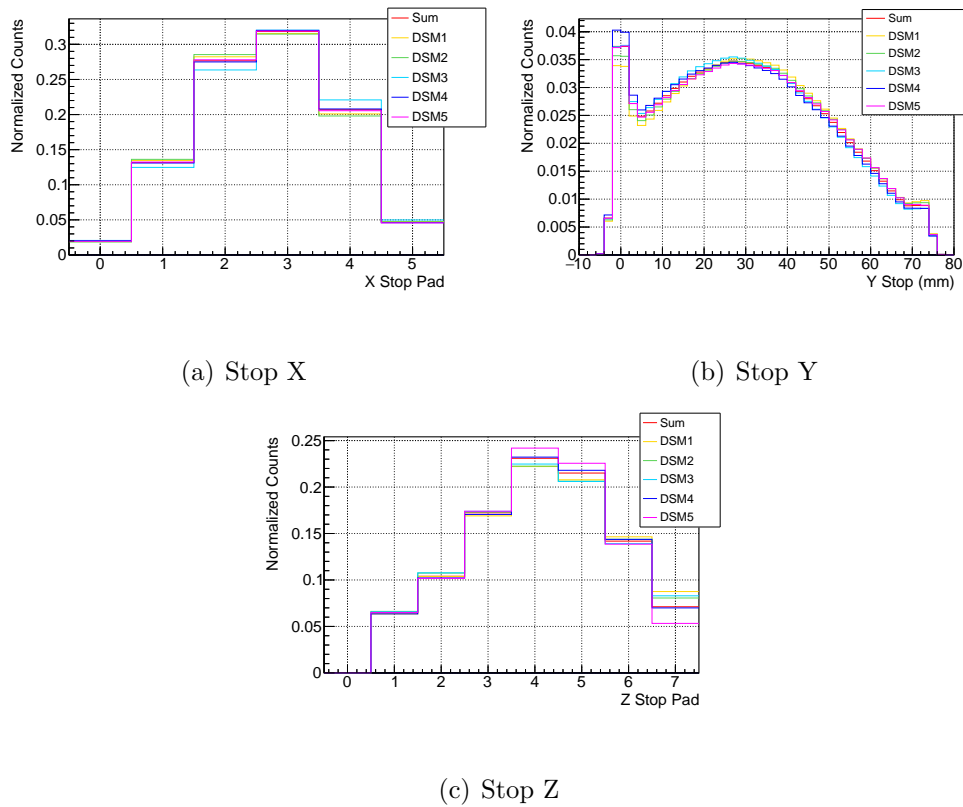


Figure 7.14: Stop distributions in (a) X, (b) Y, and (c) Z using the road tracker for each sub-dataset. The X distribution shown in (a) is asymmetric, indicating the muon beam was not centered in the TPC. The well defined peak at $Y=0$ mm in (b) results from muons stopping within the anode pad plane. The range of muons within the target gas is 80 mm, or roughly 5 Z pads, and some fraction of the stops in Z row 7 shown in (c) are punch-through muons.

can be suppressed. The Z , θ , and ϕ , as defined in the simplified diagram in figure 7.16, are determined from the location of the charge collection in both cylindrical wire chambers. Each ePC consists of anode wires, strung along the length of the chamber parallel to the beam axis, interstitial to two cathode layers of strips running with left and right helical trajectories such that a specific location in Z and ϕ can be determined via readout of the anode and cathode layers. Tracks are formed with the following procedure, first developed

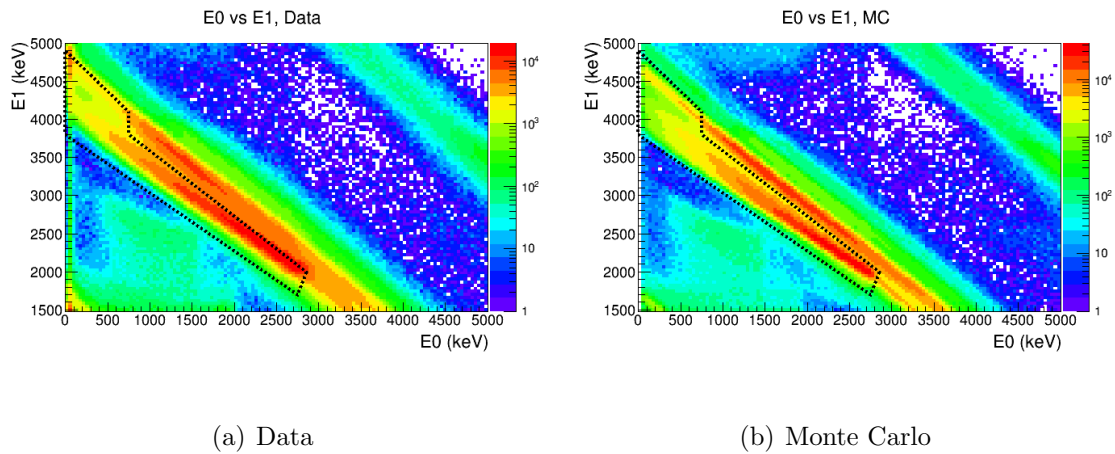


Figure 7.15: The E0 (most downstream energy) vs E1 (energy on one pad upstream) for both (a) data and (b) Monte Carlo. The cut used by the proton direction tracker to select downstream going proton events is shown in each case with the dashed black line.




Figure 7.16: Cartoon illustrating the coordinate system used for tracking the electron through the cylindrical wire chambers.

in the MuCap experiment:

1. A deadtime is applied to each individual wire/strip to mitigate the effects of afterpulsing

2. Clusters in each anode and cathode plane are created with temporally and spatially coincident hits
3. The anode clusters are joined with temporally coincidence inner cathode clusters in each ePC
4. The anode clusters are joined with temporally coincidence outer cathode clusters in each ePC1
5. The three layers of each chamber are joined with a temporal and angular (ϕ) coincidence window
6. A Z cut is used to select appropriate three layer coincidences from each ePC1
7. A time and ϕ cut is used to join the three layer ePC1 and ePC2 clusters

Tracks are considered if a coincidence exists between the anode plane and one of the two cathode planes in both chambers. This definition is referred to as *Cathode OR*. Full tracks are created by combing the track information listed above with a temporally coincident scintillator object, described in the next section. The ϕ , θ , and Z information are determined by the spatial geometry of the ePC tracks, while the scintillator is used for the timing information of each track.

A high frequency pickup presided in the chamber electronics for the majority of the data collection period. This produces oscillations in the detector hit times of each ePC sector, evident in the autocorrelation histogram shown in figure 7.17. In addition, the noise in certain sections sometimes overwhelmed the data acquisition, resulting in the need to completely mask certain wires. In initial analyses which include the ePCs in the electron definition, the fitted decay rate is highly unstable as a function of the start time of the fit, and extremely sensitive to the choice of impact parameter. Ongoing studies are attempting to understand these effects, but for the purpose of this thesis the chambers are removed from the standard analysis. However, the tracking is useful in characterizing the beam electron background, discussed in section 9.2, and is used in this respect.





Figure 7.17: Autocorrelation of hits observed in the ePC outer cathode electronics.

7.4.2 Simplified Scintillator-Only Electron Definition

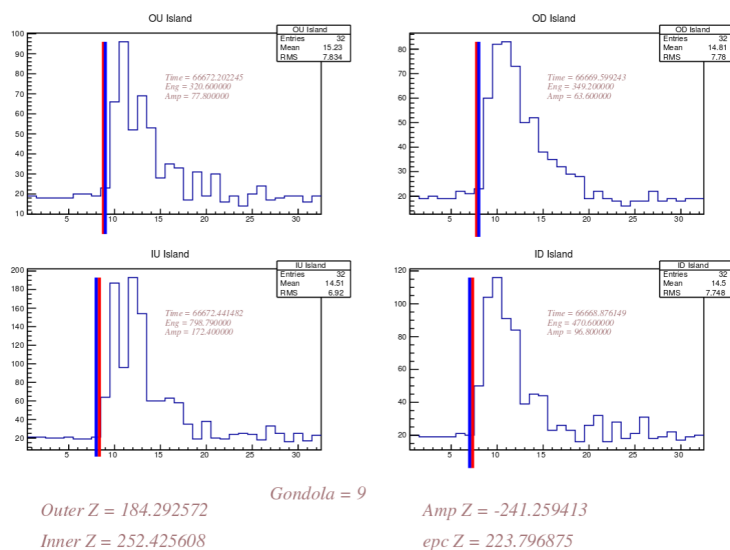
Several issues with the electron wire chambers motivated a simplified electron definition consisting only of information from the scintillators, developed by the author for the R2014 analysis. The drawback of removing the chambers from the analysis is an increased sensitivity to electron backgrounds, as no real impact parameter cut can be applied. A single eSC segment consists of two layers of scintillator, each with a photomultiplier tube on both the upstream and downstream end, as shown in figure 7.18(a), allowing for determination of the time and Z position within the scintillators, but no directional information.

Pulse Processing Electron scintillator objects are created from the digitized eSC pulses, consisting of the leading edge time of each pulse in coincidence. A constant fraction of $1/4$ of the pulse height is used to determine the time of each pulse. A typical event for all four tubes (inner upstream (IU), inner downstream (ID), outer upstream (OU), and outer downstream (OD)) of an eSC segment can be seen in figure 7.18(b). The time, amplitude, and integral of

each pulse is stored as a pulse object for each tube. However, time of each pulse must first be corrected, due to an ambiguity in the recorded WFD time.



(a) eSC Segment Diagram



(b) Waveform Display

Figure 7.18: (a) Simplified diagram of a single eSC segment and (b) typical eSC four fold event In a single eSC segment. The red band shows the pulse time as determined by the constant fraction method, the blue from the time over constant threshold method.

The eSC waveforms are recorded with custom built WFDs, operated at a sampling fre-

quency of 450 MHz. The WFDs begin recording data at the DAQ block start signal, and record 4 samples at a time. Due to a latching mechanism, the overall time of a WFD channel can be off by 4 clock ticks (8.88 ns) on a block by block basis. The inner-outer and upstream-downstream timing correlations for a single eSC segment displaying the ambiguity can be seen in figure 7.19. The highest timing resolution occurs between inner and outer eSC layers, shown in the top two histograms. The upstream-downstream correlations shown in the remaining panels are smeared by the Z dependence of the charge arrival time; electrons which pass through the scintillator more upstream in Z are seen earlier in the upstream eSC tubes and vis versa. There is a clear 9 ns jump between inner and outer coincidences on a block by block basis for the inner-outer correlations. In this particular run file, there was no clock slip between the upstream and downstream tubes. This shift in electron times will have no

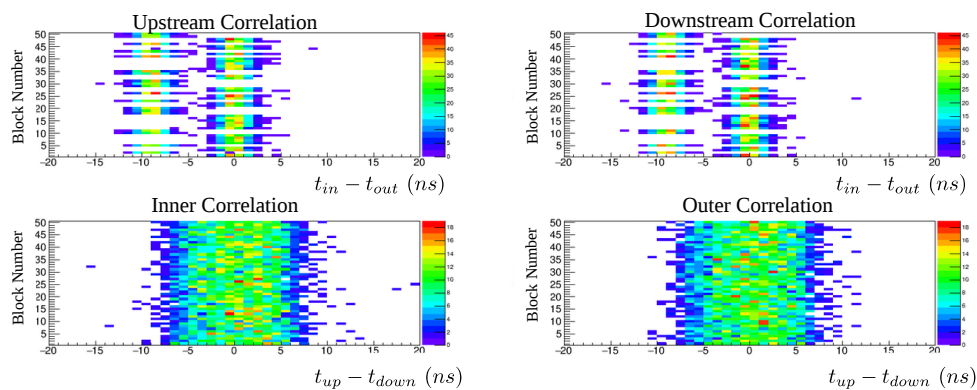


Figure 7.19: Block by block timing correlations for eSC Segment 1. The top two panels show the inner-outer correlations for the upstream (left) and downstream (right) tubes. The lower panels show the upstream-downstream correlations for the inner (left) and outer (right) eSC layers. In this particular run file, there is a clock slip between the inner and outer eSC channels.

effect on the lifetime as long as the fit starts after 9 ns. However, an accurate relative timing

between the upstream and downstream PMTs is needed in order to extract the Z position of electrons using a time of flight method. The timing is always off by exactly four clock ticks for the duration of a block, which allows for a straight-forward block by block correction. A new method was developed to correct for the relative timing ambiguities, in which correlation histograms are created on a block by block basis. The correlation histograms are fit at the end of each block, and if a mean greater than 7 ns is found, all times are shifted by four clock ticks, relative to a single tube. At the end of the block, the histogram is cleared and the process repeats. The timing correlation histograms can be seen for the first 500 blocks with this correction applied in figure 7.20.

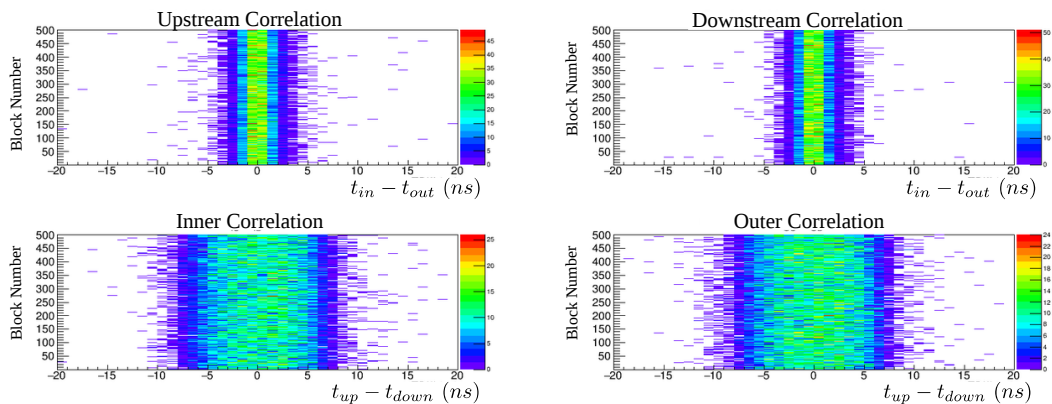


Figure 7.20: Block by block timing correlations for the first eSC segment for the same run file as displayed in figure 7.19, after application of the clock tick ambiguity correction.

N-Fold Coincidence Once the raw pulses have been processed and the clock tick ambiguity corrected, a 100 ns deadtime is applied to each tube to mitigate after pulsing effects observed in the raw autocorrelation displayed in figure 7.21. After the application of the deadtime, temporally coincident IU-ID pulses and OU-OD pulses are made into inner and




Figure 7.21: Autocorrelation of the raw eSC WFD pulse times before (blue) and after (red) application of the 100 ns deadtime, to mitigate the effects of after pulsing the eSC PMTs.

outer two-fold objects, respectively. A coincidence window of ± 30 ns is used to determine coincidences within a scintillator panel, as the scintillation light can travel as far as 900 mm for electrons which travel through the hodoscope at extreme values of Z . Next, a coincidence window of ± 20 ns is used to join the inner and outer two fold objects to create the final eSC 4-fold objects used in the analysis. Although only events with pulses in all four PMTs of an eSC segment are considered in the final analysis, 3-fold objects are created for the purpose of efficiency studies. The final eSC object contains the pulse time, amplitude, and integral information from all four tubes within an eSC panel.

As afterpulsing is a potential source of systematic error within the electron scintillators, auto correlations are created for for each level of eSC object. All electron autocorrelation histograms are created with electrons observed within the first 600 ns of the muon entrance time, to reduce any effects from the electron background, and are normalized to the number of counts at time zero. The single hit autocorrelations all four tubes of the eighth eSC segment, shown in figure 7.22, demonstrate significant afterpulsing effects. The eighth segment is displayed because it has the most prominent afterpulsing effects; all other segments show similar or less pronounced features. Autocorrelation histograms created from the times of the




Figure 7.22: Autocorrelation of the raw eSC WFD pulse times for all tubes of eSC segment 1, normalized to the number of counts at time zero. All 16 segments display similar or less pronounced features.

inner and outer two-fold objects, as well as the final four-fold electron objects from the eighth eSC segment can be seen in figure 7.23. Most features seen in the single hit correlations are greatly reduced, although still somewhat present in the two-fold correlations. All features are nearly eliminated once a four-fold correlation is required. Any residual effects on the disappearance rate from afterpulsing are studied in more detail in section 9.2.5.

Z Position from Time of Flight The Z position of the electron is derived from the time of all four signals within a 4-fold using a time of flight method,

$$Z = \frac{V_{eff}}{2}(t_u - t_d) \quad (7.1)$$

where $V_{eff} = 14.2 \text{ cm/ns}$ is the effective velocity of light within the scintillator, and t_u and t_d are the times the signal is observed at the upstream and downstream tubes, respectively.


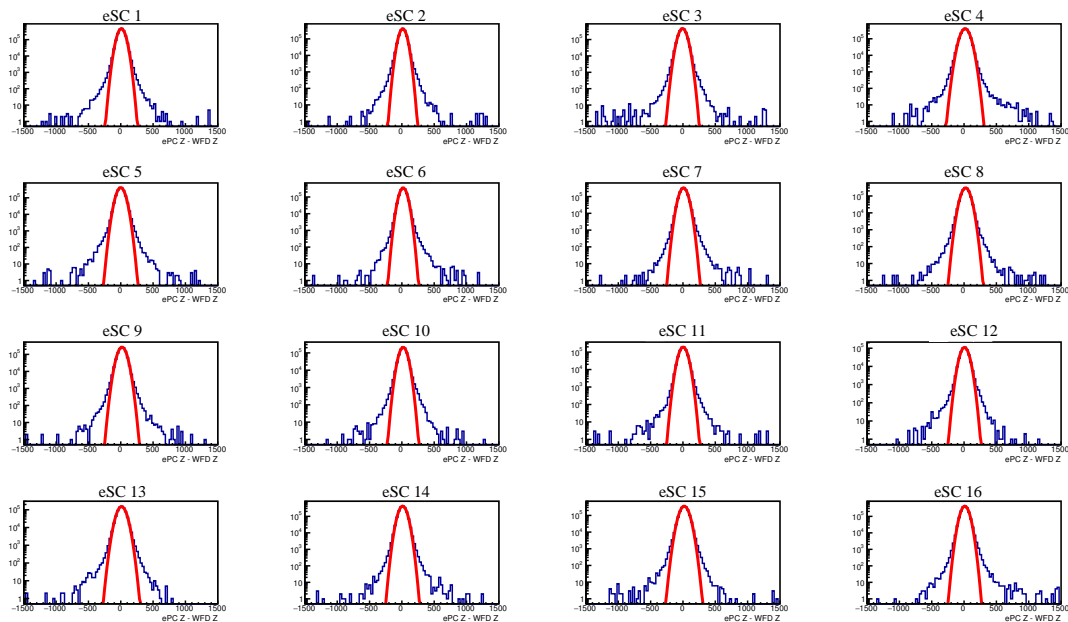



Figure 7.23: Autocorrelation of the inner two-fold, outer two-fold, and four fold electron objects for eSC segment 8. The features seen in the single hit autocorrelations are reduced in the two-fold correlations and nearly completely suppressed in the four-fold correlation.

The high spatial resolution tracks from the cylindrical proportional chambers are used to calibrate V_{eff} . An attenuation length of 126 cm was determined, in good agreement with the literature value of 140 cm or BC404. The Z position from the time of flight in the scintillator is compared to the Z position obtained from projecting the ePC track out to the scintillator hodoscope in order to determine the Z resolution of the time of flight method. The difference between the ePC projected Z position and the eSC derived Z position, along with the fitted resolution for each eSC segment can be seen in figure 7.24. An average Z resolution of roughly 50 mm is obtained using the time of flight method. As pulses from events which pass through the scintillator at extreme values of Z must travel the length of the scintillator before reaching the opposing PMT, the resolution is found to vary with the electron Z position. The difference between the ePC track position and the obtained Z position from the time of flight for several ranges of ePC Z can be seen in figure 7.25, as well as the resolution obtained for each Z range. As expected, the resolution is highest in the



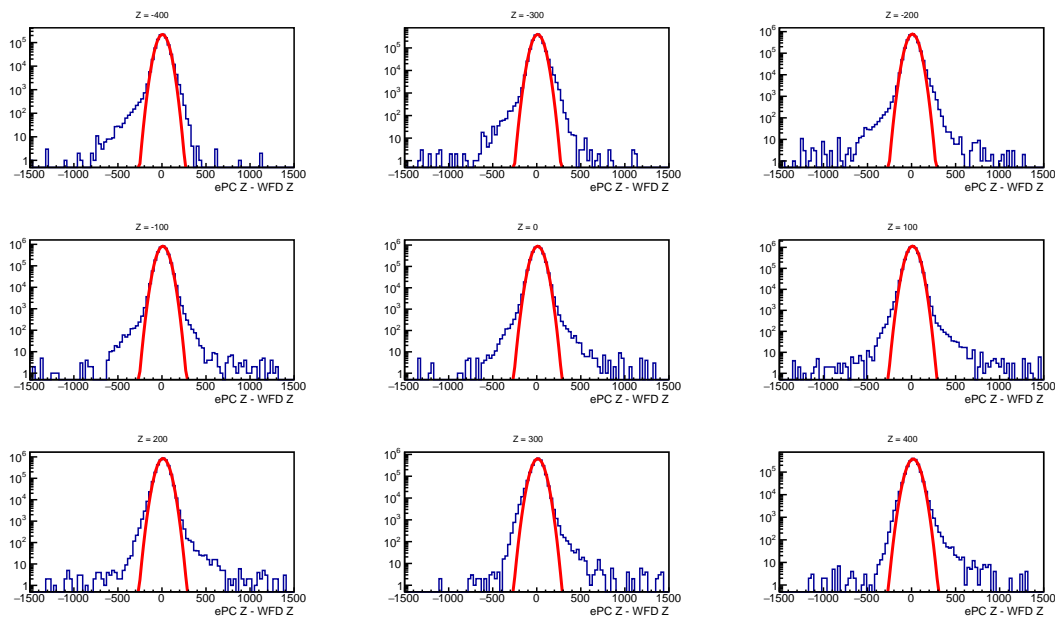
(a)




(b)

Figure 7.24: Panel (a) shows $Z_{ePC} - Z_{eSC}$ for each eSC section, ranging from segment 1 in the upper left corner to segment 16 in the bottom right corner. Panel (b) shows the resolution from the RMS (blue) and the fit sigma (green) for each of the fits shown in the left panel.

center of the eSC segment.



(a)



(b)

Figure 7.25: Panel (a) shows $Z_{ePC} - Z_{eSC}$ for slices of the Z position, as determined by the ePC, ranging from -400 mm in the upper left corner to 400 mm in the bottom right corner. Panel (b) shows the resolution from the RMS (blue) and the fit sigma (green) for each of the fits to Z slices shown in the left panel.

7.5 *Trend Histograms and Quality Cuts*

A preliminary MU pass is used to determine all calibrations and timing offsets. Trend histograms are created from each MU file in this preliminary pass to determine regions of stability, and set run by run values into a database. This database is accessed in the first step of MU to retrieve all necessary calibrations and offsets. Additionally, a set of trend histograms is created with the final data pass in order to assess data quality over the collection period. Several cuts are made based on the standard rates and values seen in various trends. If any single file fails to pass one of the quality cuts, it is removed from the final analysis.

7.5.1 *Initial Quality Cuts*

Prior to the final analysis, several basic quality cuts are applied to all data:

- Total blocks analyzed > 800 blocks
- Average block length > 5 ms
- Unkicked beam rate > 45 kHz
- Muon Stop rate > 8 kHz

The cut on the number of blocks analyzed and average block length eliminates the runs in which electron wire chamber noise saturated the data acquisition system (DAQ), causing low rates in all other detectors. The runs with low beam rate can lead to calibration issues in both the electron timing and TPC baseline determination, and have thus also been removed with the unkicked beam rate cuts. Trend histograms with the applied cuts represented as red lines can be seen in figure 7.26. Overall, 1.3% of muon entrances are cut with all quality cuts.

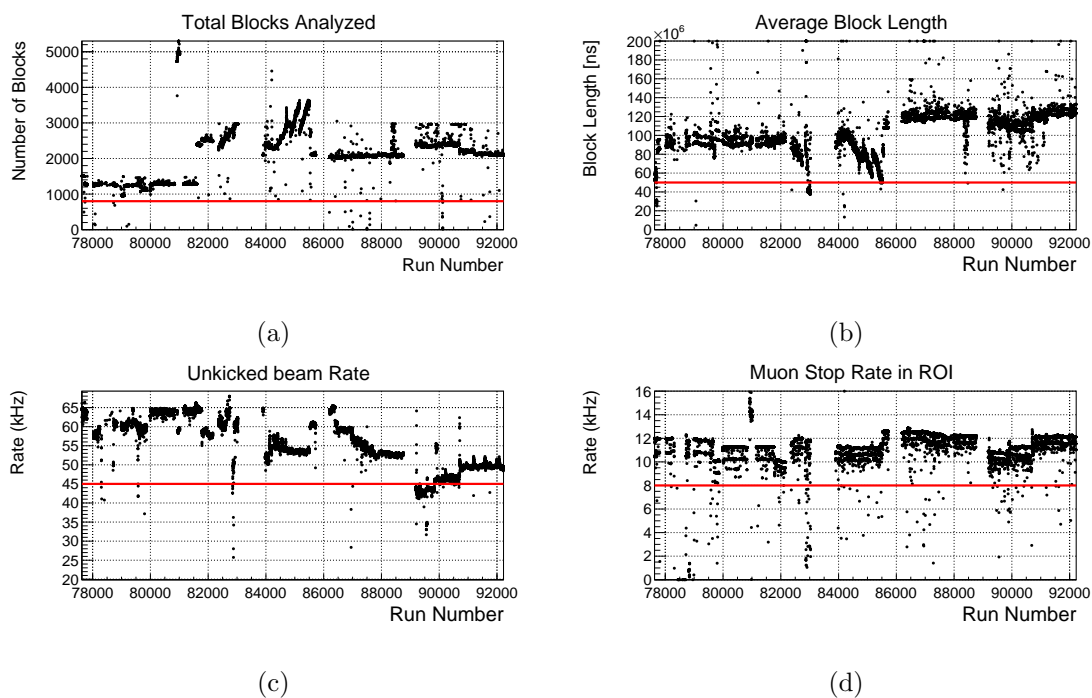


Figure 7.26: Trend histograms showing (a) number of data blocks analyzed, (b) average block length, (c) unkicked beam rate, and (d) muon stop rate in the region of interest for each run collected in the R2014 dataset. The red lines indicate basic quality cuts which must be satisfied for each file to be considered in the final analysis.

7.5.2 Error Conditions

Several error conditions are recorded to monitor data quality during acquisition. Flags are set to denote bad events and blocks, which are cut from the data in the later stages of the analysis.

MuSC Errors

As the timing precision of the experiment relies heavily on obtaining a correct time for the muon entrance, several redundancy checks are in place for quality assurance. The muSC signal is fed into two separate TDC channels as a redundancy check. Anytime one of the

MuSC signals are missing, or the two times do not agree, a **MuSC Mismatch** error flag is set. A small percentage of the data is affected by this error. As a second check on the timing, the muSC signals are routed sequentially to one of four TDC channels. If the routed signals are ever missing or out of order, a **MuSC Router** error flag is set. The fraction of blocks with MuSC errors in each data run over the course of the collection period can be seen in figure 7.27.




Figure 7.27: The fraction of data blocks in each run file containing MuSC related errors.

CAEN Errors The CAEN TDCs used to record time stamps from the muSC and muPC channels have several built in error checks. The most frequent CAEN error seen in the R2014 data is a **Trailing Edge** error, where the signal triggers on the trailing edge instead of the intended rising edge trigger. The TDCs are configured to only trigger on leading edges, so there should never be any instances of trailing edge triggers. Nonetheless, these instances are observed within the R2014 dataset. A flag is set for events with a large number of trailing edges, but it was found that cutting on these events leads to a lifetime distortion. As

such, these errors are not cut from the final data selection. The CAEN TDCS function by counting for a finite range, then rolling over and starting back at the beginning. The time of the rollover is recorded for each TDC, and if the rollover times do not agree, a **Rollover** error flag is set. The fraction of blocks in each data run over the collection period with a CAEN related error can be seen in figure 7.28.




Figure 7.28: The fraction of data blocks in each run file containing CAEN related errors.

Compressor Errors The custom built compressors which record times of each wire in the cylindrical proportional chambers also have a set of error conditions which are checked. If any of the conditions dictate an error, a **Compressor** error flag is set. These errors typically occur with abnormally high beam rates, and affect a small fraction of the data. The fraction of blocks in each data run over the collection period with a compressor related error can be seen in figure 7.29.




Figure 7.29: The fraction of data blocks in each run file containing errors related to the eSC compressors.

WFD Bank Errors Custom built waveform digitizers (WFDs) are used to record outputs from the TPC, eSC, and muSC. Before processing the banks at the first stage of the MuSun analysis, several error conditions are checked for. In particular, an error in the TPC WFD persisted throughout the run, in which the time stamp of WFD banks is incorrectly assigned, possibly due to a clock slip error. This results in a missing TPC pad for the analysis. As discussed later in the chapter, two WFDs with different gains are employed to digitize the TPC data and flags are set to denote in which blocks this error occurs for high gain (HG) and low gain (LG) channels. The percentage of data blocks within each dataset with a TPC WFD error for both gains can be seen in table 7.1. The number of TPC WFD errors per pad, normalized to the number of waveforms recorded in each pad, can be seen in figure 7.30(a) and the fraction of blocks with TPC WFD errors over the course of the run can be seen in figure 7.30(b). Some channels were more prone to errors during specific run periods. The majority of the errors observed are in DSM5 in LG pad 45.


Dataset	Blocks with HG Error	Blocks with LG Error
DSM1	5.89%	3.02%
DSM2	4.84%	0.56%
DSM3	6.54%	8.84%
DSM4	6.23%	4.51%
DSM5	8.10%	22.26%

Table 7.1: Percentage of data blocks analyzed with TPC WFD errors for each of the five datasets.


As a missing pad can lead to inefficiencies and track reconstruction issues, particularly related to the fusion interference effects presented in chapter 10, any event with a TPC WFD error in the high gain crate is rejected. The signals from the muSC and each eSC channel are also digitized by the custom built digitizers, and are vulnerable to the same errors. Although less critical, as the muSC time is primarily taken from the TDC, and missing electron signals would only decrease the overall detector efficiency without leading to any distortions of the observed lifetime, these errors are also flagged as **muSC WFD** and **ESC WFD** errors. The prevalence of these errors throughout the run can be seen in figure 7.31.

7.6 Lifetime Histograms and Fitting

The third stage of the MuSun analysis chain consists of the creation of trees containing all information to fill lifetime histograms. The standard lifetime histogram consists of the difference between the time of the electron object and the time of the muon entrance. A cut is made to select only muon entrances which satisfy the HBE condition, as described in section 7.2, and a cut is made to reject all events which stop outside of the fiducial volume.



(a)



(b)

Figure 7.30: Trend plots showing (a) the fraction of TPC WFD errors in each TPC pad and (b) the fraction of data blocks with TPC errors in each data run. The pad numbers in (a) are zero indexed, and the majority of the errors exist in pad 45, and occur mostly later in the R2014 data collection period, in DSM4 and DSM5.




Figure 7.31: The number of WFD Errors recorded in the muSC and eSC readout over the course of the run.

A simple exponential is used to fit for most analyses

$$N(t) = N\lambda e^{-\lambda t} + B \quad (7.2)$$

where t is the time difference between the muon entrance and the observed electron, although a full kinetic fit must be applied to the final result to account for the ${}^3\text{He}$ sticking losses.

7.6.1 Beam RF Oscillations

The stages of muon source provided by the High Intensity Proton Accelerators (HIPA) are discussed in more detail in section 4.2 . Because the RF beam structure is retained in both the muons and the beam electrons, an oscillation is observed in the final histogram of the time difference. Protons directed at the carbon target, which create positive, negative, and neutral pions, arrive in batches at the cyclotron RF frequency of 19.7 ns. Although the π^0 decays primarily via $\pi^0 \rightarrow \gamma\gamma$, the processes $\pi^0 \rightarrow e^+e^-\gamma$ is also permitted, albeit suppressed by a factor of 100. This results in a beam electron background. As the $8.52 \times 10^{-17} \text{ s } \pi^0$

lifetime is much smaller than the cyclotron frequency, the beam RF structure is retained by the electrons and seen in the detector systems. The muons created from the much longer lived π^- decay would be expected to smear the RF structure. However, as only pions which decay before the first quadrupole magnet, roughly 30cm after the target, are selected, this beam RF structure is preserved in the muons as well. With a combination of the 19.7 ns accelerator RF and the 40 ns binning required by the TDCs, a distinct oscillation of 26 ns is seen in the background of the lifetime fits. This oscillation, seen distinctly at late times, has little impact on the fitted rate, but wildly increases the chi-squared of the fit. Two attempts are made by the author to characterize these oscillations and recover a good chi-square; fitting the oscillations and smearing the muon times by the oscillation period.

Fitting the oscillations As a first attempt to recover a good χ^2 for the lifetime fits, an oscillatory term is added to the background in the fit functions. To get the oscillation period due to the beam electrons, the correlation between the MuSC and the MuSC-Lo is fit with a simple sinusoid. The MuSC-Lo is independent from the electron detectors, thus providing an unbiased oscillation period for the beam electron background. This fit can be seen in figure 7.32. The functional form of the lifetime fit which accounts for the RF oscillations is





Figure 7.32: Oscillatory fit to the MuSCLo time distribution used to extract the RF period with 40ns binning.

then given by


$$N(t) = N\lambda e^{-\lambda t} + A \sin(\omega t + \phi) + B \quad (7.3)$$

where ω is fixed to the oscillation frequency found from the MuSC Lo fit, and A , ϕ , and B are floating fit parameters. B is the normal background fit parameter, while A and ϕ characterize the amplitude and phase of the beam electron oscillations, respectively. A comparison of the lifetime fit parameters and residuals for the nominal exponential fit and the fit including the oscillatory term can be seen in figure 7.33.

Smearing the Muon Time As an alternative to fitting the RF oscillations, the muon time is smeared by $\pm T/2$, where T is the RF period as determined by the MuSC-MuSC Lo fit seen in figure 7.32. The nominal three parameter fit to the smeared data is displayed in figure 7.34. In order to ensure the muon smearing procedure has no effects on the fitted rate, a scan of the RF period used for smearing is performed. This scan over the various fit parameters can be seen in figure 7.35. The first bin shows the results of the fit to unsmeared data. The fit to the data smeared by the RF period as determined by the MuSC-MuSC Lo fit is in the second bin of the scan. Subsequent bins show fit results for data smeared by values $\pm 10\%$ of the RF value. Statistical error bars are shown, but the data is completely correlated, so even small deviations with different smearing periods would be significant. The stability of the fit demonstrates that the smearing method can be applied to the data without introducing any lifetime shifts, even if the accelerator RF is only known to $\pm 10\%$. Although both methods yield consistent results, the muon time is smeared in the final analysis to simplify the lifetime fitting procedure. All lifetime histograms displayed in this thesis have been smeared in this fashion.



(a) Exp Fit



(b) Exp + Osc Fit

Figure 7.33: Lifetime fits to the full R2014 data with (a) the nominal three parameter exponential fit and (b) the nominal fit plus an oscillatory background.




Figure 7.34: Three parameter fit to the muon lifetime histogram created with muon times smeared by the accelerator RF period.

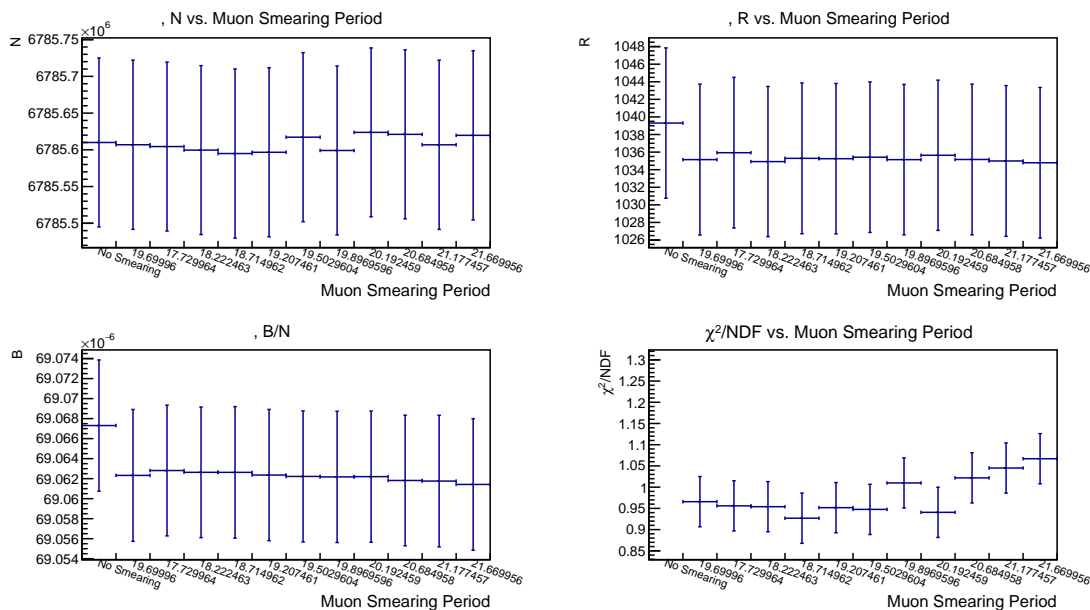


Figure 7.35: Scan of the lifetime fit parameters versus the period used to smear muon time, from $\pm 10\%$ of the fitted accelerator period. Statistical error bars are shown, although the data is completely correlated, so small deviations in the fitted rate would be significant.

Chapter 8

SYSTEMATICS

In order to achieve the precision goal of the MuSun experiment, every possible source of systematic error must be considered. Outlined below are brief explanations of all contributing sources of error, and first estimates of the effect of each on the observed disappearance rate. Two particular sources of error related to the Muon-On-Request scheme and muon-catalyzed fusion interference are studied in detail by the author and are presented more thoroughly in subsequent chapters.

8.1 Muon-On-Request scheme

The precision of the MuSun experiment relies on the ability to obtain events in which a single muon is observed in the detector system within a $\pm 25 \mu s$ time window. This is in part achieved with the first upstream beamline element encountered by incoming muons, the muon-on-request (MORE) signal [10], discussed briefly in section 4.2. Once a muon is observed in the entrance detectors, a pair of parallel plates are quickly charged to deflect the incident beam, reducing the probability of pileup events. Any event in which a second muon is observed in any of the entrance detectors within $\pm 25 \mu s$ is rejected to exclude any remaining pileup events. Two important systematic effects relating to the MORE signal are discussed in detail in chapter 9 of this thesis: unrejected pileup events due to entrance detector inefficiencies and time dependent backgrounds caused by the MORE deflection.

8.1.1 Entrance Efficiency

The inefficiency of each entrance detector is quantified in section 9.1 using three methods. First, a dedicated muon clock shift is analyzed to tag any event where a muon is seen within the TPC after full pileup protection has been applied. Second, the in-situ muon clock is used in a similar fashion, to obtain the inefficiencies for each dataset in which the μ CL signal was continuously running. Finally a method to quantify the efficiency with the data under nominal running conditions is developed to quantify each dataset. A zero extrapolation procedure is used to obtain the overall correction to the lifetime from any detector inefficiencies.

8.1.2 Electron Background

The deflection resulting from the MORE signal has a small, albeit non-zero effect on the electrons within the incoming beam, leading to a time dependent accidental structure. This effect is observed directly in the μ CL event selection, and a novel technique to enhance or suppress the background by selecting events with specific electron multiplicities is detailed in section 9.2. A background model is developed based on the observations from the entrance detector studies in combination with the electron multiplicity studies. Additionally, a comparison to the MuSun Monte Carlo simulation is presented, revealing an unexpected background contribution. The magnitude and effect on the fitted observed decay rate of several identified background sources is quantified and summarized in the final section of this chapter.

8.2 Muon Track Interference

Correct determination of the muon stop position is required to assert that each muon stops within the gaseous volume. To ensure that any event which may have stopped in surrounding materials is excluded from the final analysis, a fiducial volume cut is made. Thus, any time

dependent process which promotes mis-reconstruction of the muon stop position can lead to distortions in the observed decay distribution. This tracking interference is primarily instigated by the extra deposition of energy from two sources; muon catalyzed fusion products and decay electrons.

8.2.1 Fusion Interference

As previously discussed in section 3.2.2, a resonant formation of $dd\mu$ molecules within the target gas quickly results in muon catalyzed fusion events, resulting in either a neutron and ${}^3\text{He}$, or a proton and triton. As the muon is recycled, these events have no direct effect on the measured muon lifetime. However, observation of the fusion products in the TPC can lead to mis-reconstructions of the muon stop position, leading to a time dependent fiducial volume cut.

The energies and ranges of the various fusion products at the running conditions of the target gas are shown in table 5.2 of section 5.6. The majority of the systematic error stems from p-t mis-reconstructions, as the shorter range of all other fusion products imposes a small likelihood of travel outside the muon stop pad.

An example of a Monte Carlo event with upstream fusion migration can be seen in figure 8.1. Figure 8.1(b) shows the amplitude and energy increasing from pads 17 to 24. The true muon stop is on pad 24, outside of the fiducial volume, and a proton is emitted upstream, depositing energy on pad 23. Due to the presence of the proton, the energy on pad 23 is pushed over the threshold of the Maximum E1 for the stop threshold tracker, such that the reconstructed stop pad is 23. Thus, an upstream migration due to p-t fusion products has caused the acceptance of an event which should have been rejected. The same migration can occur on the upstream side of the fiducial volume region, leading to the rejection of an event with a true stop in pad row 2.

The frequency and direction of migration is dependent on the properties of the tracking

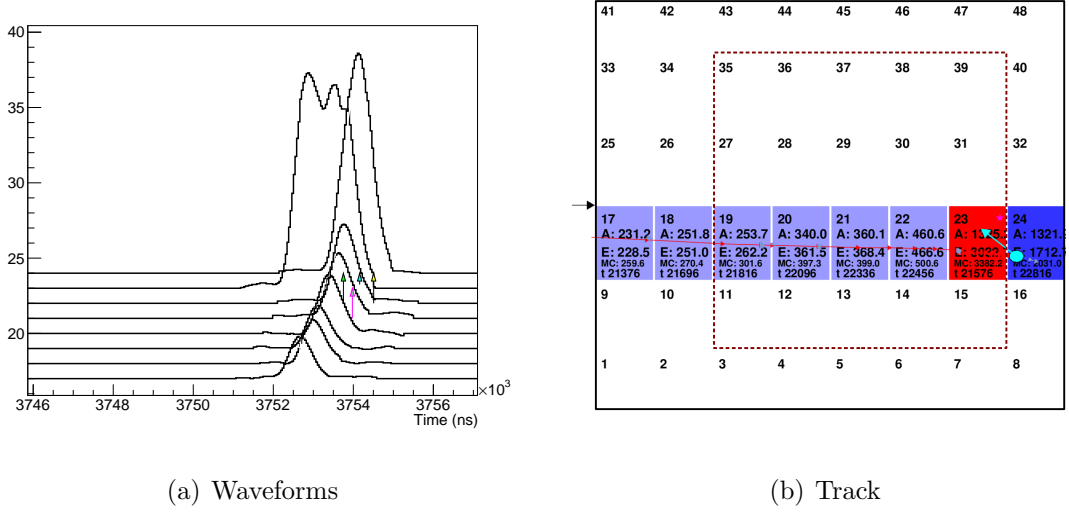


Figure 8.1: (a) TPC waveforms and (b) X and Z pulse distribution for a mis-reconstructed track resulting in an event migration.

algorithm. As illustrated in the example above, the stop threshold tracking algorithm only allows for migration in the upstream direction, as protons emitted downstream do not affect the E1 pad energy. The basic tracking algorithm, which always chooses the most downstream pad as the z stop position, has the opposite behavior, as any event with a downstream proton will cause the stop to be mis-reconstructed upstream.

As the muon must live long enough to catalyze a fusion event, the decay distribution of muons with a prior fusion is non-exponential. Any artificial acceptance or rejection of events due to the presence of fusion products leads to a distortion of the measured lifetime. Several strategies for quantifying the shift observed are presented in chapter 10. A data driven method shapes the muon stop distribution using spatial and temporal cuts on the muon entrance into the TPC to suppress events at the fiducial volume boundaries, selecting a subset of data in which muon catalyzed fusion interference is nearly absent. The lifetime effect is quantified by comparing the distributions with and without the fusion effect, and a

correction is applied to the R2014 dataset.

8.2.2 *Electron Interference*

There are two possible electron sources which could contribute to track mis-reconstructions, direction ionization from the decay electron, referred to as the Michel electron, and the secondary, energetic delta electrons. For the MuSun conditions, the Michel electrons are only expected to deposit up to 30 keV of energy per pad. Delta electrons, defined as electrons emitted with enough energy to create further ions, are modeled by an effective range-energy relation, accounting for the effect of multiple collisions, given by

$$R_p(E) = 7.67E^{1.72} \quad (8.1)$$

for the MuSun gas conditions [43]. This relation implies 100-150 keV can be deposited by delta electrons over the 16 mm pads of the MuSun TPC.

As a reminder to the reader, the standard road tracking algorithm, discussed in section 7.3.2, defined the stop position based on the first Z row in which the energy deposited is above the maximum E1 value of 1100 keV. There are thus two topologies in which the delta or Michel electron can cause a mis-reconstruction of the muon stop for the standard threshold algorithm:

1. For events which stop far enough into the last pad that the E1 energy is less than 100 keV from its maximum allowed value, the electron can deposit enough energy in E1 to push the pad over threshold, moving the reconstructed stop position upstream.
2. For events where the E1 threshold is never crossed, deposition of an electron with greater than 50 keV can trigger a downstream pad, leading to a mis-reconstruction of the stop in the downstream direction.

The extra energy from the electron does not change the acceptance of events within the fiducial volume. Thus, only events which move the stop position across a fiducial boundary are of concern.

Examples of two events with extra energy deposition from electrons can be seen in figure 8.2. In 8.2(a), the electron is emitted early enough that its ionization charge overlaps with the muon track. In 8.2(b), the electron is emitted late enough that its ionization charge cannot catch up to the drifting muon charge cloud, and the pulse from the electron is not clustered with the muon and therefore cannot interfere with tracking. As electrons must be emitted early enough to interfere with the muon track, artificial acceptance or rejection of these events leads to a time dependent effect.

In addition to the time dependence, certain event topologies are more likely to produce electron charge which interferes with the muon track. As displayed in figure 8.3, electrons which are emitted downwards are more likely to arrive at the anode plane at the same time as the muon drift cloud than upward traveling electrons. As the muon is traveling parallel to the drift field in each of these cases, the charge arrives over a time period, leading to an extended, smaller amplitude pulse. Conversely, if the electron is emitted perpendicular to the drift field, as displayed in 8.3(c), the charge is collected promptly, leading to a narrow, higher-amplitude pulse.

Due to the response of different event topologies, the effect of electron interference can be quantified by comparing the track energy in the TPC for electrons observed in various eSC segments either early or late relative to the muon entrance time. A diagram of the location of each of the eSC segments can be seen in figure 8.4. For the purpose of comparing the extra energy deposited from a Michel or delta electron, the S-Energy of each TPC track is determined. The S-Energy is defined based on a mapping in E0-E1 space, explored in greater detail in section 5.6.1, which obtains a constant value for all muon stops, regardless of the stop depth into the last pad row; $S = E0 + 2E1$. The S-Energy vs eSC segment vs

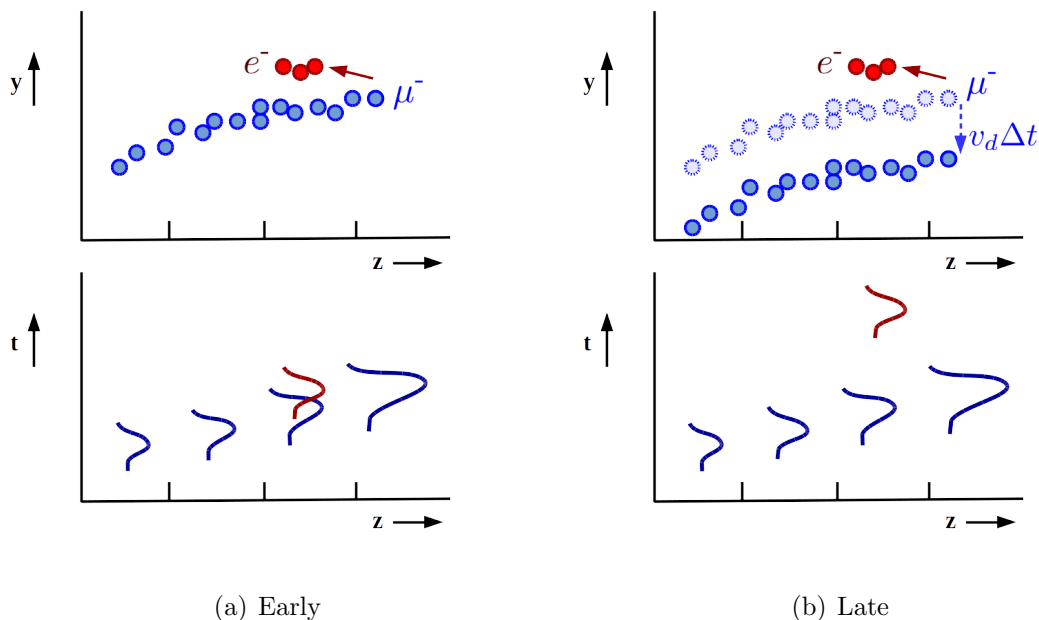


Figure 8.2: Charge arrival within the TPC for (a) early and (b) late electron decay. In the case where the electron is delayed, the dashed charge denotes the original location of the ionization from the muon stop. After a time Δt the charge has drifted by an amount proportional to the drift velocity, v_d , denoted by the solid circles. In this case, the pulse from the electron decay is well separated from the charge deposited from the muon, and unlikely to interfere with the tracking algorithm.

electron arrival time in the eSC for the full R2014 dataset can be seen in figure 8.5(a), and projections of the S-Energy vs eSC segment for early and late times are shown in 8.5(b).

For electrons observed early relative to the muon entrance time, the peak of the S-Energy distribution shifts by 0.3% in the lower eSC segments (5-12) relative to the upper eSC segments. As expected, electrons emitted downward are more likely to interfere with the muon track than upward traveling muons, and the effect is suppressed for electrons which arrive later in the eSC. A larger S-Energy shift from electrons in eSC segments 5-7 and 10-12 is also observed, resulting from events in which the electron is traveling perpendicular to the drift field. This effect is presumably due to the detector response of prompt charge arrival,

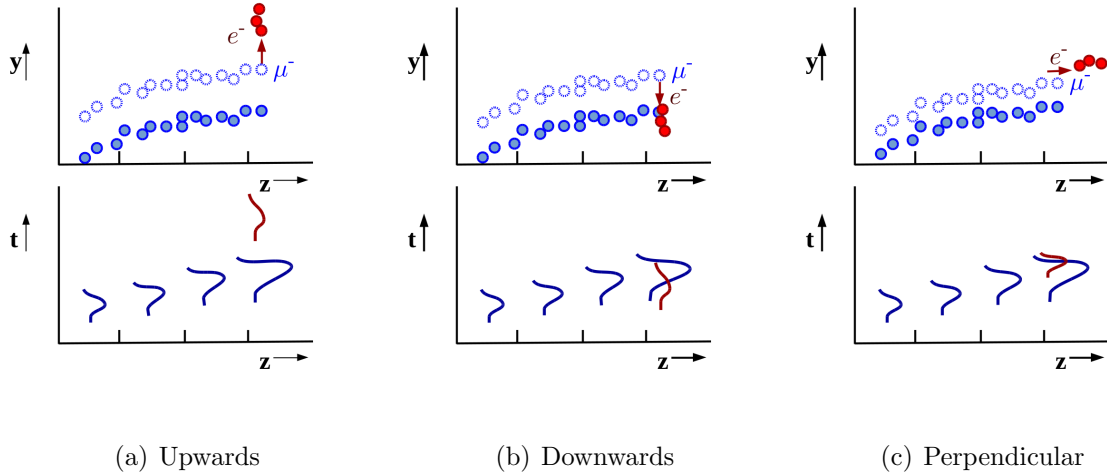


Figure 8.3: Charge arrival within the TPC for an electron emitted (a) upwards, (b) downwards, and (c) perpendicular to the drift field. In the case of (a) and (b) the electron is emitted parallel to the drift field, resulting in an extended charge arrival time and elongated, lower amplitude pulse. If the electron is emitted perpendicular to the drift field, as in (c), the charge is collected instantaneously resulting in a well peaked pulse.

yielding a sharper peak compared to electrons traveling parallel to the drift field.

The full Monte Carlo simulation mimicking the running conditions of the R2014 dataset must be used to further study the effects of electron interference. As a handle for ensuring that the Monte Carlo is accurately describing the complicated TPC response, the same S-Energy vs eSC segment vs electron arrival time in the eSC is applied to the MC data, as seen in figure 8.5(c). The same projection of the S-Energy versus eSC segment for early and late times can be seen in 8.5(d). Overall the effects of different event topologies are well captured in the simulated data. The peaks in eSC segments 7 and 11 due to prompt charge arrival are identical in magnitude to the data, and a similar time dependence is observed. However, the effect of downwards traveling electrons are larger in the MC by at least a factor of three, indicating that the integration for extended charge arrival is more efficient in the MC response than in the electronic response of the data.






Figure 8.4: Diagram of the eSC segment locations relative to the TPC drift field. Upwards going electrons will be detected in segments 1, 2, 15, and 16, whereas downwards traveling electrons will be observed in segments 7-10. Electrons emitted perpendicular to the drift field will be observed primarily in segments 4-5 and 12-13.

A scan of the lifetime fit parameters versus eSC segment can be used to obtain an upper limit on the effect of track mis-reconstructions due to electron interference on the fitted rate for the R2014 dataset. A scan of the parameters for the nominal exponential fit starting at $1 \mu s$ for various eSC segment combinations, selected to either enhance or suppress the effect based on the S-Energy scans, can be seen in figure 8.6. Although it is difficult to quantify any overall effect due to electron interference due to the reduced statistics in splitting the data for various electron topologies, a linear fit to the independent sets demonstrates a good chi squared, and a reasonable probability, indicating consistency between the various eSC combinations.


For the MC simulated data, the true stop position can be used to further quantify the effects of electron interference. A MC dataset consisting of 2.4×10^9 positive muon stops was simulated to eliminate effects of muon-catalyzed fusion effects. With the MC truth




(a) S-Energy vs Gondola vs Time R2014



(b) S-Energy vs Gondola R2014



(c) S-Energy vs Gondola vs Time MC



(d) S-Energy vs Gondola MC

Figure 8.5: (a) and (c) display the muon Track S-Energy vs eSC segment vs time of electron arrival in the eSC for 1 μ s time bins for the R2014 dataset and the MC, respectively. Two projections of this histogram for electrons emitted within 1-5 μ s of the muon entrance (blue) and 10-12 μ s of the muon entrance (red) are displayed in (b) and (d) for the R2014 data and the MC.

dataset, lifetime histograms can be created using either the reconstructed stop position or the MC truth stop position. A scan of the fitted rate for both stop definitions versus various eSC topologies can be seen in figure 8.7. The error bars shown are statistical, but as the data is highly correlated between the reconstructed value and the truth, any deviations are significant. The fitted rate of the bottom eSC segments, shifts by 3 Hz relative to the truth rate. As discussed above, this effect is much larger in the MC than the data, and the TPC




Figure 8.6: Fitted rate for the nominal exponential fit starting at $1 \mu\text{s}$ versus eSC segment for the full R2014 dataset.

response of the MC will be improved in the future for more realistic values. For now, this gives a good upper limit. The side eSC segments, which contain electrons traveling parallel to the drift field and are more accurately represented in the MC, yield a shift of 6-9 Hz. The effect is smallest in the top eSC Segments, as expected, and only a 1.5 Hz difference is observed between the fitted rates. Overall, the fitted rate obtained using the reconstructed stop position is 5 Hz lower than that obtained using the MC true stop position.

8.3 Muon Capture on Other Nuclei

8.3.1 Isotopic Purity

In the presence of an isotopic admixture of protium and deuterium, a $pd\mu$ molecule can form and undergo the fusion processes outlined in figure 8.8. This process, first observed by the Alvarez group in 1957 [5], has been studied extensively, and the reaction rates are given in




Figure 8.7: Fitted rate for the nominal exponential fit starting at $1 \mu\text{s}$ versus eSC segment for the MC positive muon dataset, shown for both the histogram created using the reconstructed and the MC truth stop information in the TPC.

table 8.1 [14]. The “Alvarez” muon emitted following a $pd\mu$ fusion is ejected with 5.3 MeV, corresponding to a range of ~ 26 cm in our chosen gas density. These muons will likely leave the deuterium volume, increasing the probability of capture on materials surrounding the target. A rough estimate of the magnitude of the effect can be made by resolving the kinetic equations including a protium concentration, as detailed in appendix C. For 100 ppm of protium, the fitted rate of the decay distribution increases by 5 Hz. Chromatography measurements taken during the R2014 data collection obtained values between 16-20 ppm of protium, which results in a 0.8 Hz shift in observed in the decay distribution.




Figure 8.8: Processes possible for an isotopic admixture of protium and deuterium gas.

Description	Symbol	Rate (s^{-1})
$pm\mu$ Molecular formation	$\lambda_{pd\mu}$	5.6×10^6
Total fusion	λ_f	0.3×10^6
Fusion, γ	$\lambda_{f\gamma}$	$\lambda_f \times 0.85$
Fusion, μ	$\lambda_{f\mu}$	$\lambda_f \times 0.15$

Table 8.1: Rates for $pd\mu$ molecular formation and fusion processes, from [14].

8.3.2 Chemical Purity

The effects of impurities within the target gas on the observed decay rate are discussed in section 3.3.2. In the presence of nitrogen, the most concerning impurity, the following reaction is possible



where C^* represents an excited state of carbon. As the capture rate on nitrogen is orders of magnitude larger than that on the deuteron, the impurity level in the TPC must be known

to the pbb level. Two techniques are used to quantify the concentration of impurities within the gas; an advanced chromatography measurement, and an in situ analysis utilizing nuclear recoil from capture events within the TPC.

Gas Chromatography In order to measure levels below 1 ppb with the chromatography, samples are taken directly from the target volume and allowed to accumulate on an adsorbent bed, resulting in a 10^3 improvement in sensitivity compared to conventional gas chromatography techniques [24]. Measurements of the nitrogen concentration for several dates throughout the R2014 data collection period can be seen in figure 8.9.




Figure 8.9: Measurements for several dates during the R2014 data collection period.

In order to reach below the ppb level, the chromatography must be carefully calibrated. Significant work to achieve this calibration was implemented only after the R2014 dataset had been collected. As such, the calibration points displayed in 8.9 did not achieve sub-ppb precision, and only offer several measurements at distinct points. A second method offering

better precision, an independent measurement, and which can confirm the sensitivity and interpolate between the chromatography measurements was thus developed, and is discussed in the subsequent section.

Capture Recoil Analysis The nuclear recoil energy from muon capture on nitrogen produces roughly 200 keV. In the 2013 data collection period, the zeolite filters were saturated, resulting in ~ 20 ppb of nitrogen impurity within the detector gas. In the case of this doped dataset, the capture recoils were clearly seen in this expected region, as depicted in the black data points of figure 8.10. The capture recoil signals are thus expected in a region typically




Figure 8.10: Comparison of the clean R2014 (red) capture recoil yield to the R2013 data (black), which contained roughly 20 ppb of nitrogen impurity according to chromatography measurements. Figure courtesy of Frederik Wauters.

dominated by the low energy tail of the ^3He fusion distribution as denoted in the annotated energy distribution of figure 8.11. This figure depicts all extra pulses following the muon which were not included within the muon stop cluster. The largest challenge of the in-situ analysis is to suppress and subtract the background well enough to measure the capture yield. Two important distinctions between capture and fusion can be utilized to select recoil events; the time of the signal relative to the muon stop and the presence of an electron.




Figure 8.11: Energy of all extra pulses which are not included in the muon stop cluster. The recoil energy from nuclear captures on nitrogen is expected in the 100-200keV region, beneath the low energy tail of the ${}^3\text{He}$ peak. The lowest energy peak is dominated by noise and electron signals in the TPC.

In order to assess the timing of impurity captures relative to fusion, the kinetic equations are re-solved including a capture term in appendix B. The time distribution for both ${}^3\text{He}$ fusion and nitrogen captures for 1 ppb of N_2 are displayed in figure 8.12(a), where the capture has been scaled by a factor of 5×10^4 . The capture products build over the first few microseconds, compared to the fusion products which quickly diminish. As such, the background from the ${}^3\text{He}$ fusion reaction can be suppressed by observing the capture recoil yield in a delayed time region. Figure 8.12(b) shows the delayed pulse amplitude versus the time of delayed pulse relative to the muon stop time.

In the case of a ${}^3\text{He}$ fusion the muon is recycled, generally resulting in a delayed electron, contrary to capture events. Thus the remaining ${}^3\text{He}$ background can be eliminated by selecting delayed events with an electron, scaling to the events without an electron, and subtracting. The difference between events with an electron and without an electron for the

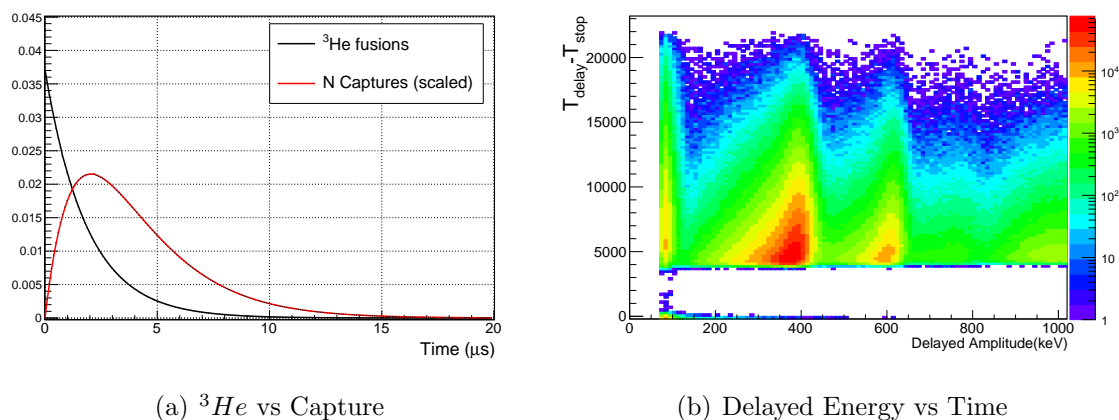


Figure 8.12: (a) The time distribution of ${}^3\text{He}$ fusions (black) compared to the scaled nitrogen capture time distribution (red). The ${}^3\text{He}$ fusion background can be suppressed relative to the captures by selecting events after $3\text{-}4 \mu\text{s}$. (b) The delayed pulse amplitude versus the pulse arrival time relative to the muon stop time.

full R2014 dataset can be seen in figure 8.13.

In order to relate the capture yield from the delayed event analysis to the impurity concentration in the TPC, a first calibration was instigated following the R2013 beam period and the TPC was doped with 0.5 and 2 ppb of nitrogen. Using this initial calibration and comparing a preliminary analysis of the R2014 data to the R2013 data, as seen in figure 8.10, it is estimated that the R2014 dataset contains 2-3 ppb of nitrogen impurity throughout the data collection period.

There are several issues with this preliminary analysis which need to be addressed. First, there were several issues with the calibration. In the initial calibration, a well defined volume of gas was diluted with a known level of impurity, and used as the target gas. This static calibration technique allows for the accumulation or adsorption of impurities from seals and wall materials, leading to an imprecise knowledge of the initial concentration. Furthermore, a mechanical failure of the TPC inhibited the ability to run the chamber with nominal




Figure 8.13: Comparison of the clean R2014 delayed energy distributions with and without a delayed electron, scaled to the ${}^3\text{He}$ peak. There is an excess of events in the expected capture recoil region for the distribution without a delayed electron.

conditions, and lower voltages were applied throughout the calibration period, leading to the necessity of energy scaling when comparing the doped and clean datasets.

A more robust calibration was performed in 2015, with a dynamic calibration technique which utilized mass flow controllers to continuously mix clean deuterium with polluted gas of a known concentration. The gas was once again doped with several known impurity concentrations. This calibration was undertaken with the same TPC used in the R2014 data collection period, with nominal densities and voltages applied. Analysis of this calibration data is currently in progress.

The second issue with the initial impurity analyses is an additional background from neutron scatters off deuterium nuclei. It is estimated that 50-70% of the events observed in the capture recoil region are from fusion neutrons which scatter and leave the same signal in

the TPC. Methods to quantify this background have been highly successful in the context of the muon catalyzed fusion analysis presented in section 10.2.1. In figure 10.3(a), the extra cluster energy is used to select the E9 distribution for n-d scatter events which are subtracted in 10.4, reducing the effect of this background for p-t event selection by two orders of magnitude. A similar technique will be applied to the impurity analysis to reduce the effect of deuteron recoils from neutron scatters on the capture yield.

Chapter 9

MUON ON REQUEST AND ACCIDENTAL STRUCTURES

As discussed in section 4.2, the Muon-On-Request (MORE) signal implements a kicked beam structure in order to optimize the rate of good events with a single muon detected within a $\pm 25 \mu s$ time window. Once a muon is detected in the entrance scintillator, the MORE signal is sent to a pair of rapidly charging parallel plates which divert the beam into the downstream collimator, as depicted in the simplified diagram of figure 9.1. The collimator lies several meters downstream at a focus within the bending magnet. More information on the beam profile can be found in appendix D. The dimensions and operating voltages of the deflection plates can be found in table 9.1.

Parameter	Symbol	Value
Plate length	l	1.5 m
Plate gap	d	0.12 m
Voltage difference	ΔV	25 kV

Table 9.1: Specifications for the MORE deflection plates.

The deflection angle of the beam from the plates is given by

$$\theta = \arctan\left(\frac{\Delta V l e}{d p \beta c}\right) \quad (9.1)$$

where ΔV is the voltage difference in volts (equivalent to eV/e), p is the beam momentum in eV/c, and $\beta \approx 0.35$ for the muons, and nearly 1 for the electrons, expressed units where the




Figure 9.1: Sketch of the muon on request method. Once a muon is detected in the entrance scintillator, a pair of rapidly charging plates divert the incoming beam into an upstream collimator, decreasing the likelihood of muon pileup. This is a simplified schematic. For the actual realization, including the beam envelope, see appendix D.

speed of light and the fundamental charge equal 1. The nominal deflection angle of 22 mrad for the muons leads to a much smaller 7 mrad deflection for the electrons.

The time of all muon entrances relative to the time of the kick can be seen in figure 9.2. A muon entrance triggers the kicker at time $t=0$ and the plates rapidly charge via a single stack of MOSFETS operating in push-pull mode. The time it takes for the signal to travel the 10 m to execute the kick, the charging time of the MOSFETS, and the emptying of the beam pipe each contribute to an 800 ns delay in the reduction of the beam rate. The muon rate is then suppressed by an extinction factor, dependent on the beam tune, for the $25 \mu s$ window in which the plates are charged. After $25 \mu s$, the plates are discharged and the nominal muon rate is observed.

As the collimator was not properly aligned for the first half of the R2014 collection period,

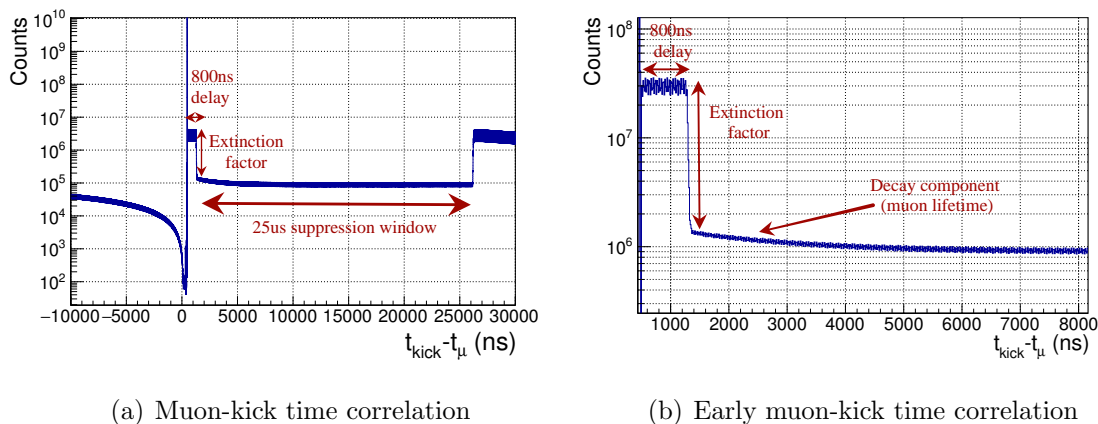


Figure 9.2: Time correlation of all muon entrance times relative to the time the kicker fires, showing the artificial time structure imposed by the MORE signal. (a) Shows the distribution from $-10 \mu s$ before to $30 \mu s$ after the kicker triggers, illustrating the $25 \mu s$ suppression window. (b) Shows the first $8 \mu s$ of the same distribution, highlighting the 800 ns delay before the MORE system suppresses the nominal beam rate. The decay component observed at early times falls with the muon lifetime, and is caused by decay electrons emitted upstream back into the muSC.

the initial beam tunes were unable to retain a high extinction factor. Once the collimator was replaced, the muon rate was suppressed by the kicker by a factor greater than 100. The extinction rates achieved as a function of the time throughout the data collection period can be seen in figure 9.3.

Several potential sources of systematic effects are associated with the implementation of the MORE beam structure. There are two cases in which pileup muons enter the detector; those seen by the entrance detectors and those unseen by the entrance detectors. If pileup muons are stopping within the gas, only a shift in the start time of the exponential decay would be observed, and the correct rate would still be extracted from a lifetime fit as long as the fit starts late enough. However, a second muon which stops outside of the target gas can capture on nuclei other than deuterium, leading to a significant distortion as described in section 3.3. In order to mitigate the effects of pileup muons which are seen by the detectors,




Figure 9.3: Suppression of the muon beam rate from the Muon-On-Request signal as a function of time within the data collection period. A better suppression is achieved once the collimator is properly aligned. The extinction factor is sensitive to the parameters of the beam tune, which changed several times throughout the data collection period, leading to the discrete jumps observed.

pileup protection is applied, and any event with a second muon within $\pm 25 \mu s$ is rejected. In order to study the effects from pileup muons which are unseen by the entrance detectors, the entrance inefficiencies and effects on the lifetime are quantified using TPC information, as discussed in the first half of this chapter.

More difficult to quantify are the effects from muons which stop upstream in various components of the beamline. These muons may decay in upstream materials with varying lifetimes, resulting in time dependent electron backgrounds which may be observed in the electron detectors either directly or from scattering. Additionally, although the deflection of beam electrons from the MORE setup is much smaller, it is non-zero which can lead to additional distortions to the observed electron background. The electron background is studied extensively using both the muon clock framework and a novel technique to enhance or suppress any such background by making event selections on electron multiplicity. Effects from potential electron background sources are described in the second half of the chapter.

9.1 Muon Pileup

It is crucial to understand the pileup inefficiency of the entrance detectors to quantify any remaining effects from pileup muons which are unseen. There are two sources of inefficiency. The muSC has an intrinsic deadtime, leading to a possible time dependent efficiency. Time independent inefficiencies may also exist in both the muSC and the muPC.

9.1.1 Time Dependent Inefficiency from muSC Deadtime

As discussed in section 7.2, the muSC signal is sent into two separate digitizers, and the coincidence of both TDCs is used to create muSC objects. The autocorrelation of both muSC copies can be seen in figure 9.4. Copy 1 has a hardware deadtime of 25 ns. The second muSC TDC copy suffered afterpulsing effects, and exhibits a hardware deadtime around 28 ns. In order to quantify the effects of the hardware deadtime on the observed decay




Figure 9.4: Autocorrelations of the muon entrance times determined by each muSC TDC copy. Copy 1 (red) suffered afterpulsing effects, and both copies exhibit a hardware deadtime of 25-28 ns.

rate, an artificial deadtime (ADT) is applied in software. The ADT is varied to understand the effect on the final result, and possibly to make a correction for the unavoidable hardware deadtime. Lifetime histograms are created for a range of applied artificial deadtimes. Each

lifetime is fit with a standard three parameter exponential fit, with a start time of $1 \mu s$. The lowest value deadtime applied is 20 ns, which eliminates the potential of afterpulsing effects, but reflects the hardware deadtime – any smaller ADT would have no effect. The shift in the value obtained for the lifetime, relative to the lifetime obtained with 20 ns ADT, can be seen for the case with solely muSC pileup protection in figure 9.5(a) and with the HasBestEntrance (HBE) requirement including muPC pileup protection in figure 9.5(b). The slope of each linear fit quantifies the change in the lifetime as the deadtime of the muSC is increased. The improvement with the addition of the HBE requirement is due to the muPC, which catches spatially separated muons which sneak in within the muSC deadtime. For the HBE definition, the lifetime shift is 0.02 Hz/ns, which implies that for our standard 30ns ADT condition, we expect a shift of 0.6 Hz to the measured lifetime.

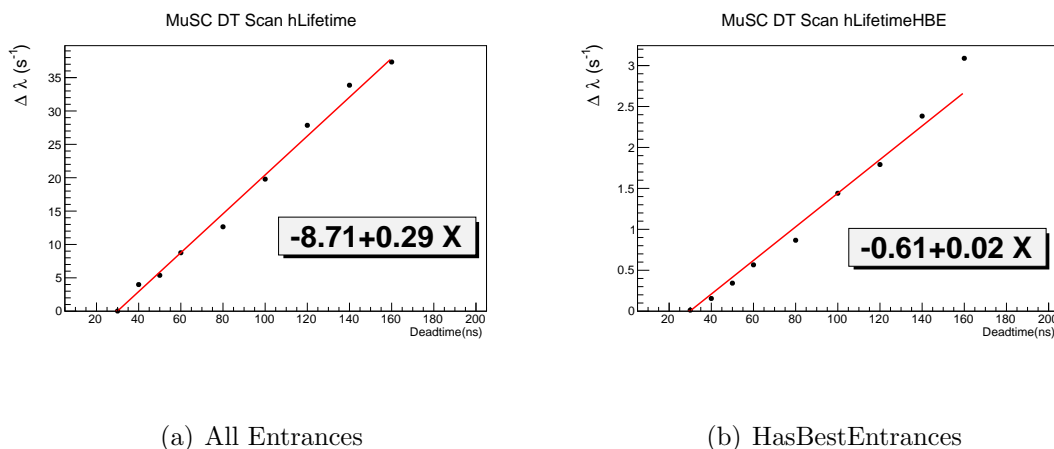


Figure 9.5: Linear fits to the artificial deadtime scans for (a) all muon entrances and (b) HasBestEntrance muons

9.1.2 Time Independent Entrance Detector Efficiencies

Three independent methods are used to obtain the inefficiencies of the various entrance detectors used for pileup protection. The most straightforward of the three methods involves

analyzing a shift of data which was taken with the muon clock signal at a frequency high enough to mimic the nominal entrance rate, and the real muSC taken out of the kicker OR. Once pileup protection is applied for these events, there should be no track within the TPC, as a clock fired the entrance trigger instead of a real muon. As the detector efficiencies can potentially change over the course of the data taking period, the same method can be applied to the in-situ muon clock for the datasets in which the muon clock was active. For the remaining datasets, a third method utilizing events with two tracks in the TPC within a single event is used to determine the pileup efficiency. In general, the detector inefficiency is given by

$$\varepsilon = \frac{N}{N_0} \quad (9.2)$$

with error

$$\delta\varepsilon = \sqrt{\varepsilon(1 - \varepsilon)/N_0} \quad (9.3)$$

where N is the number of counts observed after pileup protection has been applied for a given detector and N_0 are the total number of counts observed in the TPC. Any background which is mistaken for a muon stop within the TPC will artificially inflate N . As such, care must be taken to define a clean sample of events which are clearly muon stop signals, and the obtained inefficiency represents an upper limit. In order to reduce the background, more stringent cuts on the muon stop energy deposition are made than in the normal analysis. As will be demonstrated, the sufficiently clean selection allows for pileup suppression of several orders of magnitude to be observed.

Muon Clock Shift

A shift of data was collected with the real muon signal removed from the muSC OR, such that the kicker is only triggered by the clock. Real muons remain in the pileup definition but do not fire the kicker and no entrance objects centered on real muon signals are created. Thus, no muon stop should be present in the TPC after pileup is applied. The muon clock


frequency was increased to 30kHz to maximize the data rate. The following cuts are made to select events for this analysis:

- muCL triggered event
- kicker triggered
- $1700 < \text{SEnergy} < 2400$

The cut on the muon stop SEnergy selects for clean muon stops and is critical to remove any background signals in the TPC. The muon arrival time relative to the kicker trigger can be seen in 9.6(a). As the kicked window from the previous clock ends around $-8 \mu s$, the full beam rate is seen until the clock fires the kicker gate at time $t=0$, at which point the kicked beam rate is observed for the $25 \mu s$ gate. In the absence of pileup protection, the stop distribution of the muon clock shift should be a convolution of the muon arrival time relative to the clock time with the nominal muon stop distribution.


$$d(t_{track} - t_{Cl}) = \int f(t_{\mu} - t_{Cl}) * d_{nom}(t_{track} - t_{\mu}) dt_{\mu} \quad (9.4)$$

Where $f(t_{\mu} - t_{Cl})$ is the muon time arrival with respect to the clock and $d_{nom}(t_{track} - t_{\mu})$ is the nominal muon drift time distribution. The arrival time of muons with respect to the clock and nominal drift time distribution can be seen in panels (a) and (b) of figure 9.6. The convolution of these spectra compared to the drift time distribution observed from the muon clock shift can be seen in panel (c). Next, pileup protection is applied for each detector in the $\pm 25 \mu s$ window relative to the clock time. As the drift time of the chamber permits tracks to be seen from any muon for up to $16 \mu s$, the region from $8 \mu s$ before $t=0$, and for $25 \mu s$ after is a clean region, used to count the number of tracks remaining after pileup protection is applied. The remaining distribution after application of pileup from various entrance detector combinations, seen in 9.7(a), is used to determine an array of efficiencies. As a reminder,



(a) Muon-Clock Correlation

(b) Muon-Track correlation



(c) Clock-Track Correlation

Figure 9.6: Convolution of the muon arrival time with the nominal muon stop distribution, used to reproduce the stop distribution seen in the muCL shift data. (a) Time of the muon arrival relative to the muon clock signal, showing the full beam rate until the kicker is fired by the clock, then the suppressed rate for a $25 \mu s$ window. (b) The time of the muon stop as determined by the TPC tracking algorithm relative to the muon entrance time. The peaks at 2 and $16 \mu s$ result from muons stopping in the anode and cathode, respectively, indicating a $5 mm/\mu s$ drift velocity over the 72 mm vertical dimension of the TPC. (c) A convolution (blue) of the distributions shown in (a) and (b) yield the time of the muon stop time relative to the muon entrance time, confirmed with the observed distribution (red).

the characteristics of each detector can be found in section 4.3. Of particular interest is the muSC WFD, which can be applied with a lower threshold than the muSC TDC, as seen in

figure 7.2. A standard threshold of 100 ADC^*_{ct} , which is more sensitive to muons while still rejecting most electrons, is utilized in this analysis. The inefficiency is given by normalizing each observed distribution to the distribution without pileup protection, displayed in 9.7(b), and is nearly constant over the region of interest. As such, the overall inefficiency of each detector combination is found by integrating the distribution after pileup is applied from $[-8,24] \mu s$ and dividing by the total number of counts in this region before pileup is applied. As it is of interest to determine if the detector inefficiencies are rate dependent, the pileup inefficiency in the unkicked and kicked beam regions are also found by integrating over $[-5,15] \mu s$ and $[17,24] \mu s$, respectively. These results are shown in table 9.2. Although the statistics are lower in the kicked stretch, there is some indication that the inefficiencies of the muPC and muSC WFD are higher in this region. As the electron background is higher in the unkicked region, and both the muSC WFD and muPC are sensitive to some electrons, the efficiency in the unkicked region could be artificially inflated. It is more likely, however, that the accidental background floor in the TPC is reached with the combination of the entranced detectors used in pileup, which explains the inverse shape of the normalization distribution observed in figure 9.7(b).

Although the muon clock shift has proved useful in determining absolute detector efficiencies and possible rate dependencies of each detector efficiency, there are several motivations to develop complementary analyses. As this shift was taken with the real muon signal removed from the nominal entrance trigger chain, no viable lifetime histograms are produced, negating any chance of using the above information for a zero extrapolation to obtain the corrected muon lifetime. Additionally, detector efficiencies could change over the course of the entire data taking period, and this shift provides only a snapshot of the detector inefficiency for this particular time period and beam tune.

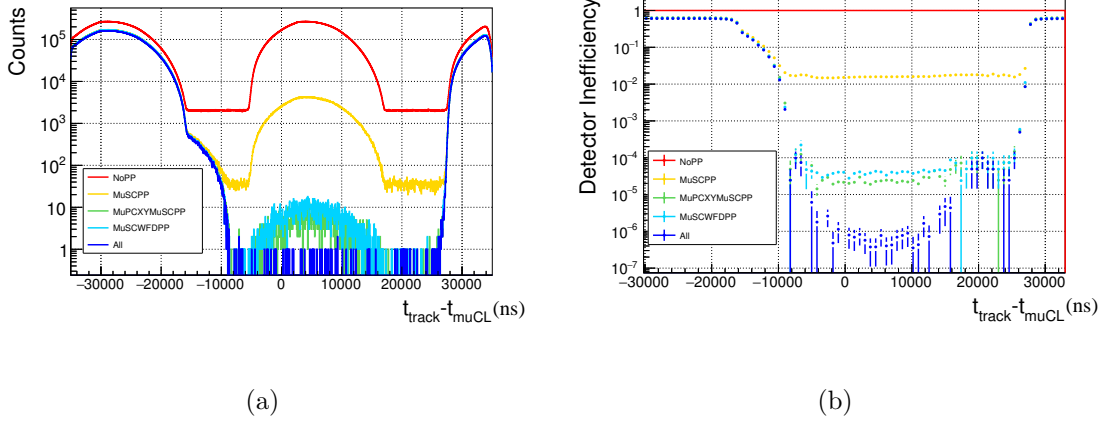


Figure 9.7: (a) Clocked muon distribution as detector pileup is applied and (b) pileup inefficiency for various pileup configurations as a function of the time relative to the kick, determined by the number of counts with pileup applied divided by the total number of counts in each bin.

Detector	Detector Inefficiency		
	All $[-8,24]\mu s$	Unkicked $[-5,15]\mu s$	Kicked $[17,24]\mu s$
muSC	1.6×10^{-2} (1.4×10^{-5})	1.6×10^{-2} (1.4×10^{-5})	1.8×10^{-2} (2.1×10^{-4})
muSC + muPC	2.4×10^{-5} (5.5×10^{-7})	2.3×10^{-5} (5.5×10^{-7})	5.5×10^{-5} (1.2×10^{-5})
muSC WFD	3.9×10^{-5} (7.1×10^{-7})	3.9×10^{-5} (7.1×10^{-7})	7.0×10^{-5} (1.4×10^{-5})
All	9.6×10^{-7} (1.1×10^{-7})	7.0×10^{-7} (9.5×10^{-8})	3.3×10^{-5} (9.6×10^{-6})

Table 9.2: Inefficiencies of several entrance detector combination from the muon clock shift data, determined by integrating the number of counts remaining after pileup is applied and normalizing to the number of counts without pileup. The first column denotes the inefficiency from each detector combination using the full pileup region. The second and third columns contain the inefficiencies determined from the kicked and unkicked regions. The values in parenthesis are statistical errors.

In-situ Muon Clock

For several later datasets, an in-situ muon clock was running in conjunction with the nominal muon trigger such that either a real muon or a clock could fire the kicker, as seen in figure 4.4. As the clocked events are recorded in a separate TDC channel, in addition to the normal muSC TDC, these events can be selected and the same procedure as above is used to determine the inefficiency of several combination of pileup detectors. However, as the real muons can now fire the kicker, any muon in the region $t \in [-25,0] \mu s$ was necessarily missed by the muSC. This unseen muon distribution takes a different shape than the time distribution of all muons, all illustrated in figure 9.8. As the time distribution differs from any




Figure 9.8: Stop distributions of in-situ muon clock triggered events for dataset DS14.

combination of detectors which records the muon seen by the muSC, the pileup inefficiencies can no longer be determined by normalizing to the distribution before pileup protection is applied. The pileup inefficiencies relative to the muSC pileup protected distribution are calculated and can be seen in table 9.3 for each dataset with the in-situ muon clock available.

The sum of all datasets in which the muon clock was running is used for a zero extrapolation. A lifetime histogram created with each combination of pileup conditions is fit to the

Dataset	Detector Inefficiency Relative to muSC		
	muSC + muPC	muSC WFD	muSC + muSC WFD + muPC
DS06	4.3×10^{-2} (1.6×10^{-2})	3.7×10^{-1} (3.7×10^{-2})	4.3×10^{-2} (1.6×10^{-2})
DS13	2.0×10^{-3} (6.5×10^{-5})	7.1×10^{-3} (1.2×10^{-4})	2.2×10^{-4} (2.2×10^{-5})
DS14	2.7×10^{-3} (1.7×10^{-4})	1.4×10^{-2} (3.9×10^{-4})	4.7×10^{-4} (7.2×10^{-5})
DS16	5.7×10^{-3} (3.1×10^{-4})	8.9×10^{-2} (1.2×10^{-3})	2.3×10^{-3} (2.0×10^{-4})
All muCL DS	2.5×10^{-3} (6.2×10^{-5})	1.6×10^{-2} (1.6×10^{-4})	4.6×10^{-4} (2.7×10^{-5})

Table 9.3: Inefficiencies of several entrance detector combinations from the in-situ muon clock data, determined by integrating the number of counts in the region from -8 to 24 μs remaining after pileup is applied and normalizing to the number of counts in this region with muSC pileup protection. The values in parenthesis are statistical errors.

nominal three parameter exponential fit and compared to the relative inefficiency of each combination, as shown in figure 9.9. The parameters obtained by the linear fit yield the following equation

$$\lambda(\varepsilon) = 4.51(1.36) \varepsilon + 1051.54(11.15). \quad (9.5)$$

As the number of events eliminated at each step of pileup protection is extremely small, the lifetime distributions are highly correlated. This correlation must be taken into account in the fit minimization. As the distributions are almost completely correlated, the covariance matrix was estimated to be

$$V_{ij} = r \sigma_i \sigma_j \quad (9.6)$$

with $r=1$ for the diagonal elements and $r=.99$ for the off diagonal elements. The least squares fit was then obtained with the python curve_fit package, which minimizes the following

$$\chi^2(\epsilon) = (\lambda - f(\epsilon))^T V^{-1} (\lambda - f(\epsilon)) \quad (9.7)$$




Figure 9.9: (a) Pileup efficiency of several combinations of entrance detector relative to the muSCPP distribution for the in-situ clock events and (b) a zero extrapolation relating the fitted lifetime of each pileup protected distribution to the obtained inefficiency. Error bars shown are statistical.

where λ is a vector of the fitted lifetime values, and $f(\epsilon)$ is a vector of the predicted lifetime values from the fit.

Real Muon Events

As the in-situ muon clock was not running for the entire data collection period, a third method is required to check the detector inefficiencies for all datasets. Additionally, any correction can be determined with smaller uncertainty if the full R2014 statistics are utilized. The time distribution of all muon stops in the TPC relative to the muon entrance times, with various pileup configurations, can be seen in figure 9.10. Any muon seen in the muSC would have triggered the kick, resulting in a distribution before $t=0$ which includes only unseen muons. As the pileup window extends to $\pm 25 \mu\text{s}$, and the drift time of the chamber limits the track arrival time to $16 \mu\text{s}$ after the muon entrance, any track seen in the region from -8 to $0 \mu\text{s}$ after pileup protection is applied indicates an event in which a muon was undetected in the pileup window. In order to avoid any tails, the region from -7 to $-3 \mu\text{s}$ is used in this

study. In addition, the distribution with no pileup protection also includes the seen muon distribution, so the normalization must be taken once more from the muSC pileup protected distribution. The pileup inefficiencies for several detector combinations can be seen for each large dataset in 9.4.

Dataset	Inefficiency ($1 - \varepsilon$)		
	muSC + muPC	muSC WFD	muSC + muSC WFD + muPC
DSM1	1.2×10^{-2} (1.1×10^{-4})	1.5×10^{-2} (1.2×10^{-4})	5.3×10^{-3} (7.6×10^{-5})
DSM2	9.8×10^{-3} (1.2×10^{-4})	1.3×10^{-2} (1.4×10^{-4})	7.9×10^{-4} (3.6×10^{-5})
DSM3	6.7×10^{-3} (9.9×10^{-5})	1.3×10^{-2} (1.4×10^{-4})	5.2×10^{-3} (8.7×10^{-5})
DSM4	1.9×10^{-3} (3.1×10^{-5})	6.7×10^{-3} (5.6×10^{-5})	2.9×10^{-4} (1.2×10^{-5})
DSM5	4.3×10^{-3} (9.2×10^{-5})	4.0×10^{-2} (2.8×10^{-4})	1.4×10^{-3} (5.4×10^{-5})
All DS	5.8×10^{-3} (3.5×10^{-5})	1.3×10^{-2} (5.4×10^{-5})	2.2×10^{-3} (2.2×10^{-5})

Table 9.4: Inefficiencies of several entrance detector combinations, determined by integrating the number of counts in the region from -7 to $-3 \mu s$ remaining after pileup is applied and normalizing to the number of counts in this region with muSC pileup protection.

The sum of all datasets is used for a zero extrapolation, to determine the lifetime correction and uncertainty from any undetected muon pileup. The inefficiency of several combinations of pileup detectors as well as the zero extrapolation from the fitted lifetime for each combination of pileup detectors can be seen in figure 9.11. The parameters obtained by the linear fit yield the following equation

$$\lambda(\varepsilon) = 4.98(1.02) \varepsilon + 1075(8.07) \quad (9.8)$$

consistent with the values from the fit of the in-situ muon clock data. Once again, the high correlation of the lifetimes was taken into account in the fit procedure. The overall lifetime




Figure 9.10: Time distribution of all muon tracks in the TPC relative to the muSC entrance time.




Figure 9.11: (a) Pileup efficiency of several combinations of entrance detectors relative to the muSCPP distribution for the real muon events and (b) a zero extrapolation relating the fitted lifetime of each pileup protected distribution to the obtained inefficiency.

shift from any remaining pileup muons can be determined by the efficiency with the most strict pileup condition, and the slope extracted in the fit. For the sum of all datasets, an

inefficiency of 2.2×10^{-3} relative to the muSC remains after all detectors are included in the pileup protection, indicating a lifetime shift of 0.011 (2.2^{-3}) Hz.

The final point used in the extrapolation in figure 9.11(b) is the muSCPP lifetime, with an inefficiency of 1 relative to the muSC. If this point is removed from the fit, a better fit is achieved, as shown in figure 9.12, with the following equation

$$\lambda(\varepsilon) = 704(111) \varepsilon + 1050(8.9) \quad (9.9)$$

resulting in a much larger slope. In this case, the remaining inefficiency after all pileup has been applied results in a lifetime shift of 1.54 (0.24) Hz. As further studies are required to assess the linearity of the correction versus inefficiency, the higher value is taken as the upper limit for now.




Figure 9.12: Fit of the observed decay rate versus pileup inefficiency relative to the muSC excluding the inefficiency of the muSC.

9.2 Electron Background

The second systematic associated with the MORE artificial beam structure is the potential for a time dependent electron background. As R2014 was the first production dataset col-

lected in the new PIE1 beamline, the magnitude of this background was unknown. Several issues, including the mis-tuning of an upstream magnet and a mis-alignment of the electron-muon beam collimator, led to an increase in the overall electron background and difficulty in reaching an optimal beam tune. Additionally, the eSC based analysis used in this thesis is more prone to electron backgrounds than one including the electron wire chambers. The electron background is studied extensively using the in-situ muon clock, and a novel technique which enhances or suppresses the electron background by selecting for event multiplicities. An 800 ns step is observed in the electron distribution due to the MORE structure, making it difficult to start the fit earlier than $1 \mu s$. To recover the 30% of the statistics which are lost with start times after $1 \mu s$, an effort is made to subtract this time dependent background. A discrepancy in fitted rates obtained for lifetime distributions created with different electron multiplicities for fit start times later than $1 \mu s$ is indicative of a second electron background effect. This is studied in the second half of this section.

9.2.1 Observing a Time Dependent Beam Electron Background

During the R2014 data collection period, a time dependent beam electron background is observed with timescales consistent with those expected from the MORE implemented beam structure. Depending on the beam configuration, the flat electron background was either suppressed or enhanced in the 800ns window before the kicker fires. This background can be seen cleanly with the in-situ muon clock. If pileup protected muCL events are selected, which have no Michel electrons, the accidental background structure is apparent. The electron background rates with respect to the muon clock time for the datasets with the muon clock running in situ, (DS06, DS13, and DSM5) can be seen in figure 9.13. Dataset DS06 was taken with the old collimator and has significantly higher electron background and different shape than the later datasets. Although the muon clock data is the cleanest way to observe the electron background, the in-situ muon clock was not running for all datasets. As an

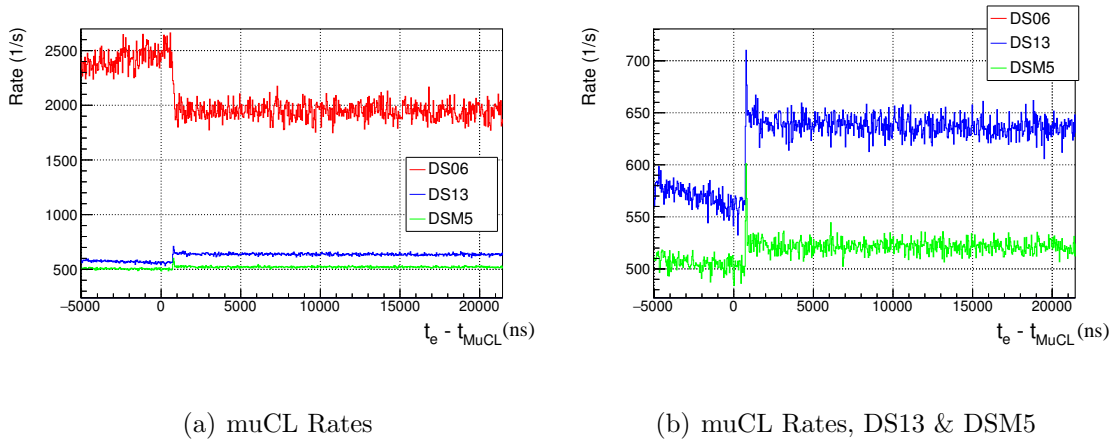
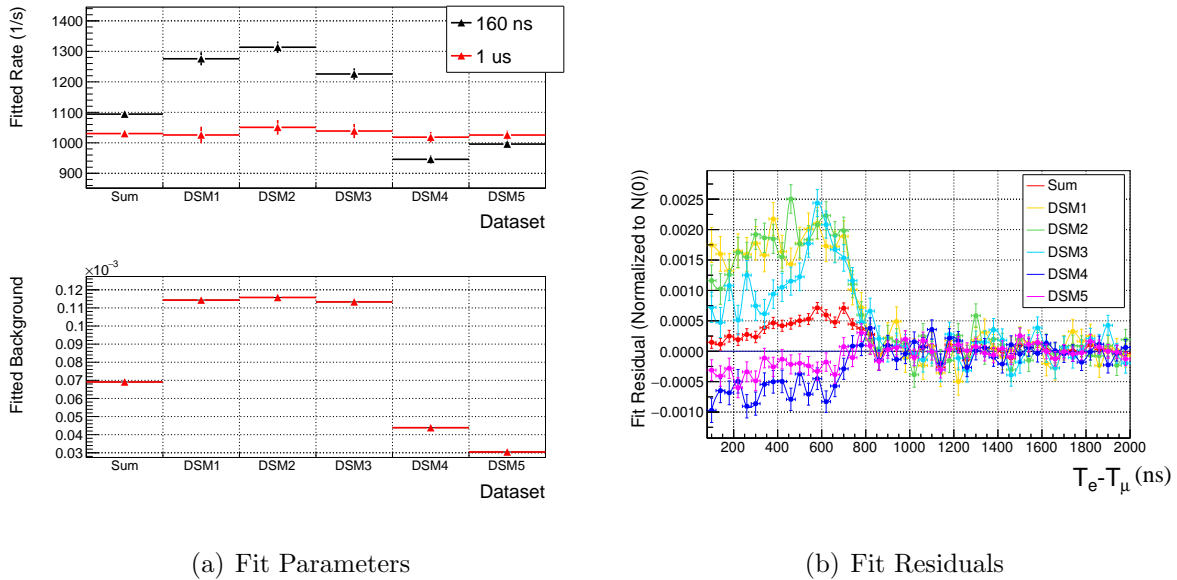


Figure 9.13: (a) Accidental background rates from the in situ muon clock. On the x axis is the time of the electron versus the time of the muon clock, which is pileup protected and fires the kicker. On the Y axis is the number of counts scaled to the number of muon entrances, yielding the background rate in 1/s. (b) A zoomed version shows the time dependent background for the later datasets.

alternative, the magnitude of the background can be extracted from the lifetime fit residuals, as displayed in figure 9.14. Figure 9.14(b) shows the difference between the data and the fit function value normalized to $N(0)$, the number of counts at $t=0$, for early times. A clear step at 800ns reveals an inconsistency between the data and the exponential fit function. The pull on the lifetime is quantified by the differences in rates obtained with an early 160 ns start time and a late 1 μs start time. It is found that the fitted rate shifts 3 Hz for a 1×10^{-5} background step relative to the bin zero content of the lifetime histogram. As a consistency check, the background rate can be extracted from the fitted muCL time distribution and compared to the rate obtained from the muon lifetime fit. Both methods should yield the same accidental background rate. The rate is determined by

$$R_B = \frac{B}{N * \Delta t}, \quad (9.10)$$



(a) Fit Parameters

(b) Fit Residuals

Figure 9.14: Lifetime fit parameters and residuals for the sum, and each of the five sub-datasets for the R2014 data.

where B is the fitted background rate, N is the number of muons used to fill the histogram, and Δt is the histogram bin width. A comparison of the calculated rates for each dataset for which the muon clock was running can be seen in table 9.5. The rate calculated via

Dataset	Fitted Background	muCL Background	Rate Difference
DS06	2183 (1.2)	2088 (3.4)	4.4%
DS13	712 (.2)	692 (.4)	2.9%
DSM5	569 (.2)	560 (.3)	1.5%

Table 9.5: Comparison of background rates calculated from the muCL and Lifetime fits.

muCL data is close, but lower than that obtained by the lifetime fit by a few percent, with the discrepancy slightly worse for the earlier datasets. The cause of this difference is

still unknown. There are two additional requirements for creation of the nominal lifetime histogram compared to the clocked distributions: muPC XY coincidence and a fiducial volume stop in the TPC. As they are essentially additional event selection cuts, either of these requirements could slightly alter either the background or the muon normalization. The rates ignoring the muPC and the TPC can be seen in table 9.6, ruling out these detectors as the cause of the discrepancy.

	No muPC			No TPC		
Dataset	Fitted BG	muCL BG	Difference	Fitted BG	muCL BG	Difference
DS06	2237 (1.2)	2145 (3.3)	4.1%	2188 (.7)	2088 (3.4)	4.6%
DS13	741 (.2)	719 (.4)	2.9%	720 (.1)	692 (.4)	3.9%
DSM5	584 (.2)	576 (.3)	1.4%	576 (.1)	560 (.3)	2.7%

Table 9.6: Comparison of background rates calculated from the muCL and lifetime fits, ignoring the muPC or the TPC.

Although the background rates do not completely agree, the rates are close enough to instill some confidence, and the muCL can be scaled directly to the lifetime distributions using an integral at late times to check whether the magnitude of the kicker step is consistent in the lifetime fit residuals and clocked data. A comparison of the lifetime fit residuals normalized to $N(0)$ and the muCL distribution scaled to match the late time integral can be seen in 9.15. The residuals have an additional downward pull at early times from muon-catalyzed fusion interference effects, described in chapter 10. For each dataset in which the muon clock was continuously running in situ, the magnitude of the increase or decrease in background rates at 800ns relative to the late time background agrees well between the muon clock distribution and the data.




Figure 9.15: Muon clock distribution and residuals normalized to the number of counts at $t=0$, $N(0)$, of the lifetime histogram for each dataset in which the muCL was continuously running in situ.

9.2.2 Electron Multiplicity Technique

The events centered around muon entrances have various electron multiplicities. In principle, the electron background can be either enhanced or suppressed by selecting events with specific electron multiplicities. In the final analysis, only events with single electrons are selected in order to suppress the electron background contributions. However, to study potential effects on the lifetime, it is also of interest to enhance the time dependent background by selecting events with multiple electrons. Particularly, as the muon clock was not running for the entire R2014 collection period, the time dependent background can instead be quantified using the multiple electron technique.

Formalism of the Electron Multiplicity Technique In events with two electrons, one is presumably the Michel electron from the muon decay and the other is an accidental, but the two cannot be distinguished. The overall time distribution can be written as the muon decay distribution plus our time dependent background

$$e_{all}(t) = \epsilon \lambda e^{-\lambda t} + R_B(t) \equiv e(t) + R_B(t) \quad (9.11)$$

If events with a single electron are selected, the decay component gets multiplied by the probability that an accidental was not also seen, and the accidental term gets multiplied by the probability that the Michel electron was not seen,

$$e_1(t) = e(t) \cdot \left(1 - \int R_B(t)dt\right) + R_B(t) \cdot \left(1 - \epsilon\lambda \int e^{-\lambda t} dt\right) \quad (9.12)$$

$$= e(t) \cdot (1 - P_a) + R_B(t) \cdot (1 - \epsilon) \quad (9.13)$$

where P_a is the probability of observing an accidental background electron and ϵ is the electron detector efficiency. Both integrals are over a fixed time range relative to the muon entrance resulting in constant values. With the small background rate and high detector efficiency, the Michel term is left nearly unchanged, while the accidental background is suppressed, as seen in the green curve in figure 9.16. Conversely, by requiring two electrons in the event window, we multiply the decay component by the probability that an accidental was also seen, and multiply the accidental term by the probability that the Michel was seen.

$$e_2(t) = e(t) \cdot \int R_B(t)dt + R_B(t) \cdot \epsilon\lambda \int e^{-\lambda t} dt \quad (9.14)$$

$$= e(t) \cdot P_a + R_B(t) \cdot \epsilon \quad (9.15)$$

In this case, the Michel term is suppressed relative to the accidental term, as shown by the red curve in figure 9.16.

In the single electron distribution, the contribution of the background term relative to the Michel term is given by

$$\chi_1 = \frac{(1 - \epsilon)}{(1 - P_a)} \quad (9.16)$$

whereas in the two electron distribution, the contribution of the background relative to the Michel term is much higher

$$\chi_2 = \frac{\epsilon}{P_a}. \quad (9.17)$$




Figure 9.16: Lifetime histograms for all events (blue), events with only a single electron seen within the region $[0,25]\mu\text{s}$ relative to the muon (green), and events with two electrons seen in the $[0,25]\mu\text{s}$ event window (red). The accidental background is reduced by requiring only a single electron, while the Michel curve is suppressed by considering only events with two electrons.

The overall enhancement of the background relative to the Michel contribution in the two electron distribution is then given by

$$\chi = \frac{\varepsilon(1 - P_a)}{P_a(1 - \varepsilon)}. \quad (9.18)$$

9.2.3 Quantifying the MORE Induced Background

In order to quantify the beam electron suppression and enhancement achieved in selecting for specific electron multiplicities, a comparison of the fitted accidental rates for each dataset can be seen in table 9.7. The background rate for each dataset is found by fitting the all electron distribution, and dividing the fitted B parameter by the number of HBE muons with TPC stops in the FV times the histogram bin width. This is displayed in the first column of table 9.7. The next two columns the table quantify the electron detector inefficiency, estimated first by comparing the fitted B parameter for the single and all electron distributions, as illustrated in equation 9.13. This can be compared to the inefficiency determined by the fraction of

events with a decay electron, N , to the total number of muon entrances, N_μ . For datasets DSM4 and DSM5, the inefficiency determined via the background ratios is smaller, indicating a time correlated background effect which is not accounted for in equation 9.13. Similarly, the probability of an accidental is estimated from the Michel suppression via comparing the fitted N parameter for the double and all electron distributions, as per equation 9.15. This is compared to the expected accidental probability determined by integrating the accidental rate over the $25 \mu s$ observation window, assuming a constant rate. The value obtained from the Michel suppression is slightly larger than that obtained from integration of the rate, once again indicating an additional correlated background term explored later in this chapter. As expected, the electron detector efficiency is nearly constant throughout the run, but the accidental probability is lower for the later datasets in which the beam quality was improved with the upgraded collimator. Using the values obtained for the sum of all datasets and inserting them into equation 9.18, the background relative to the Michel contribution is enhanced by a factor of $\chi = 82$. A comparison of the lifetime fit residuals for each dataset for specific electron multiplicities can be seen in figure 9.17. Events were selected based on the number of electrons with a $[0,25] \mu s$ window around the muon entrance. The residuals, normalized to the number of counts at time $t=0$ are shown to greatly enhance the time dependent background component when events with two electrons are selected. Conversely, the time dependent background is suppressed by selecting events with a single electron in the region of interest.

In order to quantify the lifetime pull of the time dependent step, the lifetime with each electron multiplicity selection is fit with a 160ns start time, and compared to a fit with a $1 \mu s$ start time which should be unaffected by the kicker step. The fitted rates for each dataset can be seen in figure 9.18. Figure 9.18(a) shows the fitted rates when no selection is made on the number of electrons. The pull is in the opposite direction for DSM4 and DSM5, as expected from the muCL data.

Dataset	BG Rate (Hz)	$(1 - \varepsilon) = B_1/B$	$(1 - \varepsilon) = 1 - N/N_\mu$	$P_a = N_2/N$	$P_a = R_a \Delta t$
DSM1	2003	0.29	0.30	0.08	0.05
DSM2	2022	0.29	0.30	0.08	0.05
DSM3	1957	0.30	0.31	0.08	0.05
DSM4	765	0.24	0.30	0.05	0.02
DSM5	530	0.21	0.30	0.05	0.013
DS Sum	1187	0.27	0.30	0.06	0.03

Table 9.7: Comparison of fitted accidental electron background rates, electron detector inefficiency, and probability of an accidental for each of the five sub datasets. The total background rate is determined by the B fit parameter from a fit to the full electron distribution. The detector inefficiency is determined by the ratio of the fitted B parameter from the single and all electron lifetime fits as well as comparing the number of muon entrance events, N_μ with the number of events in which an electron is detected, N . The accidental probability is found from both the ratio of the fitted N values between the two electron and all electron distributions, and the rate times the $25\mu s$ observation window. The discrepancies between the two methods of calculating the inefficiency and the accidental rate could be explained by a correlated background term, explored later in this chapter.

9.2.4 *muCL Background Subtraction*

As starting the lifetime fit at $1\mu s$ sacrifices 30% of the statistics, significant efforts were made by the author to model the 800 ns step due to the MORE induced electron background. The size of the step obtained from the muCL distributions and the fit residuals are in good agreement, as displayed in figure 9.15, implying the muCL distribution can be used to quantify the underlying background. For DS06, the muon clock distribution can be fit to a function of the form:

$$muCL(t) = B \left(1 + A * \left(1 + Erf \left(\frac{t_0 - t}{\alpha} \right) \right) \right), \quad (9.19)$$

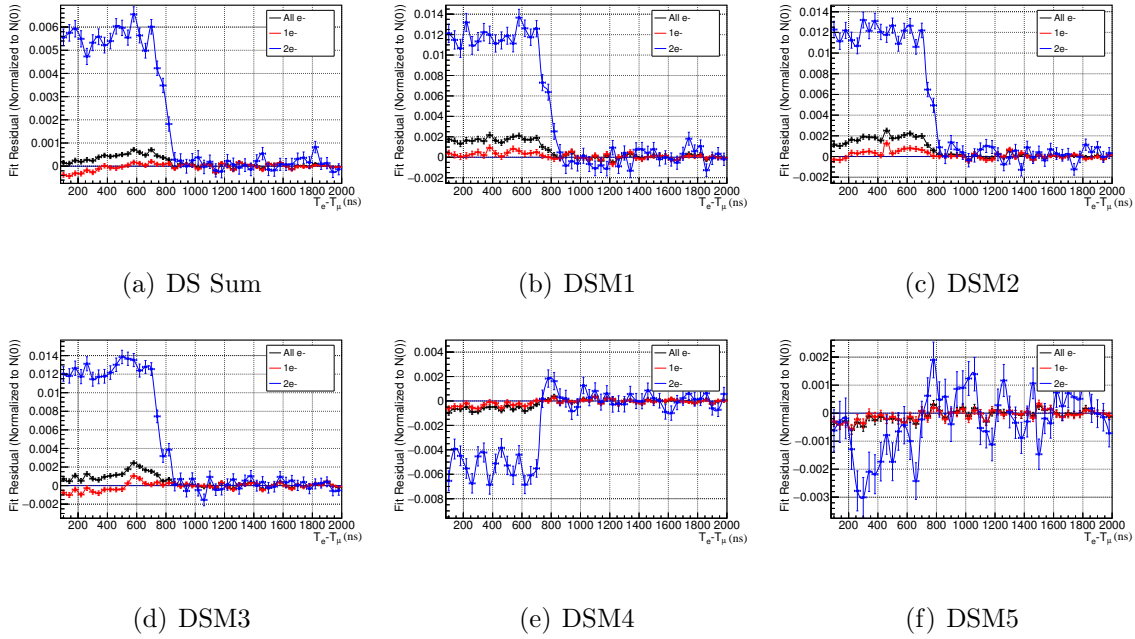


Figure 9.17: Fit residuals for a 1 μ s start time comparing all electron (black), single electron (red), and two electron (blue) event selections.

where B is the nominal background term, A is the additional early background, and t_0 and α parameterize the time and slope of the early background ‘turn off’, respectively. For DS13 and DS14 the shape of the electron background is different, presumably due to the improved beam collimator. An error function with the opposite sign is needed to account for the step. Additionally, a distinct peak is observed at 800 ns, most likely due to the charging of the kicker. A Gaussian is added to equation 9.19 for the later datasets in order to account for this additional peak. The fits to the muon clock distributions for all three dataset with the in situ muon clock running continuously can be seen in figure 9.19.

In order to eliminate the effect of the time dependent background, the function described in equation 9.19 is added to the background term of the lifetime fit, and all parameters except the background parameter, B , are fixed according to the measured muon clock step.

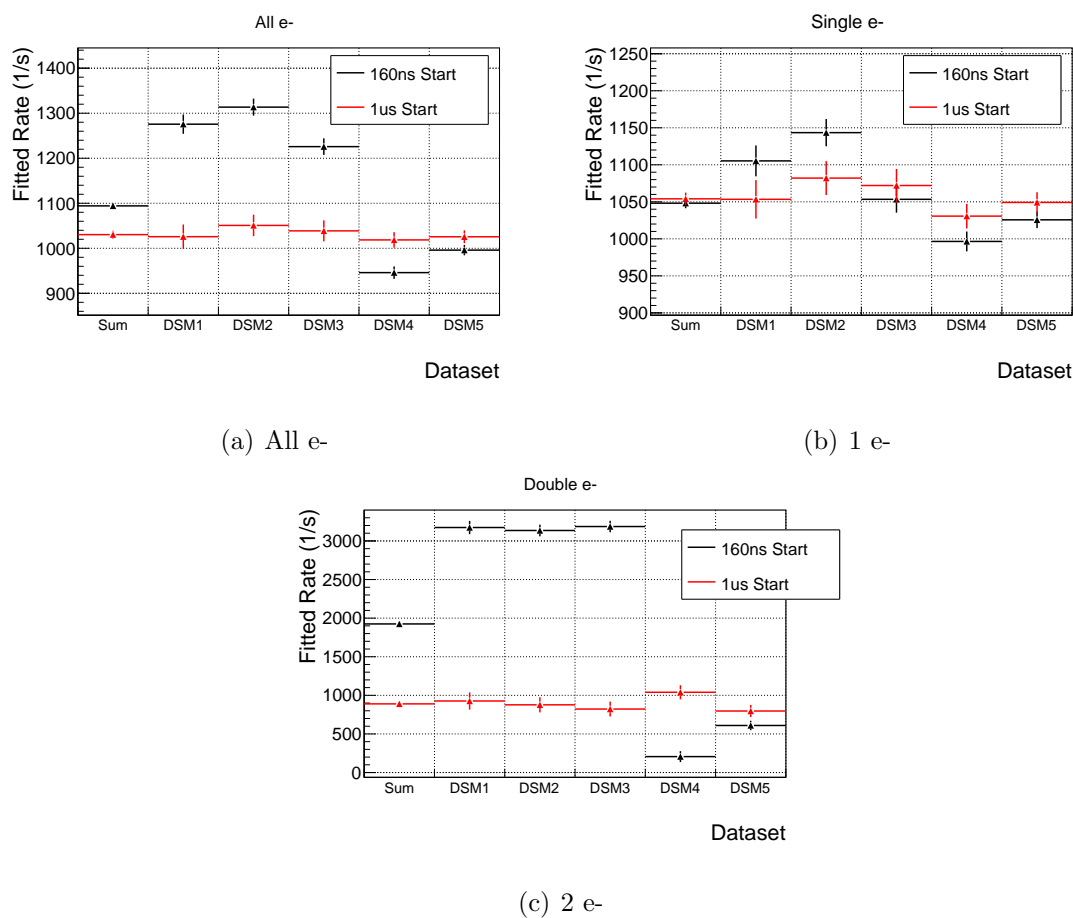




Figure 9.18: Fit parameters for a 160ns and a 1us start time comparing the single and 2+ electron lifetimes for each dataset.


The muon clock was not running continuously over all datasets, so this method cannot be applied directly to all R2014 data. It can only be assumed that all datasets before the collimator was replaced (DSM1-DSM3) have the same background distribution governed by 9.19, while DSM4 and DSM5 are modeled by the same equation with an opposite sign in the error function and an additional Gaussian. The fitted rates, chi squared, and fit residuals for each dataset, before and after applying the muCL background correction, can be seen in figure 9.20.



(a) DS06



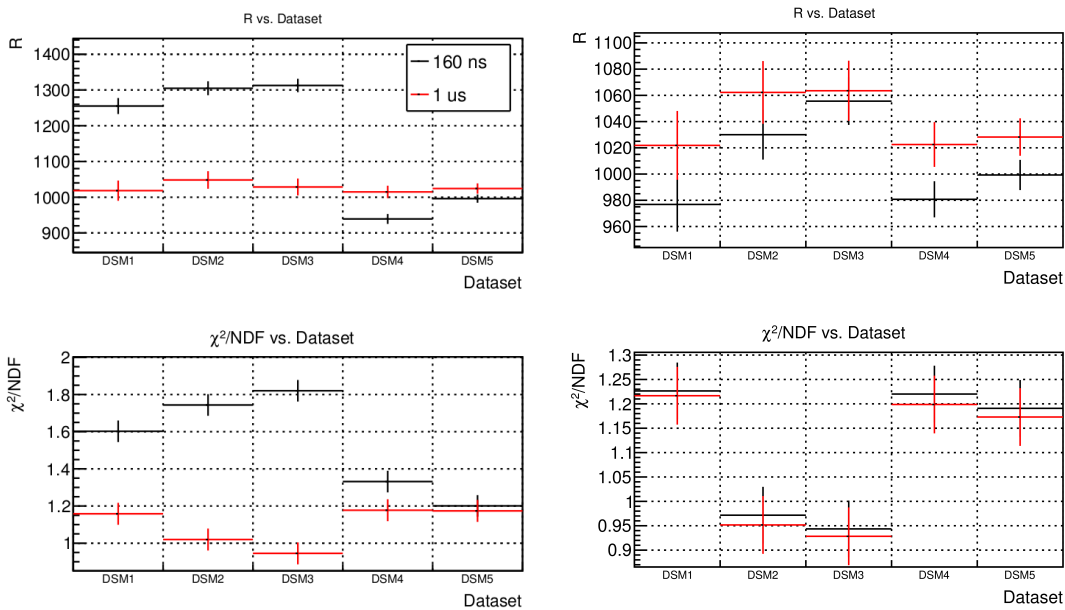
(b) DS13



(c) DS16



Figure 9.19: Fit to the eSC vs muCL time for DS06, DS13 and DS16. The collimator was misaligned in DS06, and equation 9.19 is used to fit the background shape. For DS13 and DS16, the sign of the error function is flipped and a Gaussian is added to account for the peak at 800 ns.

It is clear that adding the time dependent term as determined by the muon clock greatly reduces the residual step in the lifetime fit revealing more subtle features, such as the downward pull at early times from muon-catalyzed fusion interference. However, it is yet to be proven how accurately the 800 ns step from the muCL must be measured in order to fully correct for the time dependent background. Moreover, it cannot be ascertained whether the background had the same functional form when the in situ muon clock was not running. As such, the standard lifetime fits for this thesis start at $1 \mu\text{s}$, and the background subtraction is only used in studies to quantify early time dependencies, such as muon-catalyzed fusion



(a) Nominal Fit

(b) muCL BG Fit



(c) Nominal Fit Residuals

(d) muCL BG Fit Residuals

Figure 9.20: Scans of the fitted rate and chi squared (a) before and (b) after the muCL background step is included in the fit function for both a 160 ns (black) and 1 μ s (red) fit start time. The fit residuals before (c) and after (d) the muCL background term is added demonstrate the recovery of the sensitivity to early time effects with this technique.

interference.

9.2.5 Quantifying other Backgrounds

Assuming the time dependent background before 800 ns is the only effect from the electron background, we expect the fitted rates from the single and multiple electron distributions with a 1 μ s start time to be equivalent. This is not the case, as demonstrated in figure 9.21. The fitted rate obtained from events with two electrons is consistently lower by 100-200 s^{-1} throughout all datasets. There are two possible sources for this discrepancy which




Figure 9.21: Fitted rates for each dataset with a 1us start time for all electron (black), single electron (red) and double electron (blue) event selections.

are investigated: doubly counted events from cosmic ray muons, and scattering within the scintillators and after-pulsing in the electron scintillators. As both of these event types would be selected by the two electron distribution and rejected by the single electron distribution, either effect would be enhanced by selected events with multiple electrons.

Prompt Accidentals One can imagine several scenarios in which a single electron is double counted, including cosmic ray muons, scattering or bremsstrahlung radiation of Michel electrons, and cross talk between eSC segments. Accidentals from cosmic rays are passing through the detector at a small, constant rate, and should leave two signals in each event.

This effect should not be time dependent. Cross talk or scattering is expected to contribute to the nominal Michel distribution, and again would not lead to a significant lifetime distortion. In any case, these accidentals in prompt coincidence must be characterized in order to correctly interpret the two electron distributions. Additionally, double counted electrons will lead to error inflation in each bin of the lifetime histogram, which must be corrected before fitting. This is discussed further in Appendix E. The two dimensional histogram in figure 9.22 shows the time correlation of electrons against the electron time relative to the muon for all events with two electrons in the $[0,25] \mu s$ event window. Using this histogram, a cut of ± 10 ns can be used to select events with two electrons in prompt coincidence. A comparison of lifetime histograms with and without this cut as well as the fitted rates for each dataset can be seen in figure 9.23. The Michel component of the two electron distribution is highly suppressed when the prompt events are removed. Furthermore, the prompt coincidence electron spectra falls with the muon lifetime, which is indicative of Michel related background, as expected. The flat background is reduced by roughly 15% with the prompt cut, indicating some additional level of prompt events from cosmic ray muons. As the double electrons in prompt coincidence fall with the nominal Michel electron lifetime, including these events in the multiple electron distributions stabilize the fitted rates and suppress any effects caused by the non-prompt multiple electron events. In order to maximize the enhancement of electron backgrounds in selecting multiple electron events, the prompt coincidence distribution is cut from $2e+$ distributions for the remainder of this chapter.

After pulsing A 100 ns deadtime is applied to each eSC PMT to mitigate any after pulsing effects, as described in section 7.4. Although an after-pulsing effect can be clearly seen in the single hit correlations of the eSC, this effect is suppressed once two fold eSC coincidences are created for single eSC layers, and appear to be eliminated in the four fold correlation spectra, as displayed in figures 7.22 and 7.23 of section 7.4.2. However, as an additional check, we can fit a lifetime histogram created specifically from events with two





Figure 9.22: Time difference between the electrons vs the electron time relative to the muon for all electrons within a two electron event. The band which runs parallel with the y-axis contains events where the first electron is the Michel and the second is an accidental. The opposite is true for the diagonal band. Each band exhibits the nominal muon decay in the direction perpendicular to the band. The intersection contains events with two electrons in prompt coincidence.



(a) Lifetime Comparison

(b) Fitted Rates

Figure 9.23: Lifetimes and fitted rate comparing the two electron lifetime histograms with and without the prompt cut. A start time of $1\mu\text{s}$ was used for the lifetime fit.

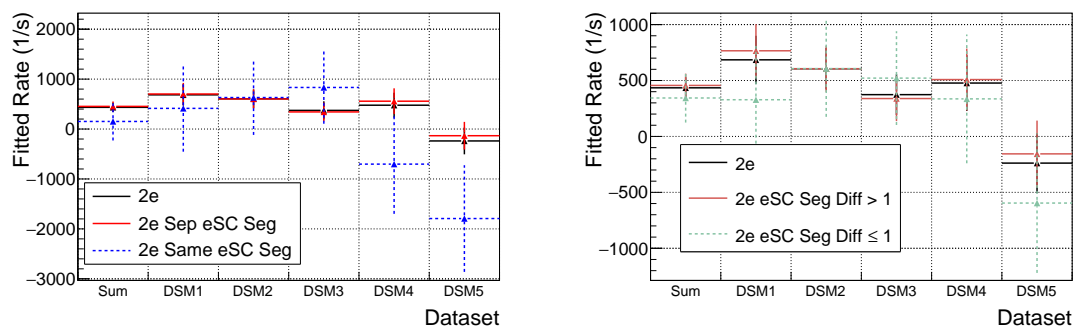
electrons in the same eSC segment, and compare to events with two electrons in separate eSC segments. Selecting events with two electrons in separate eSC segments should eliminate any after pulsing effects, whereas selecting events which contain two electrons in the same eSC segment should enhance such physical detector effects. The fraction of all events with two electrons, and the fraction of two electron events in which both electrons are seen in the same eSC segment can be seen in table 9.8. For all datasets, roughly 1/16th of the two electron events occur in the same gondola, as expected in the absence of afterpulsing effects. A scan

Dataset	Fraction of events with 2e-	Fraction of 2e- events in same eSC segment
Sum	0.015	.063
DSM1	0.027	.062
DSM2	0.027	.062
DSM3	0.026	.063
DSM4	0.0089	.063
DSM5	0.0055	.063

Table 9.8: Comparison of double electron event statistics for each dataset

over the lifetime fit parameters for all two electron events, and events with the eSC segment restrictions applied can be seen in figure 9.24(a). As there is no significant rate change in selecting events with two electrons in the same eSC segment compared to two electrons in separate segments, any after-pulsing which resides in the data likely has little effect on the observed distributions. Similarly, we can investigate the effect of cross talk between eSC segments by comparing the fitted rates of two electron events with electrons separated by at least one gondola segment to events with two electrons in the same or adjacent segments. This comparison is illustrated in figure 9.24(b), and shows no change when electrons in

adjacent eSC segments are excluded. It is therefore concluded that after-pulsing and cross talk effects are not responsible for the observed shift in the two electron lifetimes.



(a) After-Pulse Scan

(b) Cross-Talk Scan

Figure 9.24: (a) Scan over the fitted lifetime rate for all two electron events, two electron events in separate eSC segments, and two electron events in the same eSC Segment to investigate effects of after-pulsing. (b) Scan over the fitted lifetime rate for all two electron events, two electron events separated by at least one eSC segment, and two electron events in the same or adjacent eSC Segments, to investigate the effects of cross-talk between eSC segments.

9.2.6 Characterizing the Beam Electron Background

Directional Dependence of Electron Background

In trying to uncover the source of the beam electron background and build a model such that the effect on the final result can be mitigated and quantified, it is useful to study any directional dependence of the beam electron background. Here, it is useful to employ the cylindrical wire chambers, as a finer position resolution can be obtained with the electron tracking from the proportional chambers. The electron distributions in the azimuthal and z directions for both the standard scintillator only tracks, and the tracks including the wire chambers versus the decay time of the electron can be seen in figure 9.25. For the track

definition which includes the ePCs, the theta dependence is also shown. In order to enhance the effect of the electron background, events with two or more electrons are selected for each distribution.

Figure 9.26 shows the projection for each two-dimensional histogram in figure 9.25 on to the y-axis for both early and late times. At late times, the background contribution should be enhanced relative to the Michel distribution. Figures 9.26(a) and 9.26(c) show the distribution of electrons in the azimuthal direction for both electron definitions at early times and late times. There are slightly more downward going electrons at late times, but otherwise the phi distribution is relatively uniform. The deficit observed at values of $\phi = 3$ is due to the section of ePC 2 which contained a broken wire and was shunted to zero voltage for the duration of the data collection period. The electron distribution in Z, displayed in figures 9.26(b) and 9.26(d), is slightly asymmetric, tending towards more downstream values for both the eSC and ePC track definitions at early times. At later times, this distribution appears bifurcated, disfavoring tracks in the center of the electron barrel. The theta distribution remains uniform until late times, revealing a dramatic forward going nature of the beam electrons, increasing by an order of magnitude between forward and backward traveling tracks.

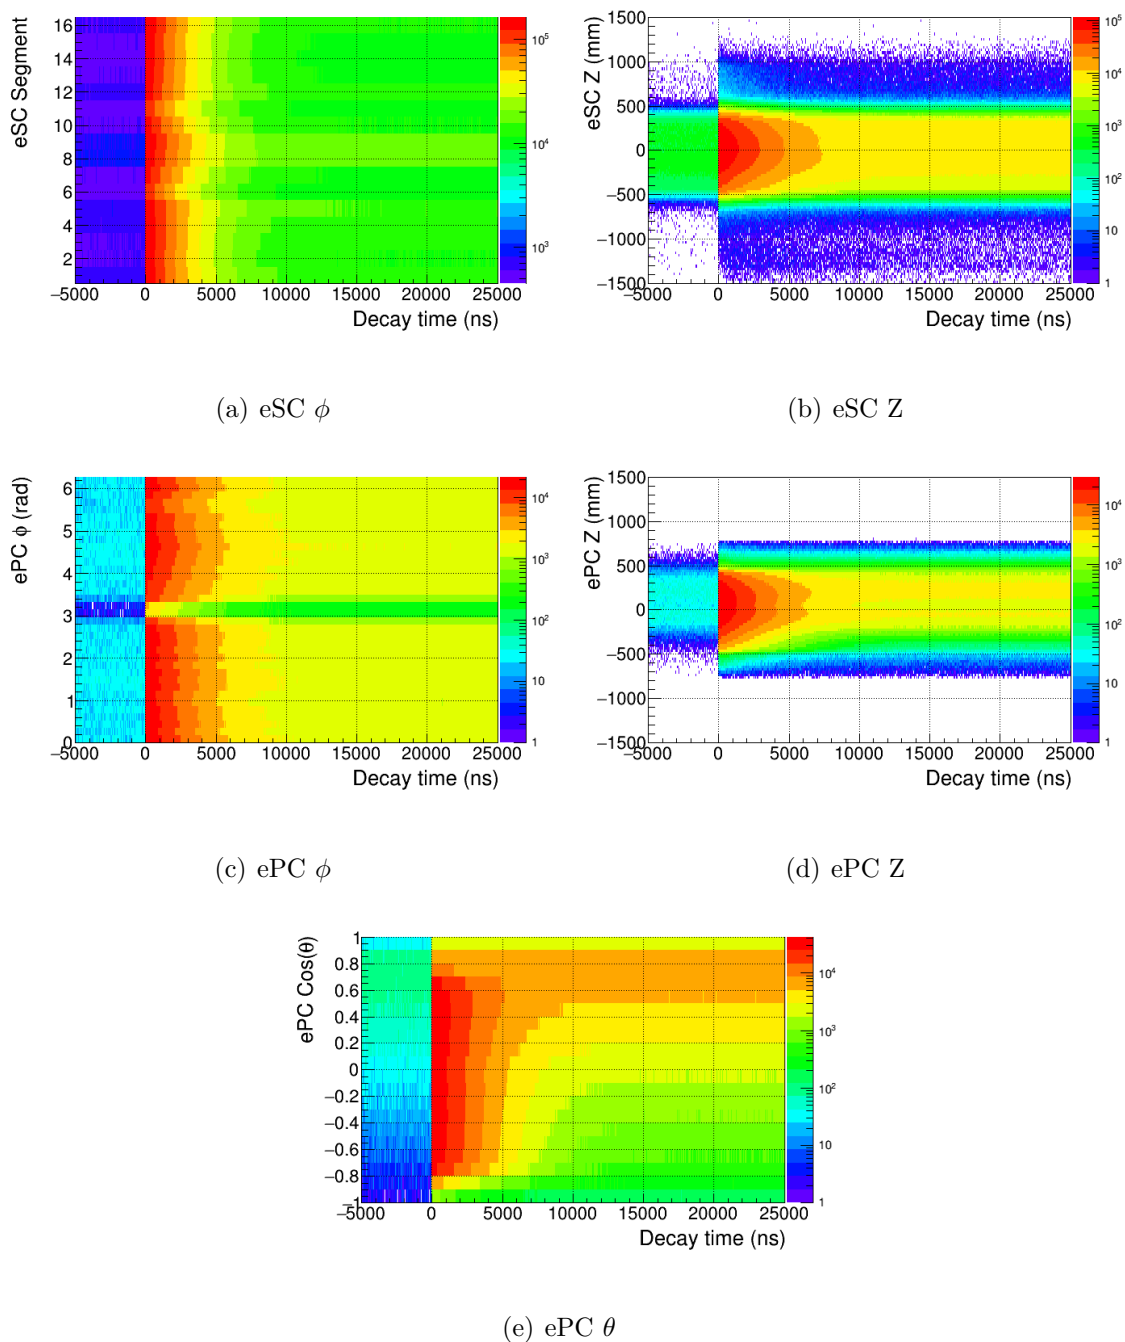


Figure 9.25: Two dimensional histograms of the two electron decay distribution versus the azimuthal and Z direction for both the eSC-only definition of the electron and that including the wire chambers. For the definition in which the ePC is included, the theta dependence can also be seen.

The correlation between the z position and the polar angle of all two electron tracks using the ePC electron definition can be seen in figure 9.27, for both early and late times. The electron background at late times exhibits increased forward going electrons, with $\text{Cos}(\theta) > 0.6$ which span a wide range of Z values, and bifurcate.

As it is clear that the electron background exhibits some directional dependence, it is of interest to scan the lifetime fit parameters as a function of eSC segment and electron Z position. A scan of lifetime fit parameters over the eSC segment can be used to observe any ϕ dependence on the beam electron background. This can be seen for the sum of all datasets in figure 9.28 for both the single and multiple electron lifetime distributions. In both cases, the fitted rate and background are stable as a function of the eSC segment.

As the electron background distribution displayed in 9.26(b) is bifurcated, any time dependence of this background should be more prevalent at certain z values, leading to an inconsistent fitted rate as a function of Z . The fit parameters comparing the single and double electron distributions as a function of the electron Z position in the eSC for the sum of all datasets can be seen in figure 9.29. Both the single and two electron distributions yield an increased fitted background parameter per muon stop at high and low values of Z , as expected. Although there is a potential outlier for electrons observed within -350 to -250 mm, there is no general trend observed in the fitted rate as function of the electron z position. The higher backgrounds at extreme values of z appear to have no effect on the fitted rate.

Time Dependence

The observed difference in the fitted rate between single and multiple electron events after $1 \mu\text{s}$ implies a time dependence other than that created by the 800 ns kicker step. By scanning the start and stop time of the lifetime fit, the lifetime pull as a function of time due to additional background effects can be explored. The value of the fitted rate and the fit chi

squared for both single and multiple electron events as a function of the start and stop time of the lifetime fit for the electron definition which does not include the wire chambers can be seen in figure 9.30. The red bands display the allowed statistical uncertainty, taking into account the complete correlation between each fit. Consistent fit parameters are obtained for all values of the fit start and stop time for the single electron distribution. The two electron distribution shows a significant downward pull on the fitted rate, particularly at start times after $4 \mu s$ and stop times after $10 \mu s$.

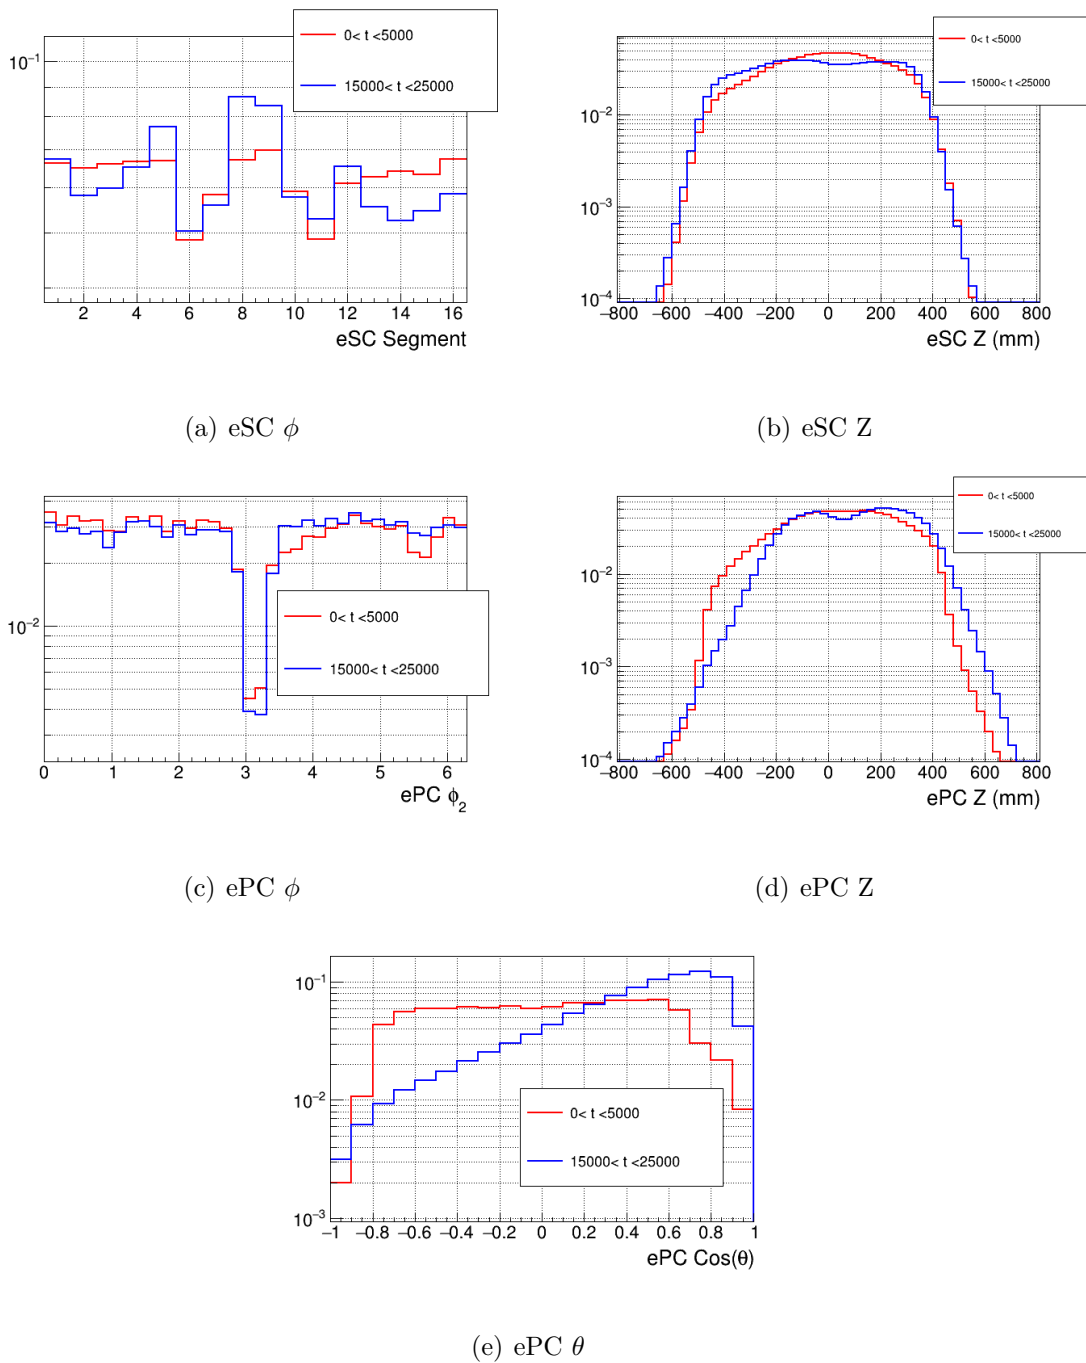


Figure 9.26: Projection of the two dimensional histograms of the two electron decay distribution displayed in figure 9.25 for early ($0 < t < 5$) μs (red) and late ($15 < t < 25$) μs (blue) electron arrival times.

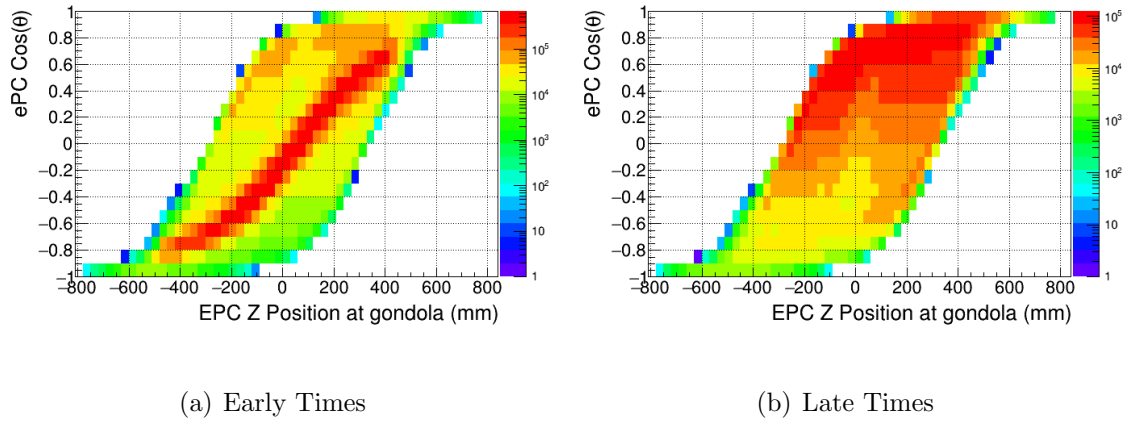


Figure 9.27: The ePC electron Z position versus polar angle for events with two electron tracks at (a) times within $5 \mu s$ after the muon entrance and (b) electrons arriving $15-25 \mu s$ after the muon entrance.

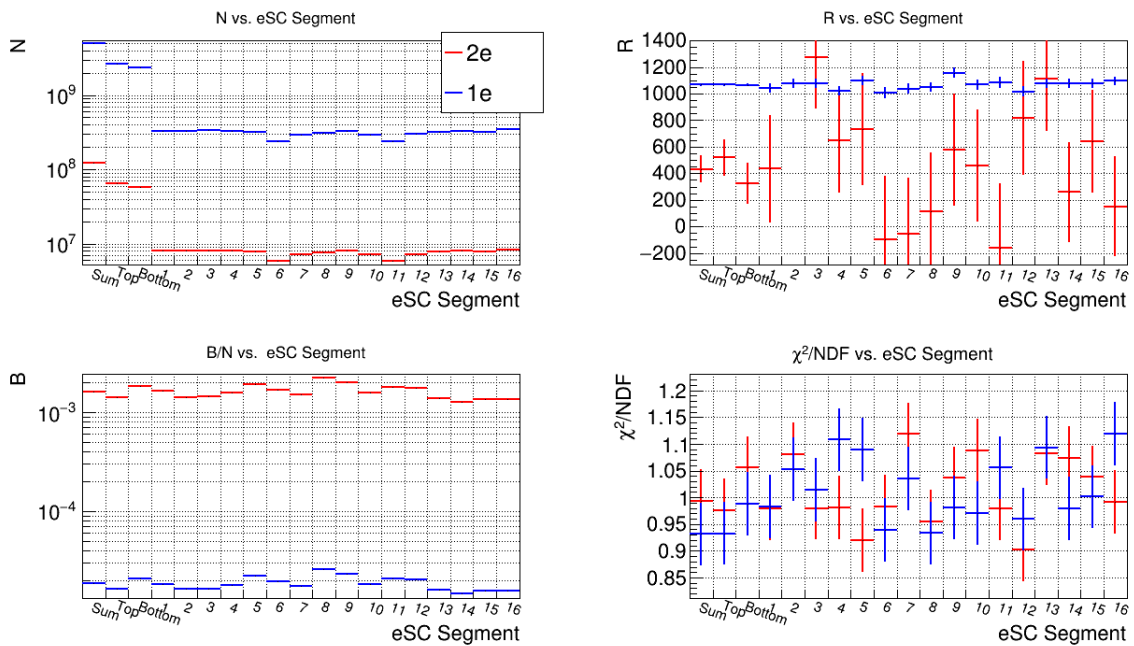


Figure 9.28: Lifetime fit parameters versus eSC segment.

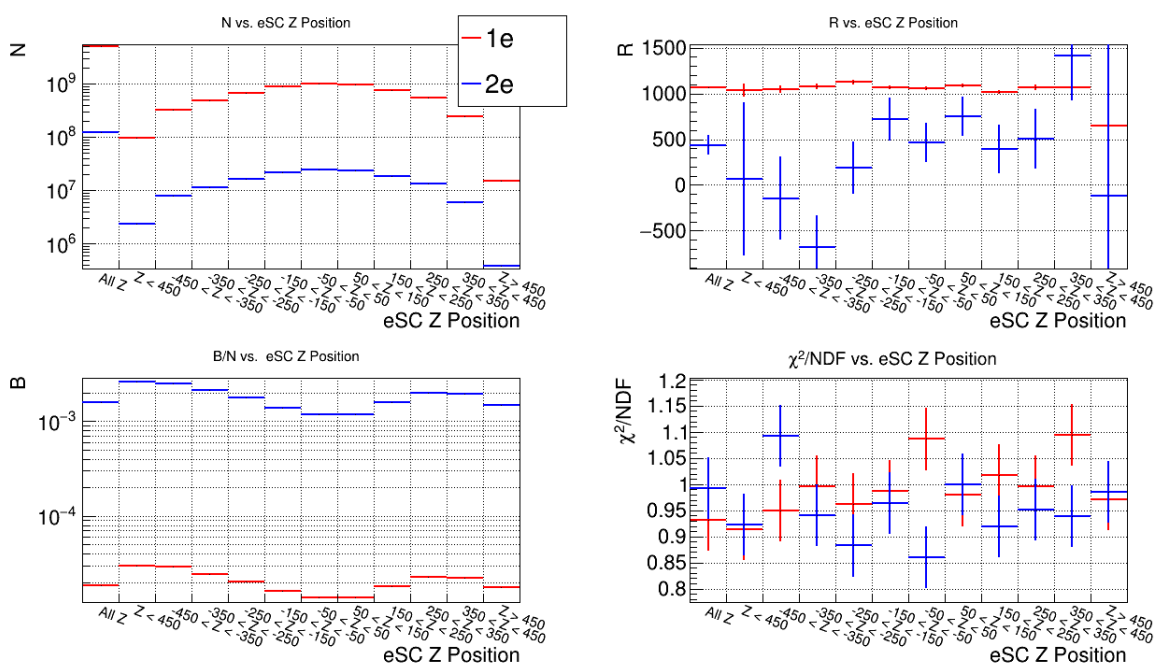
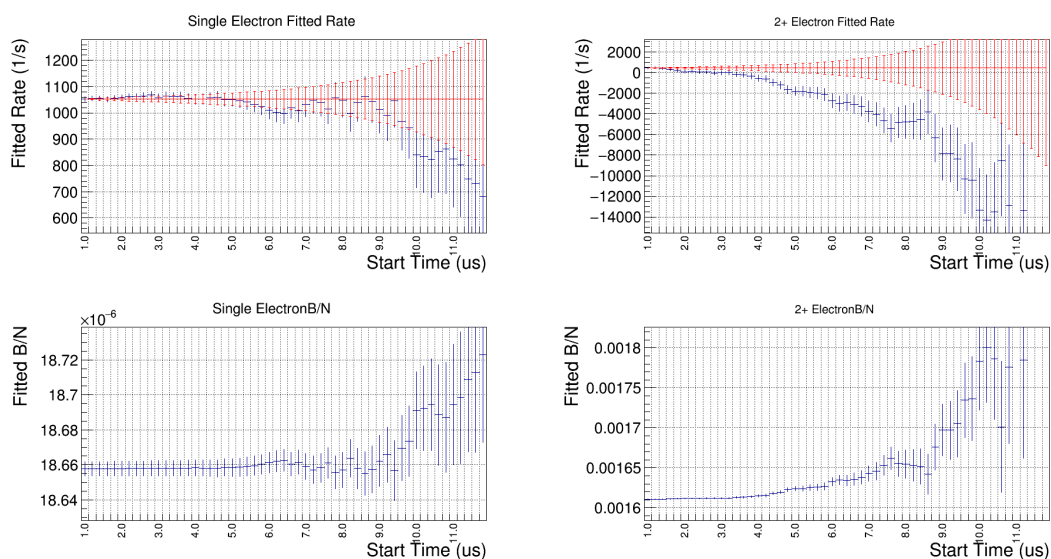
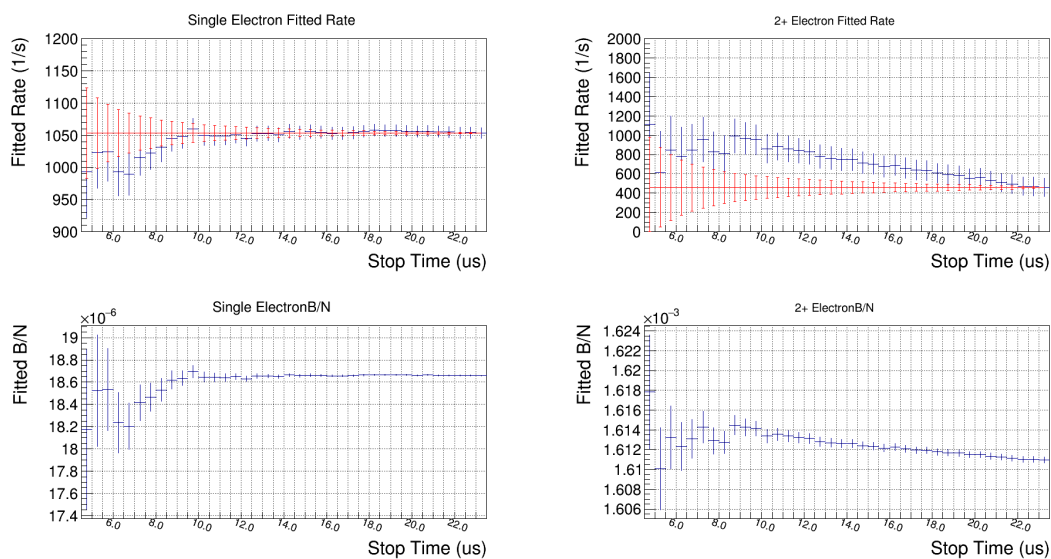


Figure 9.29: Scan of the fitted decay and background rate as a function of the electron Z position within the eSC, shown for both the single electron and double electron distributions.



(a) Start Time Scan



(b) Stop Time Scan

Figure 9.30: Scan of the fitted rate as a function of the (a) start time and (b) stop time of the lifetime fit function. On the left is the scan of the single electron distributions, showing fitted values within the allowed statistical bands (red) for all start and stop times. On the right is the same for the 2+ electron events, showing a sensitivity to both the start and stop time of the fit.

9.3 Background Model

From the observations and characterization of the data in the entrance efficiency analysis and the electron background studies, a model to characterize the effects of the MORE systematics was developed by the UW group. The model defines four predominant muon stop configurations, which are shown in the simplified diagram of figure 9.31.

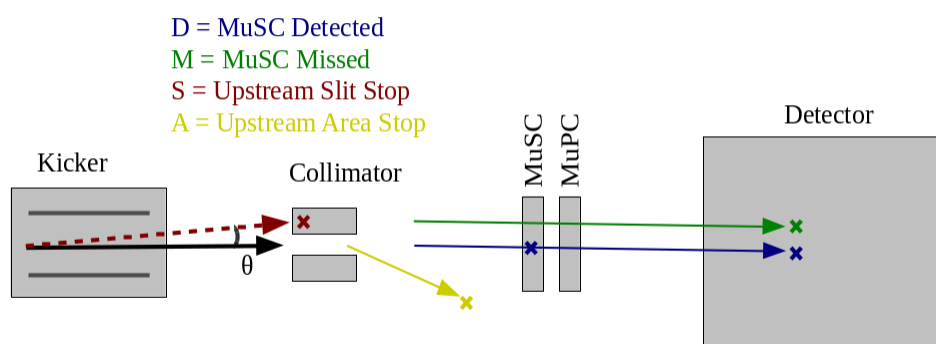


Figure 9.31: Simplified diagram of the beamline elements indicating the four potential muon stop configurations relevant for the MORE background studies.

1. μD Muons which pass through and are Detected by the muSC
2. μM Muons which pass through the muSC but are Missed
3. μC Muons which stop in the upstream beam Collimator
4. μA Muons which stop in materials in the Area upstream of the muSC

Each term is modeled based on the observations in the preceding sections.

The first two distributions, μD and μM , are characterized in the entrance efficiency studies of section 9.1.2. Muons which are detected by the muSC are not a problem, as any secondary muon within the $+/- 25 \mu s$ observation window will be eliminated by pileup protection. The second effect is fully constrained with the zero extrapolation of section 9.1.2. In any case, the next neighbor distributions for both of these terms can easily be written in terms of the unkicked beam rate, kicker extinction, and entrance detector inefficiencies determined from the data. The observed electron distributions from these terms are modeled by convolving the muon distribution with an exponential distribution falling with the nominal muon lifetime.

Based on the difference between the behavior of the step within the datasets before and after the adjustment of the upstream beam collimator, it is hypothesized that the 800 ns step is an artifact of the third stop configuration listed above, μS . For the early datasets, in which the collimator was misaligned or removed, the background rate decreases after the kick, suggesting the beam electrons were poorly separated until diverted by the kick, thereby reducing the background rate. In the later datasets, where the collimator is properly positioned, the background rate instead increases once the kick is applied. With the collimator properly in place, the electron separation is sufficient regardless of the kick configuration, eliminating the downward step seen in the background. With the high beam electron background suppressed, the small increase in the background could be due to decay electrons from muons stopping in the beam collimator. This systematic can be constrained by comparing the multiple electron distribution to the single electron distribution, where the effect is suppressed by nearly two orders of magnitude. Additionally, as long as muon decays within the collimator are fast, lifetime fits starting at $1 \mu s$ should evade any effect. The muon and electron time distributions for this effect are derived assuming the muons stop in copper, as this was the main component of the upgraded beam collimator.

The most difficult stop configuration in terms of systematic uncertainty is $\mu\mathbf{A}$, the muons which stop in some unknown material in the area without passing through the muSC. Returning to the formalism discussed in section 9.2.2, the enhancement of both the $\mu\mathbf{S}$ and $\mu\mathbf{A}$ effects in the multiple electron distribution compared to the single electron distribution can be considered. Assuming a time distribution of $e_x(t)$, where x represents either the distribution from the collimator stops or the upstream area stops, equations 9.13 and 9.15 obtain additional terms and are given by

$$e_1(t) = (1 - P_X) \cdot (1 - P_a) \cdot e(t) + (1 - \epsilon) \cdot (1 - P_X)R_B(t) + (1 - \epsilon) \cdot (1 - P_a) \cdot e_x(t) \quad (9.20)$$

for the single electron distribution and

$$e_2(t) = (P_X + P_a) \cdot e(t) + (P_X + \epsilon) \cdot R_B(t) + (P_a + \epsilon) \cdot e_x(t) \quad (9.21)$$

for the two electron distribution, where $P_X = \int e_x(t) \cdot dt$ is the probability for seeing the electron from either $\mu\mathbf{A}$ or $\mu\mathbf{S}$. The contribution from the $e_x(t)$ component relative to the real muon lifetime component is given by

$$\chi_1 = \frac{(1 - \epsilon)}{1 - P_X} \quad (9.22)$$

or

$$\chi_2 = \frac{(P_a + \epsilon)}{(P_a + P_X)} \quad (9.23)$$

for the single and multiple electron distributions, respectively. We therefore expect to see an enhancement by a factor of

$$\chi = \frac{(P_a + \epsilon)}{P_a + P_X} \cdot \frac{(1 - P_X)}{(1 - \epsilon)} \quad (9.24)$$

in the two electron distribution compared to the single electron. This is generically the same enhancement factor found in equation 9.18, as long as $P_x \ll P_a$. Taking the average beam electron background rate of 1100 Hz, and assuming an upstream stopping rate which is smaller, an enhancement factor of $\chi = 80$ is obtained. Any remaining effects from $\mu\mathbf{A}$


seen in the multiple electron distribution are therefore expected to contribute to the single electron lifetime with 80 times less severity. The muon distributions for this term are modeled assuming a rate of upstream stops of $R_x=100$ Hz. The electron distribution is obtained by convolving with an exponential which falls with the lifetime of either aluminum, iron, or carbon.

The electron distributions for each of the terms modeled based on rates from DSM5 can be seen in figure 9.32 before and after pileup protection is applied.


As discussed in section 9.1, any of the muons in the Missed, eM, (blue) distribution of figure 9.32 can distort the lifetime if they stop in materials outside of the deuterium target. However, as shown in the second panel, this term is highly suppressed once pileup protection is applied, and any remaining effects are constrained by the zero extrapolation of section 9.1.2. Muons which stop in the upstream area, eA, shown as the yellow distribution in figure 9.32, mostly likely present with a fast lifetime component, such as iron, after 800 ns which is greatly enhanced in the two electron distribution. In the data, the opposite is observed—the two electron distributions exhibit a slower fitted rate, as displayed in figure 9.21. It is thus concluded that upstream stops in the area are not a major component of the electron background in the R2014 dataset. The only term which could lead to a slower rate is the muons which stop in the collimator, eC, displayed in green. This term contributes to a slow buildup, which could result in a background which increases over time leading to a slower fitted rate. However, this is constrained in the data with the muon clock distributions in which no increase is observed over the $25 \mu s$ window with the given statistics.

9.4 Monte Carlo Background

As a complement to the simplified background model developed in the previous section, used to characterize the effects from pileup detector inefficiencies and the beam electron background induced by the MORE, it is of interest to quantify any remaining effects using



(a) Nominal



(b) Pileup Protected

Figure 9.32: Time distributions of the modeled background terms (a) before and (b) after pileup protection is applied.

the full MuSun Monte Carlo. The MC simulation does not contain any pileup muons or electron background associated with the beam. As such, the MC simulated data will contain only effects from physical processes unrelated to the MORE signal. The fitted rate obtained for various electron multiplicities for the three Monte Carlo datasets can be seen in figure 9.35. The lifetime histograms were created for this dataset using the true TPC stop position, to

remove any affects from fusion interference, but the electrons are the reconstructed electrons. The true rate expected is shown as a dashed purple line. The fits and residuals for the full




Figure 9.33: Fitted rates for each MC dataset with a 1us start time for all electron (black), single electron (red) and double electron (blue) event selections. The expected true MC rate is shown as a dashed purple line.

MC data set starting at $1 \mu s$ for the prompt subtracted two electron event can be seen in figure 9.34. The dramatic change in shape and deviation of the fit residuals at early times is indicative of muon catalyzed fusion, suggesting that fusion neutrons are observed in the eSC.

As the particle type is available with the MC data, a separate set of lifetime histograms can be created selecting only the reconstructed electrons which are coincident with a true Michel electron. In this case, events with more than one electron are eliminated, as expected, and the fitted rate agrees with that expected by the MC input, as displayed in figure 9.35. The two main particles other than the Michel electron which create four-fold events are positrons, likely from Bremsstrahlung splitting, and neutrons.

It is difficult to relate the MC studies presented here to the data. Although the deviation of the two electron lifetime from the single electron is lower, as observed in the data, there is an obvious early time effect from the neutrons in the MC simulations in the decay distribution




Figure 9.34: Lifetime fit for the full (a) single electron and (b) two electron MC dataset, with a fit start time of $1 \mu s$. These lifetime histograms were created using the MC true stop position, to eliminate fusion interference effects.




Figure 9.35: Fitted rates for each MC dataset with a $1 \mu s$ start time for Michel selected electrons. The expected true MC rate is shown as a dashed purple line.

of figure 9.34 which is not seen in the R2014 dataset, indicating that either the threshold for neutrons in the eSC is significantly lower in the MC or that the effect is enhanced in the MC because there is no stabilizing accidental background. Further studies are needed to ascertain the experimental threshold for neutrons to create four fold events in the eSC

detectors. This effect will be better constrained in later analyses, where the neutron detection will be compared by selecting events with an associated ${}^3\text{He}$ in the TPC. However, it is interesting to note that a correlated background term which was not accounted for in the development of the background model, can give rise to the behavior seen in the two electron distribution of the R2014 dataset.

9.5 Conclusions

Entrance efficiency studies in combination with a background model developed via studies of the R2014 data is used to constrain lifetime shifts and uncertainties due to the MORE signal from the contributions displayed in table 9.9, assuming a $1\ \mu\text{s}$ fit start time.

In addition, an effect from fusion neutrons observed in the electron scintillators is studied and constrained via Monte Carlo simulations. The lifetime shift and uncertainty due to pileup muons missed by the muon entrance detectors, $\mu\mathbf{M}$ are determined via a zero extrapolation of the lifetime versus detector inefficiency for several combinations of the entrance detectors, described in section 9.1.2. The contributions from the muons which stop in the beam collimator, $\mu\mathbf{C}$, and in any other upstream material in the area, $\mu\mathbf{A}$ are constrained via studies of the two electron distribution, which magnify such effects by a factor of 80 compared to the single electron distribution. Conservative upper limits on the accidental rate based on the muon clock are assumed to determine the effect on the lifetime for these terms. Although the background model was unable to reproduce the 600 Hz lower rate obtained with the two electron distribution, it is estimated that this effect is generically suppressed by a factor of 80 in the single electron distribution, yielding a 7.5 Hz effect. Monte Carlo simulations in which fusion neutrons are observed in the eSC reproduce the general trend of the reduced two electron rate, and may indicate a correlated term which was not considered in the background model.

Background Contribution	Lifetime Shift (s^{-1})	Shift Direction
μM	< 1.54	+
μA	< 0.25	+
μC	< 0.25	-
$2e$	< 7.5	-

Table 9.9: Table of the systematic errors associated with the MORE accidental structure.

Chapter 10

MUON CATALYZED FUSION INTERFERENCE CORRECTION

10.1 Fusion Formalism

The mechanism in which muon-catalyzed fusion products lead to muon stop mis-reconstruction is detailed in section 8.2.1. In order to establish the effect of the acceptance or rejection of such events on the observed decay distribution, it can be written in two components

$$e(t) = e^f(t) + e^{nf}(t). \quad (10.1)$$

where $e^f(t)$ represents all events which had a prior fusion and $e^{nf}(t)$ is the converse. In order to construct these two components, the conditional probability of a fusion at time t_f followed by a muon decay at time t must be determined,

$$e^c(t, t_f) = P_f(t_f)P_e(t - t_f), \quad (10.2)$$

The probability distribution (pdf) for a muon decay at time t takes the following approximate form

$$e_0(t) = \lambda_0 e^{-\lambda_0 t}. \quad (10.3)$$

The observed fusion pdf can be written as

$$f_{obs}(t) = f(t)e^{-\lambda_0 t} \quad (10.4)$$

where $f(t)$ is the pdf of a fusion given that the muon survives until t , which can be obtained by solving the system of differential equations presented in appendix A with λ_0 set to zero.

The conditional probability of a decay following a fusion can then be written

$$\begin{aligned}
 e^c(t, t_f) &= P_f(t_f)P_e(t - t_f) \\
 &= f_{obs}(t_f)\lambda e^{-\lambda_0(t-t_f)} \\
 &= f(t_f)\lambda e^{-\lambda t}.
 \end{aligned}
 \tag{10.5}$$

The pdf of a muon decay at time t with and without a prior fusion is then given by

$$\begin{aligned}
 e^f(t) &= \int_0^t e^c(t, t_f) dt_f \\
 &= e_0(t) \int_0^t f(t_f) dt_f
 \end{aligned}
 \tag{10.6}$$

$$\begin{aligned}
 e^{nf}(t) &= e_0(t) - \int_0^t e^c(t, t_f) dt_f \\
 &= e_0(t) \left(1 - \int_0^t f(t_f) dt_f \right)
 \end{aligned}
 \tag{10.7}$$

respectively. Because the muon must live long enough to catalyze a fusion, events with a prior fusion have a later decay time on average. The decay time distribution of all events, non-fusion events and fusion events can be seen in figure 10.1. Due to the non-exponential nature of $e^{nf}(t)$, any change in the acceptance of fusion events leads to a distortion on the measured rate. If α is defined as the nominal efficiency for accepting a muon stop and $\alpha_f(t)$ as the acceptance efficiency for muons with a fusion event, the observed distribution of muon decay events with and without a prior fusion is revised to

$$\begin{aligned}
 e_{obs}^f(t) &= \alpha e_0(t) \int_0^t \alpha_f(t_f) f(t_f) dt_f \\
 e_{obs}^{nf}(t) &= \alpha e_0(t) \left(1 - \int_0^t f(t_f) dt_f \right)
 \end{aligned}
 \tag{10.8}$$

and the total observed electron distribution is given by

$$\begin{aligned}
 \tilde{e}(t) &= e^f(t) + \tilde{e}^{nf}(t) \\
 &= \alpha e_0(t) \left(1 - \int_0^t \eta(t_f) f(t_f) dt_f \right)
 \end{aligned}
 \tag{10.9}$$




Figure 10.1: Comparison of the time distribution for events with and without a prior fusion.

where the fusion correction function has been defined as

$$\eta(t) \equiv 1 - \alpha_f(t). \quad (10.10)$$

This function quantifies the change in acceptance efficiency for events with and without a prior fusion, and is expected to differ from zero at early times, when the fusion products can interfere with muon tracking, and fall to zero as fusions become separated enough in time from the muon stop that they cannot interfere.

10.1.1 Characterization of the fusion correction function

There are several scenarios where the acceptance of an event is altered by the presence of a fusion

1. Migration events lost or gained at the fiducial volume (FV) boundaries, as illustrated in Sec 8.2.1. The presence of fusion products within the TPC can lead to mis-reconstruction of the muon track, pushing events inside or outside of the fiducial volume. This is the dominant effect.

2. Changes to muon detection efficiency within the fiducial volume. The addition of fusion products within the TPC can lead to changes in event topologies and energies within the fiducial volume. To avoid coupling event acceptance to these changes, very loose cuts are applied to the muon stop.
3. Coupling to electronic inefficiencies. Fusions could increase the acceptance of events by filling gaps in pads with low thresholds, where a muon pulse might otherwise be lost. This mechanism is studied by comparing lifetimes including and excluding events with gaps in the z tracks, and the effect is determined to be insignificant.

The fusion correction function has been studied with Monte Carlo simulations in the thesis of Michael Murray [38]. The findings presented in detail in his thesis are briefly summarized here. As expected, migration occurs at early times when the fusion products may interfere with the muon tracking and falls to zero as the fusion comes late enough in time relative to the muon stop. Additionally, the fusion acceptance can be written as a magnitude and time dependent shape

$$\eta(t_f) = \eta_0 \bar{\eta}(t). \quad (10.11)$$

The shape of $\bar{\eta}(t)$ is independent of the choice of volume, while the magnitude η_0 varies with variable migration fractions. For small η_0 , the change in the observed decay distribution due to fusion interference is then approximated by a linear relation such that the lifetime change due to fusion interference is proportional to the correction function

$$\Delta\lambda \propto \eta_0. \quad (10.12)$$

10.1.2 Lifetime correction strategies overview

At this point, it is established that the observed shift in the lifetime due to fusion effects is proportional to $\bar{\eta}(t)$. In order to develop a correction strategy, this must be related to physical observables. The main mechanism which produces a non-zero $\bar{\eta}(t)$ is fusion migration. The

distortion to the observed rate must therefore be proportional to the number of migration events, M .

$$\Delta\lambda \propto \eta_0 \propto M \quad (10.13)$$

The net migration can be quantified using several observables. For the rest of the chapter, observable quantities are denoted with a tilde, while untilded variables refer to the expected value in the absence of fusion migrations.

First, M is equal to the difference in the expected and observed number of fusion events. Relating to the formalism developed in the previous section,

$$M \equiv F - \tilde{F} = \alpha N \int_0^T f(t_f)(1 - \alpha_f(t_f))dt_f \quad (10.14)$$

where N and F are the true number of muon stops and fusions in the absence of migrations, respectively, and \tilde{F} is the observed number of fusions including migration effects. The integral includes times up to the point where $\eta(t)$ drops to zero at time T .

This can also be phrased in terms of the true and observed fusion fractions

$$\epsilon = \frac{F}{\alpha N} = \int_0^T f(t_f)dt_f \quad (10.15)$$

$$\tilde{\epsilon} = \frac{\tilde{F}}{\alpha N} = \int_0^T f(t_f)\alpha_f(t_f)dt_f. \quad (10.16)$$

The change in the measured rate due to fusion migrations, which is proportional to $\eta(t)$ and thus scales with the number of migrations, can now be expressed in terms of the total migrations, or the difference between the true and observed fusion fractions

$$\Delta\lambda = \kappa \left(1 - \frac{\tilde{\epsilon}}{\epsilon}\right) = \kappa \frac{M}{F}. \quad (10.17)$$

where κ is a constant of proportionality which must be determined empirically. The effect on the lifetime can then be determined either by asserting the ratio of the true and observed fusion fractions, $\frac{\tilde{\epsilon}}{\epsilon}$, or, determining the number of migrations. The uncertainty of such a determination is given by

$$\delta(\Delta\lambda) = \delta\kappa \frac{M}{F} \oplus \kappa \delta\left(\frac{M}{F}\right). \quad (10.18)$$

The constant κ is given by the shape of the $\eta(t)$ function. In order to reduce the uncertainty, $\frac{M}{F}$ should be minimized via robust tracking algorithms.

Three potential correction strategies are presented, although only one method is explored in detail in this thesis.

Monte Carlo Migration Determination

The first attempt to correct for fusion interference, detailed in the thesis of Michael Murray[38], sought to calculate the number of migration events using a Monte Carlo formulation of the migration probability coupled with the muon stop distribution observed from data. For a given z row in the TPC, the number of migrations can be written

$$M_i = \alpha N \epsilon \int_{z_i}^{z_i+L} p^{mig}(z) n(z) dz \quad (10.19)$$

where $p^{mig}(z)$ is the probability of migration as a function of the z position within a pad, $n(z)$ is the normalized muon stop distribution, L is the pad length, and ϵ is the fusion fraction defined in equation 10.15. The migration probability is calculated via Monte Carlo simulations, and the muon stop distribution estimated from data. The precision of this method is limited by the uncertainty in interpolating $n(z)$ within each pad, extrapolating $n(z)$ beyond the fiducial boundary, and the reliance on Monte Carlo. Additionally, this method neglects potential effects in the X and Y dimensions.

Fusion Fraction Determination

In this thesis, a data driven method is developed based on a new idea. The muon stop distribution is shaped using a series of cuts on the energy deposition within the first pad row, suppressing the events at the fiducial volume boundaries. The shaped distribution highly suppresses migration due to fusion, and can be used to determine the fusion fraction in the absence of interference effects.

As illustrated in equation 10.17, the lifetime shift from fusion migration can be quantified by determining the ratio of the true and observed fusion fractions. There are then three steps to determining the lifetime correction due to fusion interference:

1. Determine the true fusion fraction in the absence of fusion migrations, ϵ
2. Determine the observed fusion fraction, $\tilde{\epsilon}$
3. Determine the constant of proportionality, k

The most difficult task lies in the first step. As fusion migration occurs only at fiducial volume boundaries, a stop distribution in which the number of stops on the fiducial boundaries is suppressed yields an approximation of the true fusion fraction, ϵ . The stop distribution can be shaped to suppress stops on the fiducial boundary by making X, Y, and energy cuts on the first pad row, which cannot be reached by fusion products. For this shaped distribution, the fusion fraction is defined

$$\epsilon_{sh} = \frac{F_{sh}}{\alpha_{sh}N_{sh}}. \quad (10.20)$$

such that the difference between the true and shaped fusion fraction is

$$\epsilon_{sh} - \epsilon = \frac{F_{sh}}{\alpha_{sh}N_{sh}} - \frac{F}{\alpha N}. \quad (10.21)$$

where α is the efficiency for detecting a muon stop. If we assume that the migrations are small in the shaped distribution, and the event reconstruction efficiency is unchanged by the shaped distribution, such that $\alpha_{sh} = \alpha$ we can assert that the fusion fraction is equivalent to the true fusion fraction, up to a small correction from the remaining migrations in the shaped distribution.

$$\epsilon_{sh} = \epsilon + \frac{M_{sh}}{\alpha N} \quad (10.22)$$

For a given stop distribution, we have the following observables:

- $\tilde{F} + B_F$, the observed fusions in the presence of migrations
- $\tilde{N} + B_\mu$, the observed muon stops in the presence of migrations.

where B_f and B_μ are background events which are misinterpreted as a muon stop or a fusion, respectively. As three consecutive pad rows in z are required for a muon stop, the likelihood of false muon tracks is low, and B_μ is taken to be zero. The fusion background term B_f is comprised of muon events with extra energy, such that the event is incorrectly tagged as a fusion. This is found to occur mostly from deuteron recoil events following a fusion-neutron scatter. A background subtraction method, described later in this chapter, is used to suppress this background, so B_f is also taken to be zero.

In order to determine $\tilde{\epsilon}$ as defined in equation 10.16, the true number of muon stops, N , is required. If it is once again assumed that α is equivalent in all stop distributions, an iterative approach can be used to obtain the true number of muon stops via

$$\begin{aligned} N &= \tilde{N} + M \\ &= \tilde{N} + F\left(1 - \frac{\tilde{\epsilon}}{\epsilon}\right) \end{aligned} \tag{10.23}$$

where $\tilde{\epsilon}$ is first obtained using \tilde{N} . We then find the true and observed fusion fractions from the physical observables to get the fusion fraction ratio

$$\frac{\tilde{\epsilon}}{\epsilon} = \frac{\tilde{F} N_{sh}}{N F_{sh}} \tag{10.24}$$

Finally, the value of κ is obtained by altering the distribution to obtain a variety of observed fusion fractions, determining the lifetime shift from each, and performing a linear fit.

This method is clearly advantageous in its purely data driven nature, and ease of uncertainty estimates. However, it relies upon the following assumptions:

- The detection efficiency of muon stops is the same in all distributions
- The detection efficiency of fusions is the same in all distributions.

As the shaping of the stop distribution necessarily shifts the fraction of stops across the pad rows of the TPC, any inconsistencies in efficiency between the pads of the TPC would lead to errors in the estimation of the number of muon stops or fusions. Thus it is essential to achieve uniform calibration across the TPC, and carefully identify any pads with lower thresholds or electronic issues leading to inefficiencies.

Migration Determination via Delayed ${}^3\text{He}$

A third method to determine the number of fusion migrations directly by using delayed ${}^3\text{He}$ as an indicator of the true muon stop position is currently being developed. As the range of the ${}^3\text{He}$ is rather limited, its position reveals the true muon stop position, and can be used to directly measure fusion migration. For all PT events followed by a delayed ${}^3\text{He}$, the reconstructed stop position is compared to the position of the delayed ${}^3\text{He}$, and any event in which the pad difference is non-zero is counted as a mis-reconstruction. Several new challenges need to be addressed with this method:

1. The larger undershoot from PT events in the electronics chain affects the efficiency for detecting a subsequent ${}^3\text{He}$ pulse. As the undershoot is proportional to the size of the preceding pulse, care must be taken to choose a time window in which the ${}^3\text{He}$ detection efficiency is not coupled to the PT energy deposition.
2. As events with both a PT and a ${}^3\text{He}$ fusion are selected, available statistics could be a potential issue. Integrating the kinetics suggests that 10^{10} fiducial stops should yield 5×10^5 PT fusion events followed by a delayed ${}^3\text{He}$ after $2 \mu\text{s}$. For 10% of stops in the fiducial volume boundary, this still allows for a fusion migration determination to a few percent
3. Events in which the track was mis-reconstructed due to non-fusion related effects need to be subtracted.

As the sources of uncertainty are completely different, this method would serve as a good complement to the fusion fraction method, and can be used for a completely independent determination of the migration correction, or determination of the small correction for the remaining migrations present in the shaped stop distribution.

10.2 Application of Fusion Fraction Method to Data


Three quantities are necessary in determining the lifetime shift via the fusion fraction method: the observed fusion fraction, $\tilde{\epsilon}$, the true fusion fraction, ϵ , and the constant of proportionality, κ . This section first describes how each of these necessary components is obtained. Then the method is applied to the R2014 dataset to obtain a lifetime correction.

10.2.1 Determining $\tilde{\epsilon}$

As shown in table 5.2, a total of 4000 keV is deposited in the TPC from a proton-triton fusion events, which makes tagging the events relatively easy. The E9 energy is defined as the sum of the energy deposited on the stop pad and the eight surrounding pads, as described in section 5.6.2. The distribution of the E9 energy for the R2014 dataset, can be seen in 10.2(a). Summing the surrounding pads ensures that the proton and triton energies are included, even if the stop pad is mis-reconstructed. A muon stop without pileup has an E9 energy between 1000 and 1800 keV, depending on the depth the muon travels into the final pad. With the addition of the proton and triton, events with a p-t fusion generate an E9 energy peaked around 5000 keV. Fusion events can thus be tagged by making an E9 energy cut in the region between the muon and p-t peaks.

To optimize the cut in E9 space, one must consider the background which lies between the muon and p-t peaks. The Monte Carlo E9 energy distribution including each event type can be seen in figure 10.2(b). The neutron from ${}^3\text{He}$ fusion can scatter off deuterium nuclei in the TPC gas, depositing a recoil energy up to 2200 keV. The efficiency of p-t tagging is

optimal with an E9 threshold around 3 MeV, but the n-d scatter events residing on the low energy side of the p-t tail lead to a background of 1-2%.



(a) R2014 E9 Distribution

(b) MC E9 Distribution

Figure 10.2: (a) E9 distribution for the full R2014 dataset and (b) Monte Carlo E9 energy distribution including each event type. The deuteron recoil distribution sums the n-d scatter events following a ${}^3\text{He}$ fusion.

The majority of fusion events occur within the clustering time of the tracking algorithm, such that any extra energy from fusion products or recoils is added onto the muon energy. However, a fraction of fusion events come delayed enough from the muon stop that an extra cluster is created from the additional pulses. By making a cut on the energy of the extra cluster, we can isolate the n-d scatter background for a subtraction. The extra cluster energy of delayed events is plotted versus the E9 energy for the R2014 dataset in figure 10.3(a). The various labeled regions correspond to the following event topologies:

1. Events with a delayed ${}^3\text{He}$ fusion. The two bands with extra cluster energy at 600 keV and 400 keV denote fusions with and without the muon sticking to the ${}^3\text{He}$, respectively.
2. Neutron scatter events, resulting in a recoil energy up to 2200 keV. The ${}^3\text{He}$ or $\mu{}^3\text{He}$ may deposit an additional 400-600 keV if also contained within the delayed cluster.

When summed with the muon stop, this leads to an E9 energy up to 4000 keV.

3. Delayed p-t fusion events. A delayed p-t fusion has cluster energies ranging from 2700-4700 keV, leading to E9 energies peaked around 5000 keV.
4. ${}^3\text{He}$ p-t events. The delayed cluster energy contains that of the ${}^3\text{He}$, but the addition of the p-t inside the muon cluster leads to an E9 energy above 5000 keV.
5. 2pt events. Events with two pt fusions, one within the cluster one outside. This results in an extra cluster energy from 3000-4000 keV, and an E9 energy above 7000 keV.

The projection onto the X axis, showing the E9 energy from all events in comparison to events with an extra cluster energy less than 2600 keV, is shown in figure 10.3(b). By making a cut on the extra cluster energy at 2600 keV, the p-t fusion events are suppressed and the high energy tail of n-d scatter background is obtained up to E9 energies of 4500 keV.

Applying this method to the R2014 dataset, the background subtracted E9 energy distribution shown in figure 10.4 is achieved. The fusion fraction is determined by

$$\tilde{\epsilon} = \frac{\tilde{F}}{N} \quad (10.25)$$

where \tilde{F} is the number of fusions above an E9 cut of 3300 keV in the background subtracted spectrum and N is the total number of muons, taken as the total integral of the E9 distribution without background subtraction.

E9 Scans to obtain uncertainty in $\tilde{\epsilon}$

As a check on the selection of the E9 cut at 3300 keV, the analysis is repeated with a range of E9 cuts and the effect on the final correction observed. The number of fusions in various 100 keV slices, normalized to the total number of fusions in the 3300 keV cut, for both the nominal and shaped stop distributions are shown in figure 10.5. The effect of the background

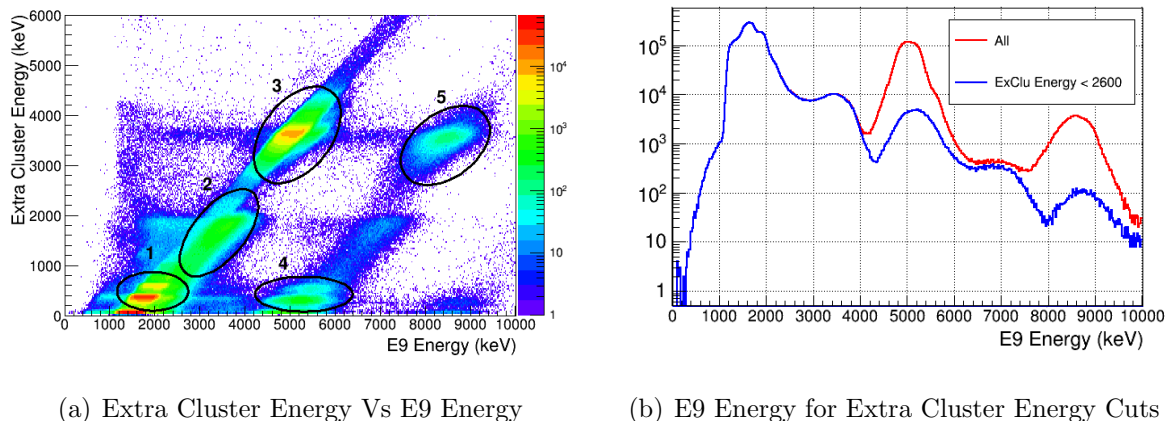


Figure 10.3: (a) Extra cluster energy vs E9 energy for extra cluster events. Events below a cut of 2600 keV on the Y-axis are used for a background subtraction, to eliminate the n-d scatter events in the E9 distribution used to tag PT events. (b) E9 energy distribution for all delayed extra clusters (red) and extra cluster with energy < 2600 keV (blue). Both panels are for the R2014 dataset.

subtraction is demonstrated in comparison of the red and blue curves. The cut of 3300 keV is chosen as the lowest point between the muon and fusion tails, such that the efficiency for tagging PT events is maximized while keeping the background from fusion-less events to a minimum. Figure 10.6(a) shows the fusion fraction obtained for several choices of the E9 cut. The effect of the background subtraction on the nominal stop distribution is seen in the difference between the red and blue curves, and the same for the shaped distribution for the green and black curves. With the background subtraction, the fusion fractions obtained are stable for cuts in the region from 2800 keV to 3800 keV to the 2×10^{-3} level. Similarly, figure 10.6(b) shows the variation in the ratio of the nominal and shaped fusion fraction as a function of the E9 cut. This is stable to the 10^{-3} level, even for the fusion fractions obtained without background subtraction, indicating that the nd scatter background is present in both the nominal and shaped distributions, leading to a cancellation in the ratio used for




Figure 10.4: Comparison of the full E9 distribution (black), reconstructed E9 background from neutron scatter events after applying the extra cluster energy cut to the data (blue), and the background subtracted version of the E9 energy distribution, to be used for determination of the fusion fraction.





Figure 10.5: Fraction of PT fusion events in each slice of the E9 distribution, before and after background subtraction, displayed for both the nominal and shaped stop distributions. The upper curves show the fraction of fusion events before background subtraction. The removal of background events mistagged as fusions via the subtraction technique is evident in the lower curves.

the final fusion correction.



(a) E9 Slice

(b) E9 Cut

Figure 10.6: Fusion Fraction for various (a) slices and (b) cuts on the E9 energy, for both the nominal and shaped stop distributions.

10.2.2 Determining ϵ

In order to obtain the true fusion fraction, ϵ , a version of the distribution with highly suppressed fusion migrations must be obtained. Migrations occur only on the fiducial volume boundary, and primarily in the Z direction. As such, one can suppress the migrations by depleting the stops around the fiducial volume boundary. No cuts can be made on the stop position, as this couples to fusion interference. However, as tracks must exceed three pads in length, fusion products from valid tracks cannot reach the first Z pad row. Cuts on the energy deposited on the first row can constrain the stop depth of the muon into the detector, to shape the stop distribution in Z, depleting events at the fiducial volume boundary. Similarly, cuts on the X and Y position of the pulse on the first pad row can be used to deplete boundary events in the X and Y directions. The first row X (FRX) versus the stop row in X can be seen in figure 10.7(a). The stop distribution in X for no cut, and two increasingly strict cuts can be seen in figure 10.7(b). This is an unnormalized stop distribution; requiring

the track to start on the two center X pads of the TPC reduces the stops on the fiducial volume boundary in X by a few orders of magnitude without reducing the statistics inside the fiducial volume. The same can be seen for the Y direction in figures 10.7(c) and 10.7(d). Figure 10.7(e) shows the energy of the pulse on the first Z row (FRE) versus the stop Z pad. In Z it is more difficult to suppress the fiducial boundaries without losing events; in the normalized distributions of figure 10.7(f), only 15% of the original statistics remain in the shaped distribution.

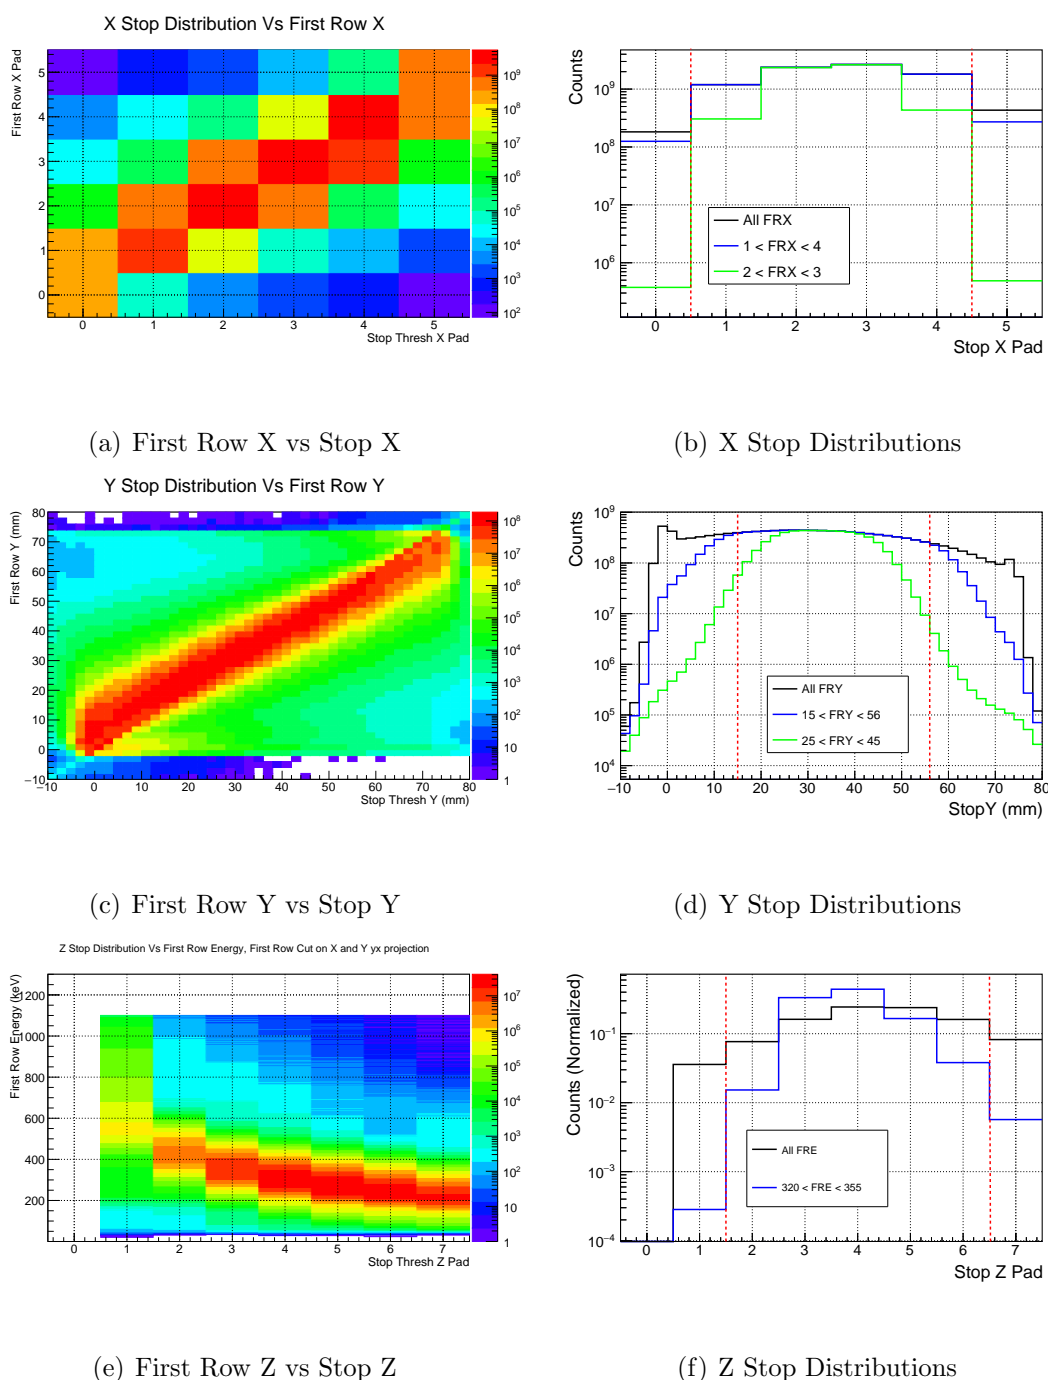


Figure 10.7: (a) First row X (FRX) vs reconstructed stop X using the stop threshold tracking algorithm and (b) stop X distributions for various cuts on the FRX. The same is shown in (c) and (d) for the Y dimension and (e) and (f) for the Z dimension. The fiducial volume boundaries are denoted in each case by the dashed red lines.

From the shaped stop distribution obtained using the strict version of all cuts shown in figure 10.7, the fusion fraction is determined in the same way as the observed fusion fraction,


$$\epsilon = \frac{F}{N} \quad (10.26)$$

where F is the number of events above the E9 cut of 3300 keV in the shaped distribution with the same background subtraction method applied for the deuteron recoils, and N is the total number of muons in the shaped distribution. For the full 2014 dataset, a true fusion fraction of 2.67×10^{-2} is obtained.


First Row Energy Scans to obtain uncertainty in ϵ

The uncertainty in the true fusion fraction, ϵ , is dominated by the remaining fusion migrations. With the stop threshold tracker, migrations only occur in the upstream direction. By selecting first row energy cuts to fix the population in bin 7 to 1% and slowly reduce the population in bin 2, we can quantify the sensitivity of our true epsilon to the remaining migrations in bin 2. The variety of stop distributions created to scan bin 2, as well as the corresponding values of the fusion fraction can be seen in figure 10.8. As the bin 2 population is decreased, the value of the fusion fraction stabilizes, indicating that migration losses out of bin 2 have been eliminated in the shaped distribution.

Similarly, we fix the population in bin 2 to 1% and reduce the population in bin 7 to quantify the sensitivity of the fusion fraction to the remaining migrations in bin 7. The set of distributions created via first row energy cuts to scan the bin 7 population can be seen in figure 10.9(a), while the corresponding fusion fractions obtained from each distribution can be seen in figure 10.9(b). Unlike the scan over bin 2, no saturation is observed in the fusion fraction as the bin 7 content is depleted. The cause of this unexpected behavior is unknown, but could be an artifact of unaccounted for losses within the TPC due to pad inefficiencies, threshold mis-matches, or calibration issues. For the purpose of this thesis, a conservative uncertainty estimate is made for the value of ϵ to account for this behavior.




(a) Stop Z Distributions




(b) Fusion Fraction Vs Bin 2 population

Figure 10.8: (a) A variety of stop distributions created with first row energy cuts to scan the fusion fraction as a function of the bin 2 population, as the bin 7 population is held constant. The fusion fraction determined from each stop distribution can be seen in (b).



(a) Stop Z Distributions



(b) Fusion Fraction Vs Bin 7 population

Figure 10.9: (a) A variety of stop distributions created with first row energy cuts to scan the fusion fraction as a function of the bin 7 population, as the bin 2 population is held constant. The fusion fraction determined from each stop distribution can be seen in (b).

10.2.3 Determining κ

Once the true and observed fusion fraction are obtained, the constant κ is needed to translate the difference between true and observed fusion fractions into a lifetime shift. It is necessary

to determine a set of conditions with variable fusion fractions for which the lifetime fit can be applied. This is obtained in two independent ways; using a scan over the stop Z position in the TPC and a scan of various first row energy cuts.


It was previously shown that the number of migrations in a bin is determined by the stop distribution, such that each Z row in the TPC contains a different fusion fraction [38]. The observed fusion fraction for each Z row can be seen in figure 10.10. By determining the true and observed fusion fractions as outlined above for each Z row, and fitting the lifetime for each accordingly, a linear fit to the fusion fraction ratio vs lifetime for each z row can be used to extract the constant of proportionality, κ . For the R2014 dataset, a value of $\kappa = 596 \pm 82.6$ is determined from the fit, as displayed in figure 10.10(c).

Alternatively, the stop distribution can be shaped to obtain various values of the fusion fraction, as shown in figure 10.11. The first row energy cut can be used to dramatically shift the stop distribution. The stop distribution and corresponding fusion fractions for a variety of first row cuts can be seen in 10.11(a) and 10.11(b). By fitting the lifetime to each of these distributions and fitting the lifetime versus fusion fraction, the constant of proportionality is determined. A value of $\kappa = 441 \pm 96$ is obtained, just within the allowed error to be considered consistent with the previous method. The value of $\kappa = 596 \pm 82.6$ is used in the final fusion correction, as the fit error is smaller.

10.3 Applying correction to Data


10.3.1 Nominal Stop Distribution

The stop distribution and observed fusion fraction for the sum and each sub-dataset for the 2014 data can be seen in figure 10.12. By comparing the observed fusion fraction for each dataset to the true fusion fraction and multiplying by the κ value obtained from the fitted fusion fraction versus z, a lifetime correction for each dataset is obtained. The lifetime for each dataset before (solid) and after (dashed) correction can be seen in figure 10.13.



(a) Z Stop Distribution

(b) Fusion Fraction vs Z Stop Row

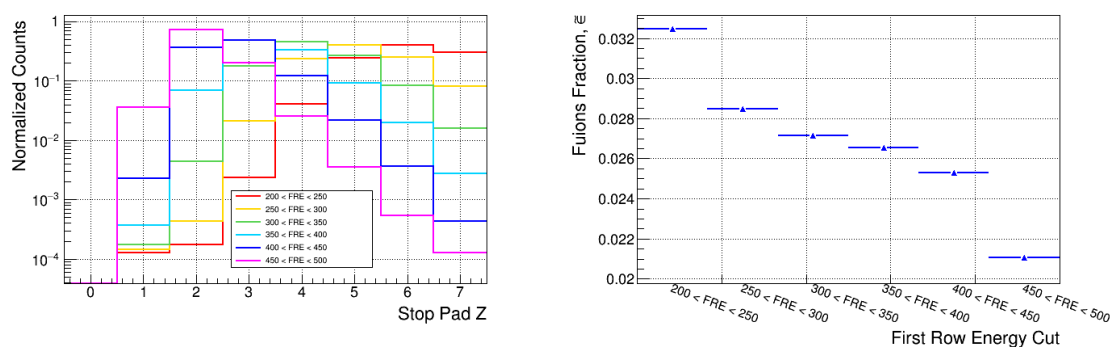


(c) Fitted Lifetime Vs Fusion Fraction

Figure 10.10: (a) The Z stop distribution for the summed R2014 dataset. (b) As fusion migration is a function of the shape of the stop distribution, each z row contains a different fusion fraction. (c) By scanning the z rows, a lifetime vs fusion fraction curve is created and used to determine the value of κ .


10.3.2 Balanced Stop Distribution

The nominal stop distribution is particularly unbalanced in z, leading to a large correction. It is therefore advantageous to use the first row energy cut to slightly shape the distribution in z, resulting in a more balanced distribution at the cost of 15% of the original statistics. The following cuts are applied to the to create this balanced distribution



(a) Stop Z Distributions

(b) Fusion Fraction for each stop distribution

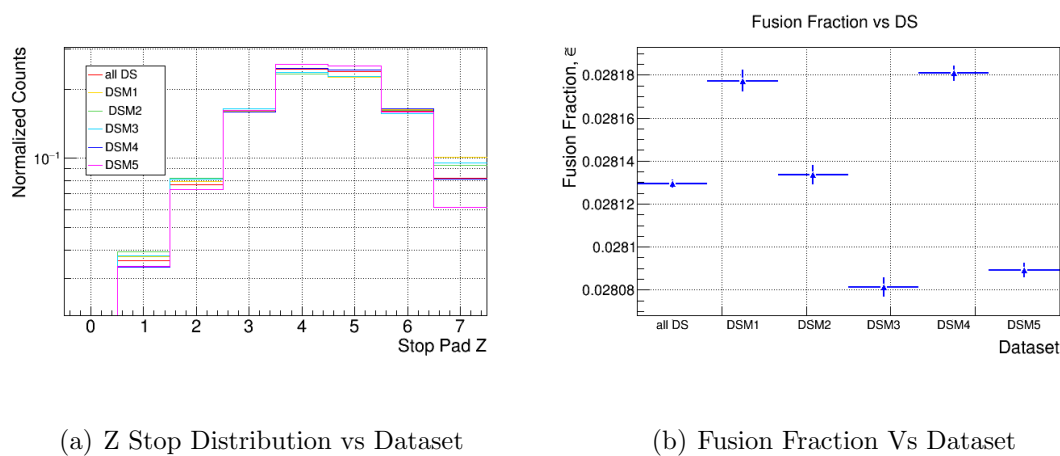


(c) Fitted lifetime Vs Fusion Fraction

Figure 10.11: (a) Several independent, skewed stop distributions produced by making cuts on the first row energy. (b) The fusion fraction of each of the stop distributions shown in (a). (c) The lifetime versus fusion fraction created from scanning the first row energy cuts is fitted to obtain the value of κ .

- $0 < \text{FR X} < 5$
- $15 < \text{FR Y} < 56$
- $245 < \text{FRE} < 1300$

The stop distribution and observed fusion fraction for the sum and each sub-dataset for the 2014 data can be seen in figure 10.14 for the balanced version of the distribution. The



(a) Z Stop Distribution vs Dataset

(b) Fusion Fraction Vs Dataset

Figure 10.12: (a) Stop Z distribution for each dataset and (b) observed fusion fraction for each dataset.


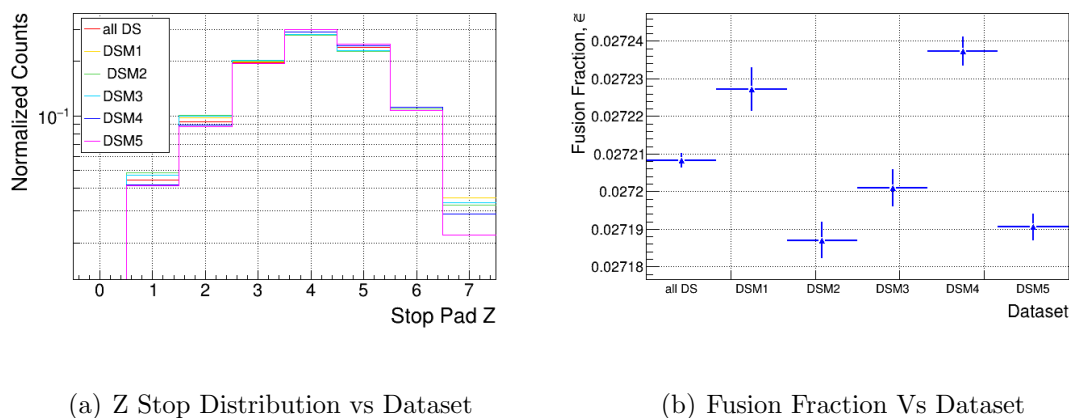


Figure 10.13: Nominal fitted rates (black) and fusion corrected (red) rates for a lifetime fit at 1 μ s with the nominal stop distribution.

lifetime corrections for the each dataset with the more balanced distributions can be seen in figure 10.15.



(a) Z Stop Distribution vs Dataset

(b) Fusion Fraction Vs Dataset

Figure 10.14: (a) Stop Z distribution for each dataset using the first row energy cut to create a more balanced z distribution, and (b) the observed fusion fraction for each balanced dataset.




Figure 10.15: Nominal fitted rates (black) and fusion corrected rates for a lifetime fit at 1us with the balanced stop distribution.

10.4 Uncertainty Estimates and Conclusions

As our lifetime shift is proportional to the difference between the true and observed fusion fractions,

$$\Delta\lambda = k \Delta\epsilon \quad (10.27)$$

where $\Delta\epsilon = 1 - \frac{\tilde{\epsilon}}{\epsilon}$. The total uncertainty in the lifetime correction is given by

$$\delta(\Delta\lambda) = k \delta(\Delta\epsilon) \oplus \delta(k) \Delta\epsilon. \quad (10.28)$$

The uncertainty in the constant, κ , relating the fusion fraction difference to the lifetime shift, is simply the uncertainty in the fit displayed in figure 10.10(c). A value of $\kappa = 596 \pm 82.6$ is obtained via scanning the z rows of the full R2014 dataset.

Both the true and observed fusion fraction determination contribute to the uncertainty in $\Delta\epsilon$, given by

$$\delta(\Delta\epsilon) = \frac{\partial(\Delta\epsilon)}{\partial\tilde{\epsilon}} \oplus \frac{\partial(\Delta\epsilon)}{\partial\epsilon} = \frac{\delta\tilde{\epsilon}}{\epsilon} \oplus \frac{\tilde{\epsilon}\delta\epsilon}{\epsilon^2}. \quad (10.29)$$

From the stability of the E9 cut in figure 10.6, an uncertainty of 2×10^{-3} is taken as the contribution in finding the observed fusion fraction. The uncertainty in $\Delta\epsilon$ is dominated by the linear dependence of the shaped fusion fraction on the remaining bin 7 population, quantified in figure 10.9. This effect is much smaller for a scan of the remaining population in bin 2, shown in figure 10.9. As this behavior is currently not understood, a conservative uncertainty estimate of 1.5% is assigned to $\delta\epsilon$. The fitted and corrected lifetimes for each dataset, as well as the uncertainty in the correction as determined by the procedure stated above, can be seen for both the nominal and balanced distributions in table 10.1.

10.5 Application of Fusion Fraction Method to Monte Carlo

One of the strong points of the technique illustrated above is its model independence; the method was developed with a completely data driven approach. It is however advantageous to apply the method to Monte Carlo data in order to validate its accuracy. Three Monte Carlo datasets are generated with GEANT4, with stop distributions akin to DS13, DS14 and DS16 of the R2014 dataset and a fusion fraction of exactly 3%. The GEANT4 simulation is then sent through a second stage to generate all detector responses and produce an output format identical to the data. This Monte Carlo generated data can then be processed through the same analysis chain as the data, and the truth information is retained for comparison.

	Nominal stop Distribution			Balanced Stop Distribution		
Dataset	R_{fit}	Correction	Uncertainty	R_{fit}	Correction	Uncertainty
DSM1	1055.36	34.17	10.57	1035.72	10.58	8.28
DSM2	1079.27	33.19	10.50	1101.65	9.79	8.25
DSM3	1068.12	32.02	10.41	1083.83	10.07	8.25
DSM4	1025.73	34.24	10.57	1039.20	10.79	8.29
DSM5	1046.98	32.19	10.42	1062.28	9.85	8.26
all DS	1051.15	33.09	10.49	1065.30	10.21	8.26

Table 10.1: fitted rates and corrections due to fusion interference for both the nominal and balanced stop distributions. All units are (1/s).

10.5.1 Determining $\tilde{\epsilon}$

The method of isolating the n-d scatter tail was applied to the Monte Carlo to check the validity of the method. Figure 10.16 shows the full E9 energy distribution, the reconstructed background found from looking at the E9 distribution from extra clusters with energy between 0 and 2600 keV, and the Monte Carlo truth distribution for the recoils. There is good agreement between the reconstructed recoil tail (blue) and the truth distribution (green) up to energies of 4200 keV.

10.5.2 Determining ϵ

The Monte Carlo stop distribution is shaped in the same manner as the data to suppress events at the fiducial volume boundaries and obtain the true fusion fraction. A fusion fraction of 0.03003 is obtained with the shaping procedure, indicating the value of epsilon is correctly obtained to the 10^{-3} level. In order to compare with the data, it is useful to repeat the





Figure 10.16: Comparison of the Monte Carlo full E9 distribution (black), reconstructed E9 background from neutron scatter events after applying the extra cluster energy cut to the Monte Carlo (blue), and the Monte Carlo truth distribution of scatter events (green).

scans which quantify the change in the fusion fraction as a function of the depletion of the bin contents on either end of the fiducial volume. The scan over the fraction of counts in bin 2 and as bin 7 is held constant can be seen for the simulated data in figure 10.17. The general trend of saturation near low bin 2 counts is consistent with that observed in the data, displayed in figure 10.8.


In contrast, a different trend is observed between the R2014 data and the simulated data when the bin 7 population is scanned while bin 2 is fixed. In the scan for the R2014, displayed in figure 10.9, the value of ϵ is still dramatically decreasing for small bin 7 fractions. In the simulated data of figure 10.18, this trend is less extreme, and appears to reach a finite value. This discrepancy between data and Monte Carlo is still not understood, but is perhaps a hint at non-uniform detector effects which are not present in the simulated data.

10.5.3 Determining κ

The constant of proportionality relating the difference between the true and observed fusion fraction to the lifetime shift is again obtained in two ways; using the first row cuts to




(a) Stop Z Distributions




(b) Fusion Fraction Vs Bin 2 population

Figure 10.17: (a) A variety of stop distributions created with first row energy cuts to scan the fusion fraction as a function of the bin 2 population, as the bin 7 population is held constant. The fusion fraction determined from each stop distribution can be seen in (b).



(a) Stop Z Distributions




(b) Fusion Fraction Vs Bin 7 population

Figure 10.18: (a) A variety of stop distributions created with first row energy cuts to scan the fusion fraction as a function of the bin 7 population, as the bin 2 population is held constant. The fusion fraction determined from each stop distribution can be seen in (b).


create several distributions which scan various values of ϵ and using the lifetime in each z row, which relies on the properties of the stop distribution to produce variations in the

observed fusion fraction. The values of ϵ and fitted rates obtained from each Z row as well as the fit used to extract κ can be seen in figure 10.19.



(a) Z Stop Distribution

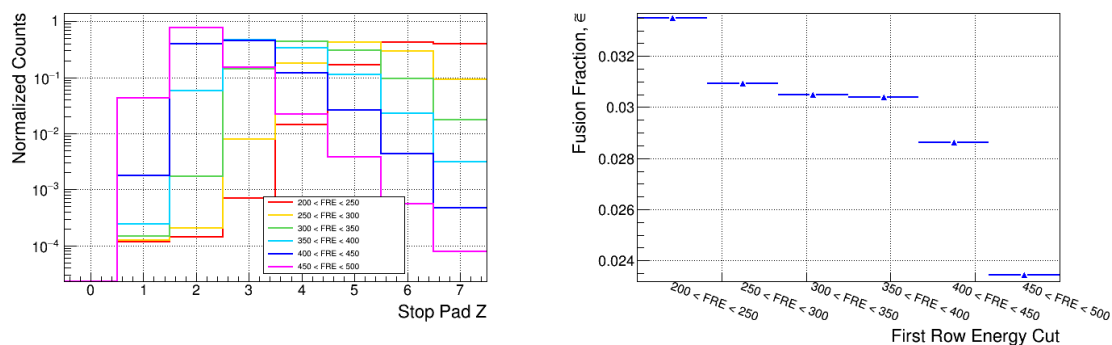
(b) Fusion Fraction vs Z Stop Row



(c) Fitted Lifetime Vs Fusion Fraction


Figure 10.19: (a) The Z stop distribution for the Monte Carlo data. (b) As fusion migration is a function of the shape of the stop distribution, each z row contains a different fusion fraction. (c) By scanning the z rows, a lifetime vs fusion fraction curve is created and used to determine the value of κ .

Similarly, a range of stop distributions producing a variety of fusion fractions and the lifetime associated with each can be seen in figure 10.20.



(a) Stop Z Distributions

(b) Fusion Fraction for each stop distribution



(c) Fitted lifetime Vs Fusion Fraction

Figure 10.20: κ determination procedure applied to the Monte Carlo data. (a) Several skewed stop distributions produced by making cuts on the first row energy. (b) The fusion fraction of each of the stop distributions shown in (a). (c) The lifetime versus fusion fraction created from scanning the first row energy cuts is fitted to obtain the value of κ .

10.5.4 Monte Carlo Correction

A comparison of the fitted and corrected rate, as well as the MC truth rate obtained from a lifetime histogram created using the MC true stop position in the TPC, can be seen in figure 10.21. The uncorrected rate from the lifetime fit is 20 Hz lower than the rate obtained when using the MC true stop position. The fusion correction brings the rate within 8 Hz of the

truth value. It is unclear whether the remaining discrepancy is due to electron interference. A critical next step will be running the MC histograms in the absence of electron interference in order to disentangle the effects.




Figure 10.21: Comparison of the fitted rate, fusion corrected corrected rate, and the rate obtained by fitting the Monte Carlo when the TPC truth information is used to determine the stop position.

Chapter 11

CONSISTENCY CHECKS

As a high degree of precision is required for the MuSun measurement, rigorous consistency checks must be performed in order to expose any unforeseen systematic sources of error. The following chapter illustrates several important scans to determine the stability of the fitted decay rate with variations on the fit range, muons stop position, and decay electron position within the scintillators.

A set of standard cuts and conditions, justified in previous chapters, are applied to each histogram. In order to reduce the effects of accidentals and pileup muons, only events with a single electron are used to produce the final lifetime histograms. The effect of this requirement is studied in detail in chapter 9. In chapter 10, several cuts are applied to the stop position and energy deposited on the first row of the TPC in order to produce a more peaked stop distribution and mitigate effects of muon catalyzed fusion interference. This reduces the fusion correction by a factor of three, at a cost of 15% of the statistics. Unless otherwise noted, all lifetime fits are a standard three parameter fit using the range $[1,24] \mu s$.

11.1 Dataset Consistency

The R2014 data is split into five sub-datasets, characterizing five distinct beam periods, as discussed in section 6.2. The first two datasets, DSM1 and DSM2, were taken with the misplaced beam collimator, leading to an increased beam electron background. In DSM2, the MuSCa was also mis-aligned. DSM3 was taken with the beam collimator at the bending magnet completely removed, and the final two datasets were taken with the upgraded collimator. As an optimum beam tune was difficult to achieve with the mis-placed slit, the

parameters of upstream beam elements were frequently changed throughout the data collection period, especially in the first three datasets. It is crucial to check the consistency of the sub-datasets to study the possibility of differences in the observed disappearance rate correlated to changes in the beam conditions.

The fitted rate and background parameters can be seen in figure 11.1 for each of the five large datasets. The fitted background parameter has been normalized to the number of muons, and quantifies the change in the accidental rate before and after the collimator was properly aligned. A linear fit to the five datasets, displayed in red, yields a fit probability of 0.28 and chi squared indicating good consistency within the five datasets.




Figure 11.1: Fitted rate and background parameter for each of the five large datasets. The background parameter, B , has been normalized to N , the number of muons determined from the fit. A linear fit to the five datasets, displayed in red, yields a reasonable probability of 0.28 and a good chi squared indicating consistency between the datasets.

11.2 Fit Variations

A standard range must be selected for the lifetime fits which makes physical sense and does not bias the result. Ideally, it is advantageous to choose the largest possible range to increase the statistics included in the fit. However, due to the 800ns step in the accidental rate from the MORE signal, the fit can only start after $1 \mu s$, and as the MORE pileup window is finite, the fit must end before $25 \mu s$. A scan of the fit parameters as a function of the start and stop time of the fit can confirm that the choice of the fit range does not bias the results.

11.2.1 Lifetime Vs Fit Start Time

As several systematic effects are more significant at early times, such as wall stops in high Z materials and muon catalyzed fusion interference, those effects would be enhanced with early fit start times. In order to test the possibility of such effects and justify using the earliest possible start time to retain statistics, a scan for a variety of fit start times is made and the consistency of the fitted rate between early and late start times is established.


As each fit uses a subset of the same data, the fit results are highly correlated, and the statistical error bars do not describe their correlation. A separate figure of merit must be derived for determining whether the scan points are statistically consistent. Assuming two lifetime fits obtaining rates R_1 and R_2 , where the second start time is later, such that N_2 is a subset of N_1 , the statistical uncertainty between the two rates can be written

$$\delta(\Delta R) = \sqrt{\left(\frac{\partial(\Delta R)}{\partial R_1}\right)^2 \delta_1^2 + \left(\frac{\partial(\Delta R)}{\partial R_2}\right)^2 \delta_2^2 + 2\frac{\partial(\Delta R)}{\partial R_1} \frac{\partial(\Delta R)}{\partial R_2} \text{cov}(R_1, R_2)} \quad (11.1)$$

where δ_1 and δ_2 are the statistical uncertainties given by $\frac{\lambda}{\sqrt{N_1}}$ and $\frac{\lambda}{\sqrt{N_2}}$, respectively, and $\Delta R = R_1 - R_2$. In this case, where the second fit is a subset of the first, $\text{cov}(R_1, R_2) = \delta_1^2$ such that the statistical error on the difference in fitted rates is given by

$$\delta(\Delta R) = \sqrt{\delta_1^2 + \delta_2^2 - 2\delta_1^2} = \sqrt{\delta_2^2 - \delta_1^2} \quad (11.2)$$

The start time scan for the full R2014 data can be seen in figure 11.2. The stop time time of each fit is fixed at $24\mu s$. The red shaded area represents the allowed statistical deviation of each point from the first fit at $1\mu s$, using equation 11.2. For all datasets and the sum, consistent results are obtained for each fit start time out to $10\mu s$.



(a) Full Scan

(b) First $5\mu s$

Figure 11.2: (a) Fitted rate for a range of fit start times. The shaded red area denotes the allowed statistical 1σ deviation from the full fit starting at $1\mu s$. All fits are stopped at $24\mu s$. (b) The first $5\mu s$ of the same scan, zoomed to detail the early start time results.

11.2.2 Lifetime Vs Fit Stop Time

Accidental background dominates at later times, when the signal to noise ratio is poor. It is also of interest to scan the stop time of the fit while keeping the start fixed at $1\mu s$. Dependence of the fitted rate on the stop times could indicate a deviation from the assumed flat accidental shape. Once again, the bands showing the allowed statistical deviation compared to the full fit stopping at $24\mu s$ are displayed in red.




Figure 11.3: Fitted rate for a range of fit stop times. The shaded red area denotes the allowed statistical deviation from the full fit stopping at $24\mu s$. All fits start at $1\mu s$.

11.3 Lifetime vs TPC Stop Position

As an additional consistency check, it is advantageous to determine the observed disappearance rate as a function of the muon stop position within the TPC. It is expected that muon stops near the boundaries of the TPC are more likely to capture on wall materials, leading to a distortion of the decay rate. As such, scans over the stop position serve as a measure to check the validity of the fiducial volume cuts. Furthermore, as fusion interference is an artifact of the fiducial boundaries, scans over the stop position can aid in quantifying the effects of this important systematic. Although work is still ongoing to separate these two systematic sources of error, this study can set the scale of their combined effect.

11.3.1 Lifetime vs Stop X

The fit parameters obtained for a scan over the X stop position for the full R2014 data can be seen in figure 11.4. Fiducial volume cuts have been applied in the Y and Z dimensions in

order to isolate effects in the X dimension. As expected, the fitted rate increases dramatically on X rows 0 and 5, where punch through muons stop in wall materials. The pad rows 1-3 demonstrate stability in the fitted rate, and are consistent with the fitted rate when all fiducial volume cuts are applied. There is a slightly increased rate in X row 4. As the

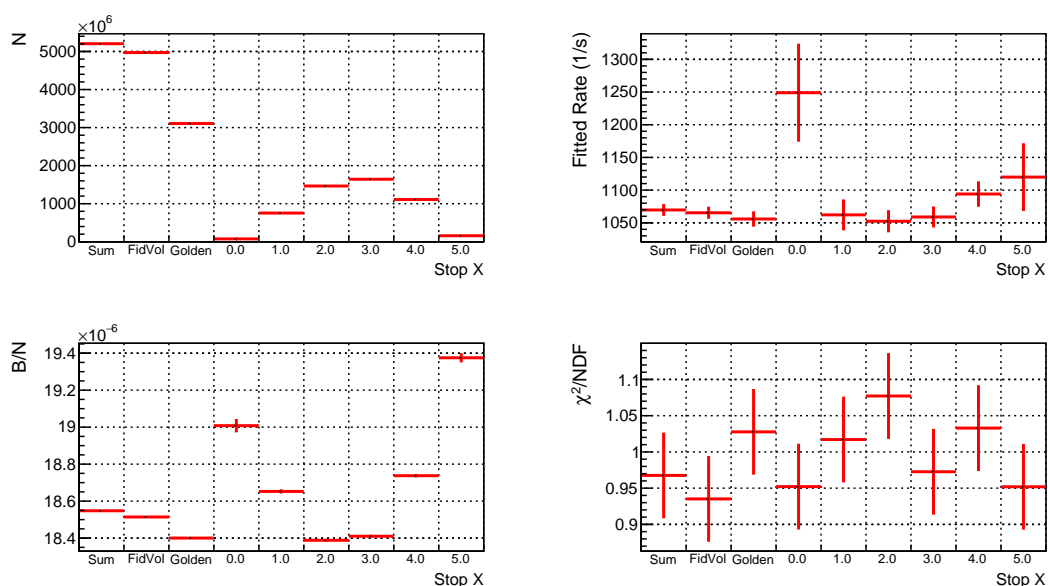



Figure 11.4: Lifetime fit parameters versus the muon stop X row in the TPC. Fiducial volume cuts have been applied in the Y and Z dimensions.

stop distribution may change with the beam conditions, it is of interest to compare the stop distributions in X as well as the fitted rate versus X stop position for each dataset. The stop distribution in X for each dataset, as well as the stop X lifetime scan can be seen in figure 11.5. The stop distribution shown is normalized to the total number of counts, and remains fairly consistent throughout the R2014 data collection period. The anomalous rate observed in X row 4 is primarily from DSM2. This should be investigated in more detail. In any case, the fiducial volume choice of excluding events which stop in X bins 0 and 5 are sufficient to eliminate large effects from stops in wall materials. Furthermore, as no rate change is seen

across the interior X bins, it can be concluded that muon-catalyzed fusion interference with the X stop position is sufficiently small.



(a) X Stop Distribution

(b) Fitted Rate Scan Over Stop X

Figure 11.5: (a) The similarity in the stop distribution in X for each dataset. (b) The fitted rate obtained for each X bin for each datasets. Fiducial volume cuts have been made in Y and Z. The golden cut is a more strict variation of the fiducial volume cut, which includes only the inner most pads of the TPC.

11.3.2 Lifetime vs Stop Y

As the Y position is determined from the charge arrival time, there are no discrete bins for the Y stop location. In order to retain adequate statistics to observe deviations in the observed disappearance rate, the Y stop position is recorded in 10mm bins for the purpose of scanning the fitted rate. The fit parameters for a scan over the Y stop position for the full R2014 data are displayed in figure 11.6. Once again, the fitted rate is higher at extreme values of Y due to stops in the silver cathode and pad plane. Otherwise, the observed disappearance rate is consistent for all values of Y within the fiducial volume. Figure 11.9 shows the normalized Y stop distribution, as well as the scan of the fitted rate for each dataset. The standard fiducial volume cut includes all Y stops in the range 15-56mm. As such, the first and last

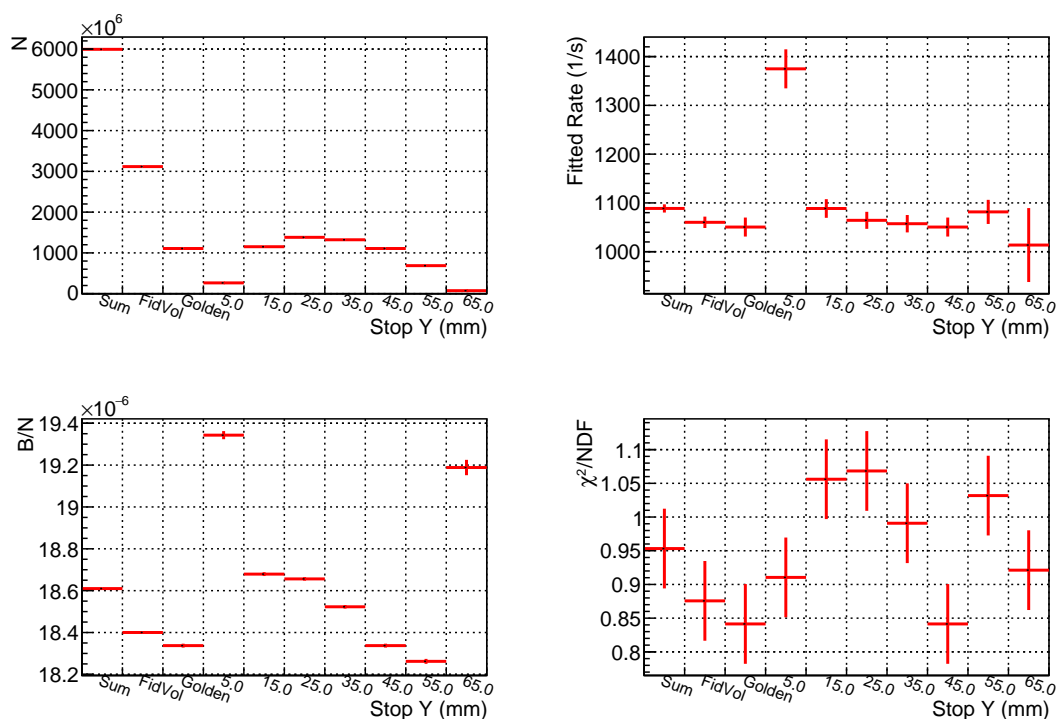
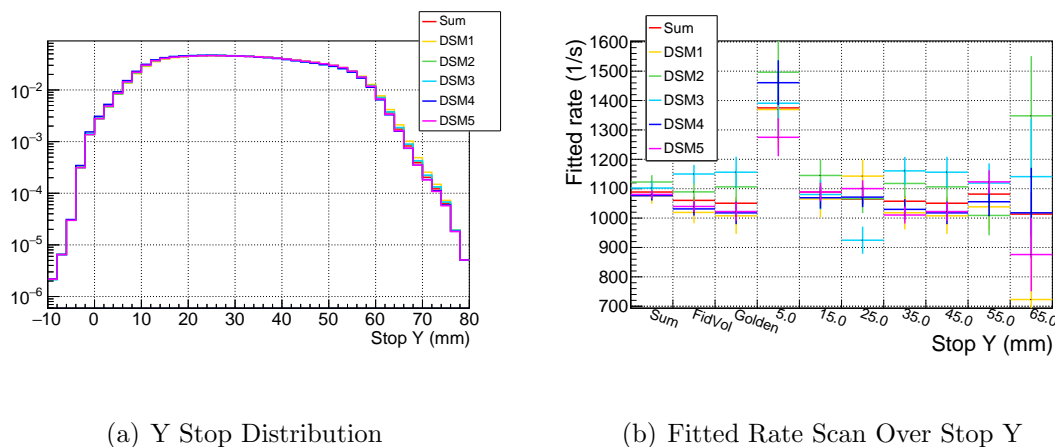


Figure 11.6: Lifetime fit parameters versus muon stop Y in the TPC. Fiducial volume cuts have been applied in the X and Z dimensions.

bins of the scan shown are excluded, once again confirming the choice of fiducial volume, as all interior bins yield consistent decay rates. A slight downward trend in rate may indicate a small fusion interference effect in the Y direction.

11.3.3 Lifetime vs Stop Z

In the Z direction, a significant number of stops in the final Z bin result from punch through muons, which likely capture on wall materials. Additionally, the largest contribution from muon-catalyzed fusion interference is expected in the Z direction, as the proton readily changes the Z stop position in all three tracking algorithms. The size of this effect is evident



(a) Y Stop Distribution

(b) Fitted Rate Scan Over Stop Y

Figure 11.7: (a) The stop distribution in Y for each dataset. (b) The fitted rate obtained for each Y bin for each datasets. Fiducial volume cuts have been made in X and Z.

in figure 11.8. As migration only occurs in the upstream direction from the threshold tracker, and depends on the population in neighboring bins, the net migration into each bin increases for more upstream bins, leading to a reduction in the observed decay rate.

The Z direction appears to be most sensitive to changes in beam tune, as evidenced by the comparison of the stop distributions in each dataset seen in figure 11.9(a). This leads to different slopes in the change of rate versus Z bin displayed in 11.9(b). A method for correcting the fusion interference contribution to the disappearance rate is detailed in chapter 10, which removes the Z dependence based on the observed number of fusion products. A comparison of the Z dependence of the fitted rate for the sum of all R2014 datasets before and after this correction is shown in figure 11.10.

11.4 Lifetime vs Electron Studies

In the interest of exploring potential systematic effects due to electron background or detector inefficiencies, it is useful to scan the fitted rate as a function of the electron position in the electron scintillator hodoscope.

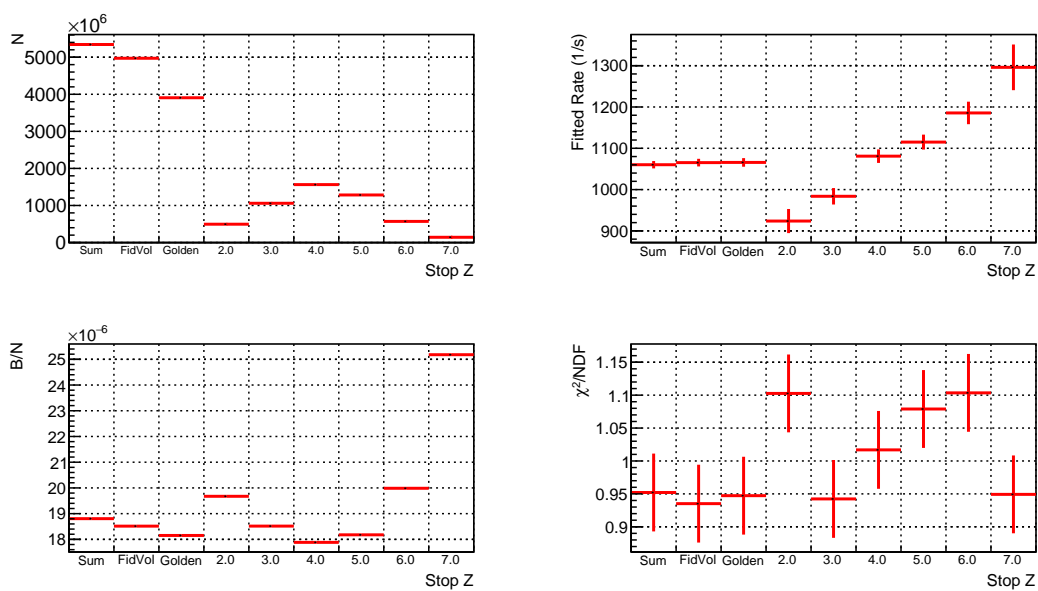
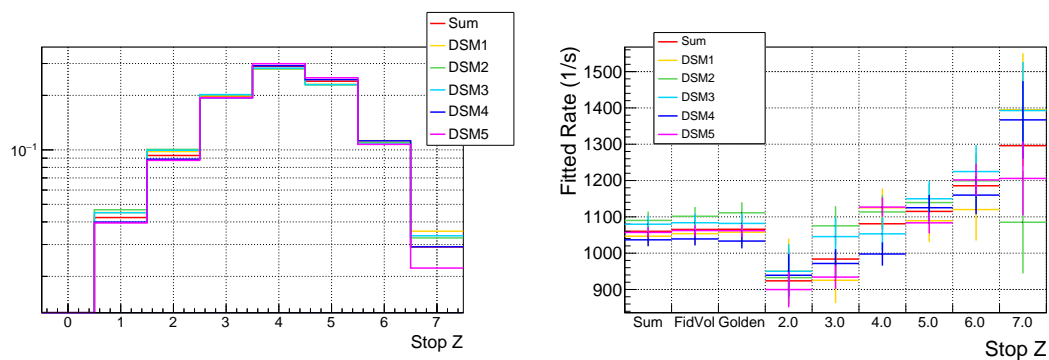


Figure 11.8: Lifetime fit parameters versus the muon stop Z row in the TPC. Fiducial volume cuts have been applied in the X and Y dimensions.



(a) Z Stop Distribution

(b) Fitted Rate Scan Over Stop Z

Figure 11.9: (a) The stop distribution in Z for each dataset. (b) The fitted rate obtained for each Z bin for each datasets. Fiducial volume cuts have been made in X and Y.




Figure 11.10: Dependence of the fitted rate on the muon stop Z position before (red) and after (black) the fusion migration correction described in chapter 10 is applied.

11.4.1 Lifetime vs eSC Segment

The azimuthal angle of the electron arrival in the scintillators can be explored via a scan of the fitted rate versus the eSC segment, as displayed in figure 11.11. The sum of all segments, as well as four combinations of distinct detector topologies can be seen in the first five bins of the scan. A fit of the rate obtained from each of the 16 individual eSC segments is also shown in the scan, and demonstrates a good chi squared, indicating a consistent rate is obtained over the azimuth.

11.4.2 Lifetime vs electron Z Position

The Z position of the electron in an eSC segment is determined by a time of flight method, described in detail in section 7.4.2. A scan of the lifetime fit parameters versus electron Z position is displayed in figure 11.12. As the physical scintillator is only 900 mm in length, any Z position less than -450 mm or greater than 450 mm is due to a mis-reconstruction from the time of flight method. Only a small number of events are reconstructed outside of the physical range. The statistics are too low to obtain an adequate fit for events with Z position greater




Figure 11.11: Fitted rate for each eSC segment. The first bin shows the fitted rate for the sum of all segments, while the subsequent four bins show the sum of the top (1-4 and 13-16), bottom (5-12), left (1-8), and right (9-16) segments. A linear fit through the 16 individual segments demonstrates good chi squared and reasonably high fit probability, indicating consistency in the fitted rates.

than 450 mm, and the fit to events below the -450 mm range show reasonable consistency, such that no systematic effect is believed to manifest from this mis-reconstruction.

The fitted values of N and B/N suggest that the TPC is centered within the scintillator hodoscope, with a beam electron background distributed more upstream. A fit of the observed rates to a constant, shown in figure 11.13, retains a good chi square and high fit probability, indicating consistent values of the decay rate for all values of Z .

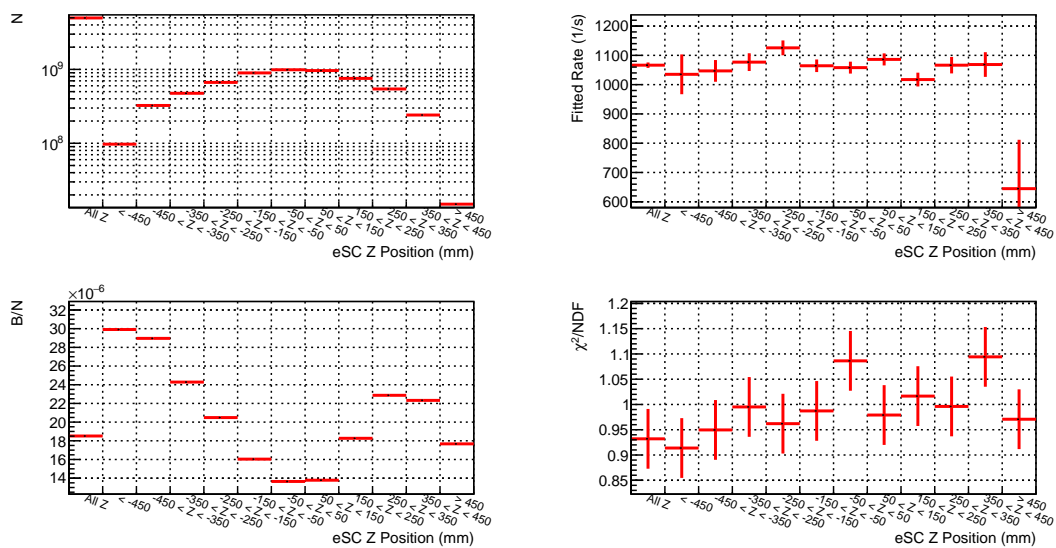


Figure 11.12: Fit parameters versus electron Z position in the scintillator.




Figure 11.13: Fitted rate versus Z position of the electron within the eSC segment.

Chapter 12

CONCLUSIONS

This chapter outlines the key conclusions of the dataset consistency, systematic assessment, and outlook for obtaining a final result with the R2014 dataset.

12.1 Quality Cuts and Data Consistency

For the final analysis, a set of standard quality cuts was developed. In order to suppress background effects by a factor of 80 or more, only events with a single observed electron within a $25\ \mu\text{s}$ window after the muon entrance are accepted. The implications of this selection criteria are discussed in detail in section 9.2. To curb the effects of muon-catalyzed fusion interference, cuts are made on the first Z row of the TPC, in particular on the energy deposition, to produce a more balanced muon stop distribution. Several important consistency checks are presented in chapter 11. The fit to the observed decay rates of all independent sub-datasets reveals good consistency, indicating a stable rate throughout the data collection period. This stability demonstrates the robustness of the analysis method, as the first half of the run was taken with beam conditions far from optimal, which were partially repaired during the run. Stable start and stop time scans imply no time distortions from systematic effects which have not already been addressed. Scans of the fitted rate versus muon stop position display consistent values within the fiducial volume in the X and Y direction. The Z dependence of the rate confirms the assumption that muon-catalyzed fusion interference predominantly affects only the Z dimension, and is nearly eliminated with the application of the fusion correction. Although a pronounced directional dependence of the electron background is established in section 9.2.6, scans of the fitted rate versus electron

position in the eSC indicate no effect on the fitted rate. Overall, analysis of the muon decay with an electron observed in the eSC shows good consistency for the R2014 dataset.

12.2 Systematics

The main systematic effects and their impact on the observed disappearance rate can be seen in table 12.1. For the majority of the effects, an upper bound has been determined, and further quantification and uncertainty estimates are needed in future analyses. The precision goal of the MuSun experiment is $4 s^{-1}$ statistical plus $4 s^{-1}$ systematic uncertainty. For the R2014 dataset, a statistical uncertainty of $9.23 s^{-1}$ is obtained from the fitted rate starting at $1 \mu s$, with all standard cuts applied.

Systematic		Rate Shift (s^{-1})	Shift Direction
MORE BG	Missed Muons	< 1.54	+
	Upstream Muon Stops	< 0.25	+
	2e Deviation	< 7.5	-
Muon Track Interference	Electrons	< 5	-
	Fusion Products	10.2 ± 8	-
Gas Purity	Isotopic	0.8 ± 0.2	+
	Chemical	< 6	+
Other	Wall Stops	small	+

Table 12.1: MuSun sources of systematic error and the effect of each on the observed rate. For the fusion correction and isotopic purity, an uncertainty in the observed shift is given. For all other systematics, only an upper limit on the size of the effect has been determined. The last column denotes the direction of the shift; positive indicates an increase in the fitted rate.

The systematics associated with the MORE induced background are investigated in chap-

ter 9. Muons which are missed by the entrance detectors may stop in materials surrounding the target, resulting in a distortion to the observed rate. The effect on the rate is determined via a zero extrapolation detailed in section 9.1.2. Muons which stop upstream of the entrance detectors, either in the area or the copper collimator, may result in a similar distortion. A background model is developed in section 9.3 to quantify the effect of upstream muon stops on the observed rate. Upstream stops in materials composing the beam pipe, such as iron, are expected to exhibit a fast component in the decay distribution after the 800 ns kicker step, and stops in the collimator may contribute to a rising background. Neither of these time dependencies are observed in the data, constraining the upstream stop rates to small values. The shift associated with these effects in the table is determined from a background model using these constrained rates. In the final analysis, only events with a single observed electron within a $25 \mu s$ window after the muon entrance are considered in order to suppress background contributions. The events with more than one electron, which enhance the background contribution significantly, are analyzed independently, and found to result in a rate $600 s^{-1}$ lower than the single electron distribution. This effect is still not understood, but will be suppressed in the single electron distribution by at least a factor of 80, based on the electron multiplicity formalism developed in 9.2.2, indicating an upper bound of $7.5 s^{-1}$.

There are two main sources which lead to mis-reconstructions of the muon stop position within the TPC – decay electrons and muon-catalyzed fusion products. The combination of fiducial volume cuts and time dependencies of each of these sources results in a distortion of the observed time spectrum. A first look at electron interference is offered in section 8.2.2. In comparison with the data, it is evident that the mechanism of electron interference is well modeled by the MC except for the extended charge arrival response of the TPC, which results in a much larger effect. This allows for the determination of a $5 s^{-1}$ upper bound on the rate shift in the data, which is expected to become smaller as the MC response is developed to more realistically characterize the data. An in depth study of muon-catalyzed

fusion interference is presented in chapter 10. A correction technique is developed which uses energy and position information of the first TPC Z row to create a shaped stop distribution in which fusion inference effects are highly suppressed. By comparing the fraction of fusion events in the nominal stop distribution to those in the shaped, migration-less distribution, a correction to the observed rate is assigned. This technique works well on the MC, but an inflated uncertainty is assigned due to the not yet understood mismatch in behavior between MC and data when using upstream stop information to deplete stops in pad row 7.

The ejection of 5.4 MeV muons following the fusion of $pd\mu$ molecules, present if trace amounts of protium exist within the deuterium target, can lead to lifetime distortions if these long range muons capture in materials outside the target. The advanced techniques used to create and maintain an isotopically pure deuterium target mitigate effects from the $pd\mu$ fusion, (cf. 8.3.1). The rate shift and uncertainties are determined by fitting the solutions to the kinetic equations, presented in appendix C, including the protium concentrations measured by the chromatography throughout the 2014 collection period. Muons will readily transfer to any impurities within the gas, the most concerning being nitrogen. As the capture rate scales with Z^4 , the presence of such impurities can lead to significant distortions of the decay rate. The effect was calibrated with N_2 doped data in a previous collection period, which determined that 1 ppb of nitrogen impurity leads to a 3 Hz shift in the observed decay rate. Two techniques are used to determine the concentration of nitrogen impurity within the R2014 target gas – a sensitive chromatography measurement and an in-situ capture recoil analysis. Chromatography measurements throughout the data collection period suggest nitrogen concentrations of 1 ppb. However, as the system was not rigorously calibrated until after the R2014 data run, a second method to confirm this result is of great importance. The capture recoil analysis, which directly observes recoil energies from nitrogen captures within the TPC, implies 2-3 ppb of N_2 . This is likely inflated due to a background of neutron-deuterium scatters which deposit a similar energy. However, this first result is taken as an

upper limit on the chemical purity systematic.

12.3 Outlook

Overall, the analysis of the R2014 is nearly complete. Scans of the fitted rate under several spatial and temporal conditions reveal excellent consistency, and an initial upper limit has been obtained for nearly all sources of systematic error. The remaining tasks focus on further constraining the lifetime shifts and determining the uncertainty for each of the systematics in table 12.1.

All systematic sources of error which contribute to the MuSun result have been discussed except one – muon stops in wall materials. An analysis which uses the neutron detectors to quantify the effects of muon stops in the materials surrounding the detector, and validate the efficiency of the fiducial volume cuts, has been developed by UW Ph.D. student Ethan Muldoon using the R2015 data. Once completed, this analysis will be applied to the R2014 dataset to constrain the remaining systematic error.

The capture recoil analysis which was applied to determine the nitrogen concentration in the target gas has an additional background from the deuteron recoil following a neutron scatter. Based on MC simulations, this is expected to artificially inflate the determined concentration by at least 50%. This background can be quantified by selecting isolated events in the expected energy range which are spatially separated from the muon track, as nitrogen capture events occur on the stop pad. This analysis is similar to the n-d background subtraction technique developed in this thesis for tagging p-t fusion events (cf. sec 10.2.1) and will be implemented for the R2014 data in the near future.

The electron interference systematic can be further reduced in two ways. First, the MC simulations must be better tuned to reproduce the response of the TPC to extended charge arrival. Second, an in depth analysis of the positive muon decay data is still in progress. Once completed, scans over the eSC segments can reveal electron interference effects in the

absence of muon-catalyzed fusion effects. Analysis of the positive muon decay data will also serve as an important benchmark, as this can be compared to the precisely determined MuLan result.

Perhaps the most difficult systematic to constrain in the MuSun analysis arises from the muon-catalyzed fusion. As this is the largest effect with the greatest uncertainty, effort should be made to better understand the short-comings of the current correction procedure. In addition, a complementary analysis which uses the ${}^3\text{He}$ position to correctly determine the muon stop position is currently under development, and will be applied to the R2014 dataset once finished.

Finally, more work is needed to understand the significant deviation of the fitted rates obtained with different electron multiplicities. The addition of the electron wire chambers, which were mostly excluded in this thesis, could help tremendously in further constraining electron backgrounds. However, the instability of the start time scans and impact parameter cuts when including the ePCs must first be addressed.

12.4 Summary

MuSun is a muon lifetime experiment based on an active, ultra-pure deuterium target operating as a cryogenic TPC. The author of this thesis contributed to upgrades of this essential and novel instrument, assumed a leading role during data collection, and has completed an advanced analysis of the first 40 TB dataset.

The upgrades and commissioning of the TPC proved successful, as it worked flawlessly with stable voltages throughout the data collection period. The TPC provides a wealth of data including muon stops and muon-catalyzed fusion products, which need to be carefully disentangled. The factor of three improvement in energy resolution yields sensitivity to subtle effects, including the nuclear recoil of muon captures on impurities which can be used to determine the chemical purity of the target gas (cf. sec 8.3.2). This also allows for a

background subtraction of n-d scatter events and efficient tagging of proton-triton fusions (cf. sec 10.2.1).

A total of 7.4×10^9 negatively charged muon decay events were collected using the upgraded TPC during the R2014 data collection period. This includes all selection cuts, except for the TPC first row energy cuts used to center the stop distribution, which reduce the total statistics to 6.3×10^9 . Several measures were taken to ensure data quality, including an in depth study of the beam profile in which it was determined that an upstream collimator was misplaced. This was remedied for the second half of the collection period, but the mis-tuning of a quadrupole magnet (cf. appendix D) remained, leading to significant challenges in the data analysis.

Tools were developed to display trends and make quality selection cuts. Several upgrades to the analysis software were implemented, including an improved muon entrance definition, a new electron analysis, and muon stop analyses based on careful studies of the TPC response. Data consistency has been verified and nearly all sources of systematic error have been addressed. With only a few analyses remaining, the R2014 dataset is nearing completion, and unblinding is expected this year!

BIBLIOGRAPHY

- [1] 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 Mhlbauer, 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.
- [2] J. Adam, M. Tater, E. Truhlck, 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):93 – 100, 2012.
- [3] 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. TM. 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. VanDeven-der, 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.
- [4] I. Alekseev, Ev. Arkhipov, S. Bondarenko, O. Fedorchenko, V. Ganzha, K. Ivshin,

- P. Kammel, make 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, December 2015.
- [5] 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.
- [6] 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.
- [7] D Balin, V A. Ganzha, S M. Kozlov, Evgeny Maev, G E. Petrov, M A. Soroka, G N. Schapkin, G G. Semenchuk, V A. Trofimov, A A. Vasiliev, A A. Vorobyov, N I. Voropaev, Claude Petitjean, B Gartner, B Lauss, Johann Marton, Johann Zmeskal, T Case, K M. Crowe, and M P. Faifman. High precision study of muon catalyzed fusion in d2 and hd gas. *Physics of Particles and Nuclei*, 42:185–214, 03 2011.
- [8] G. Bardin, J. Duclos, A. Magnon, J. Martino, A. Richter, E. Zavattini, A. Bertin, M. Piccinini, A. Vitale, and David F. Measday. A novel measurement of the muon capture rate in liquid hydrogen by the lifetime technique. *Nucl. Phys.*, A352:365–378, 1981.
- [9] 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.
- [10] 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, Oct 2004.
- [11] A. Baroni, R. Schiavilla, L. E. Marcucci, L. Girlanda, A. Kievsky, A. Lovato, S. Pastore, M. Piarulli, Steven C. Pieper, M. Viviani, and R. B. Wiringa. Local chiral interactions, the tritium gamow-teller matrix element, and the three-nucleon contact term. *Phys. Rev. C*, 98:044003, Oct 2018.

- [12] V. Bernard, N. Kaiser, and Ulf-G. Meissner. Qcd accurately predicts the induced pseudoscalar coupling constant. *Phys. Rev. D*, 50:6899–6901, Dec 1994.
- [13] A. Bertin, A. Vitale, A. Placci, and E. Zavattini. Muon capture in gaseous deuterium. *Phys. Rev. D*, 8:3774–3793, Dec 1973.
- [14] 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.
- [15] O. Bunemann, T. E. Cranshaw, and J. A. Harvey. Design of grid ionization chambers. *Canadian Journal of Research*, 27a(5):191–206, 1949.
- [16] Paul Bttiker and Ulf-G Meissner. Pionnucleon scattering inside the mandelstam triangle. *Nuclear Physics A*, 668:97–112, 03 2000.
- [17] 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. *United States: World Scientific Pub Co.*, 1989.
- [18] Carl Carlson. The proton radius puzzle. *Progress in Particle and Nuclear Physics*, 82, 02 2015.
- [19] Chia C. Chang, Enrico Rinaldi, A. N. Nicholson, E. Berkowitz, N. Garron, D. A. Brantley, H. Monge-Camacho, C. Monahan, C. Bouchard, M. A. Clark, B. Joo, T. Kurth, K. Orginos, P. Vranas, and A. Walker-Loud. A percent-level determination of the nucleon axial coupling from quantum chromodynamics. *Nature (London)*, 558, 5 2018.
- [20] Eugene D. Commins and Philip H. Bucksbaum. *Weak Interactions of Leptons and Quarks*. Cambridge University Pres, 1983.
- [21] 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.
- [22] Arnaud Dupays. Isotopic effects in the muon transfer from $p\mu$ and $d\mu$ to heavier atoms. *Phys. Rev. Lett.*, 93:043401, Jul 2004.
- [23] Nadia Fettes, Ulf-G. Meiner, and Sven Steininger. Pion-nucleon scattering in chiral perturbation theory (i): Isospin-symmetric case. *Nuclear Physics A*, 640(2):199 – 234, 1998.

- [24] 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.
- [25] 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.
- [26] Doron Gazit, Sofia Quaglioni, and Petr Navrátil. Three-nucleon low-energy constants from the consistency of interactions and currents in chiral effective field theory. *Phys. Rev. Lett.*, 103:102502, Sep 2009.
- [27] J. Goldstone. Field theories with superconductor solutions. *Il Nuovo Cimento (1955-1965)*, 19(1):154–164, Jan 1961.
- [28] Jan Govaerts and Jose-Luis Lucio-Martinez. Nuclear muon capture on the proton and ^3He within the standard model and beyond. *Nuclear Physics A*, 678:110–146, 09 2000.
- [29] 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, 05 2018.
- [30] 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.
- [31] 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.
- [32] 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, Frederick Gray, Z Hartwig, D W Hertzog, B Johnson, P Kammel, B Kiburg, and B Wolfe. Measurement of the positive muon lifetime and determination of the fermi constant to part-per-million precision. *Physical review letters*, 106:041803, 01 2011.

- [33] R Machleidt and D Entem. Chiral effective field theory and nuclear forces. *Physics Reports-review Section of Physics Letters - PHYS REP-REV SECT PHYS LETT*, 503, 05 2011.
- [34] Ruprecht Machleidt. Chiral Symmetry and the Nucleon-Nucleon Interaction. *Symmetry*, 8(4):26, 2016.
- [35] L. E. Marcucci, A. Kievsky, S. Rosati, R. Schiavilla, and M. Viviani. Chiral effective field theory predictions for muon capture on deuteron and ^3He . *Phys. Rev. Lett.*, 108:052502, Jan 2012.
- [36] L. E. Marcucci, A. Kievsky, S. Rosati, R. Schiavilla, and M. Viviani. Erratum: Chiral effective field theory predictions for muon capture on deuteron and ^3He [phys. rev. lett. 108, 052502 (2012)]. *Phys. Rev. Lett.*, 121:049901, Jul 2018.
- [37] L. E. Marcucci, M. Piarulli, M. Viviani, L. Girlanda, A. Kievsky, S. Rosati, and R. Schiavilla. Muon capture on deuteron and ^3He . *Phys. Rev. C*, 83:014002, Jan 2011.
- [38] Michael Murray. *Muon Catalyzed Fusion Effects in the Precision Measurement of Muon Capture on the Deuteron*. PhD thesis, University of Washington, 2017.
- [39] Yoichiro Nambu. Quasi-particles and gauge invariance in the theory of superconductivity. *Phys. Rev.*, 117:648–663, Feb 1960.
- [40] 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*, 466:213216, 2010.
- [41] Randolph Pohl, François Nez, Luis Fernandes, F Amaro, François Biraben, Joo Cardoso, D Covita, Andreas Dax, Satish Dhawan, Marc Diepold, Adolf Giesen, A.L. Gouvea, Thomas Graf, Theodor Haensch, Paul Indelicato, Lucile Julien, Paul Knowles, Franz Kottmann, Eric-Olivier Le Bigot, and Aldo Antognini. Laser spectroscopy of muonic deuterium. *Science*, 353:669–673, 08 2016.

- [42] 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, 2014.
- [43] Fabio Sauli. Principles of operation of multiwire proportional and drift chambers. page 92 p, Geneva, 1977. CERN, CERN. CERN, Geneva, 1975 - 1976.
- [44] 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.
- [45] M. Tanabashi, K. Hagiwara, K. Hikasa, K. Nakamura, Y. Sumino, F. Takahashi, J. Tanaka, K. Agashe, G. Aielli, C. Amsler, M. Antonelli, D. M. Asner, H. Baer, Sw. Banerjee, R. M. Barnett, T. Basaglia, C. W. Bauer, J. J. Beatty, V. I. Belousov, J. Beringer, S. Bethke, A. Bettini, H. Bichsel, O. Biebel, K. M. Black, E. Blucher, O. Buchmuller, V. Burkert, M. A. Bychkov, R. N. Cahn, M. Carena, A. Ceccucci, A. Cerri, D. Chakraborty, M.-C. Chen, R. S. Chivukula, G. Cowan, O. Dahl, G. D'Ambrosio, T. Damour, D. de Florian, A. de Gouvêa, T. DeGrand, P. de Jong, G. Dissertori, B. A. Dobrescu, M. D'Onofrio, M. Doser, M. Drees, H. K. Dreiner, D. A. Dwyer, P. Eerola, S. Eidelman, J. Ellis, J. Erler, V. V. Ezhela, W. Fetscher, B. D. Fields, R. Firestone, B. Foster, A. Freitas, H. Gallagher, L. Garren, H.-J. Gerber, G. Gerbier, T. Gershon, Y. Gershtein, T. Gherghetta, A. A. Godizov, M. Goodman, C. Grab, A. V. Gribsan, C. Grojean, D. E. Groom, M. Grünewald, A. Gurtu, T. Gutsche, H. E. Haber, C. Hanhart, S. Hashimoto, Y. Hayato, K. G. Hayes, A. Hebecker, S. Heinemeyer, B. Heltsley, J. J. Hernández-Rey, J. Hisano, A. Höcker, J. Holder, A. Holtkamp, T. Hyodo, K. D. Irwin, K. F. Johnson, M. Kado, M. Karliner, U. F. Katz, S. R. Klein, E. Klempt, R. V. Kowalewski, F. Krauss, M. Kreps, B. Krusche, Yu. V. Kuyanov, Y. Kwon, O. Lahav, J. Laiho, J. Lesgourgues, A. Liddle, Z. Ligeti, C.-J. Lin, C. Lippmann, T. M. Liss, L. Littenberg, K. S. Lugovsky, S. B. Lugovsky, A. Lusiani, Y. Makida, F. Maltoni, T. Mannel, A. V. Manohar, W. J. Marciano, A. D. Martin, A. Masoni, J. Matthews, U.-G. Meißner, D. Milstead, R. E. Mitchell, K. Mönig, P. Molaro, F. Moortgat, M. Moskvic, H. Murayama, M. Narain, P. Nason, S. Navas, M. Neubert, P. Nevski, Y. Nir, K. A. Olive, S. Pagan Griso, J. Parsons, C. Patrignani, J. A. Peacock, M. Pennington, S. T. Petcov, V. A. Petrov, E. Pianori, A. Piepke, A. Pomarol, A. Quadt, J. Rademacker, G. Raffelt, B. N. Ratcliff, P. Richardson, A. Ringwald, S. Roesler, S. Rolli, A. Romaniouk, L. J. Rosenberg, J. L. Rosner, G. Rybka, R. A. Ryutin, C. T. Sachrajda, Y. Sakai, G. P. Salam, S. Sarkar, F. Sauli, O. Schneider,

- K. Scholberg, A. J. Schwartz, D. Scott, V. Sharma, S. R. Sharpe, T. Shutt, M. Silari, T. Sjöstrand, P. Skands, T. Skwarnicki, J. G. Smith, G. F. Smoot, S. Spanier, H. Spieler, C. Spiering, A. Stahl, S. L. Stone, T. Sumiyoshi, M. J. Syphers, K. Terashi, J. Terning, U. Thoma, R. S. Thorne, L. Tiator, M. Titov, N. P. Tkachenko, N. A. Törnqvist, D. R. Tovey, G. Valencia, R. Van de Water, N. Varelas, G. Venanzoni, L. Verde, M. G. Vincter, P. Vogel, A. Vogt, S. P. Wakely, W. Walkowiak, C. W. Walter, D. Wands, D. R. Ward, M. O. Wascko, G. Weiglein, D. H. Weinberg, E. J. Weinberg, M. White, L. R. Wiencke, S. Willocq, C. G. Wohl, J. Womersley, C. L. Woody, R. L. Workman, W.-M. Yao, G. P. Zeller, O. V. Zenin, R.-Y. Zhu, S.-L. Zhu, F. Zimmermann, P. A. Zyla, J. Anderson, L. Fuller, V. S. Lugovsky, and P. Schaffner. Review of particle physics. *Phys. Rev. D*, 98:030001, Aug 2018.
- [46] E Tognelli, Scilla Degl’Innocenti, L.E. Marcucci, and Pier Giorgio Prada Moroni. Astrophysical implications of the proton-proton cross section updates. *Physics Letters B*, 742, 11 2014.
- [47] Timo van Ritbergen and Robin G. Stuart. On the precise determination of the fermi coupling constant from the muon lifetime. *Nuclear Physics B*, 564(3):343 – 390, 2000.
- [48] I-T. Wang. Muon capture by deuterons. *Phys. Rev.*, 139:B1539–B1544, Sep 1965.
- [49] 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.

Appendix A

KINETIC SOLUTIONS

The kinematic process for muons in deuterium gas can be seen in Fig. A.1, and the corresponding rates at various temperatures can be seen in Table A.1. Here, we write down and solve the system of coupled differential equations to obtain the relative populations of μd atoms from both hyperfine states as well as the population of $\mu^3\text{He}$.

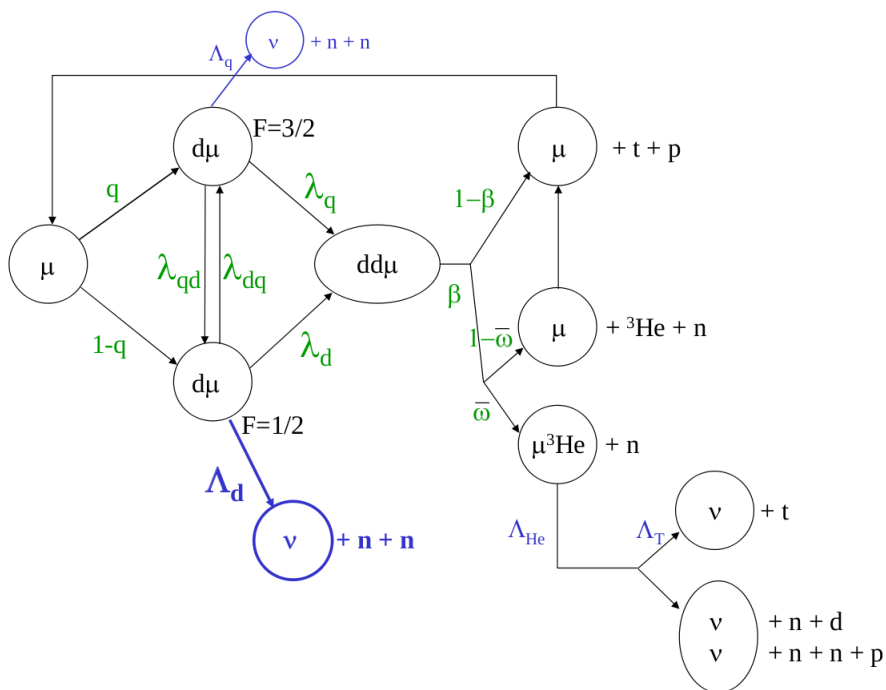


Figure A.1: Simplified diagram of Muon Kinetics within deuterium gas.

Process	Symbol	Value (μs^{-1})	
<i>Temperature Dependent</i>		300K	30K
hf transition q \rightarrow d	λ_{qd}	35(5)	37.0(4)
hf transition d \rightarrow q	λ_{dq}	10.5	0
Quartet dd μ formation	λ_q	3.75	3.98(5)
Doublet dd μ formation	λ_d	2.549(23)	0.053
Effective Fusion Fraction	β (unitless)	0.590(6)	0.517(15)
<i>Temperature Independent</i>		Value (s^{-1})	
3He total Capture Rate	Λ_{He}	2216(70)	
3He Partial Capture Rate	Λ_T	1496.0(40)	
μ d Quartet Capture Rate	Λ_q	10	
μ d Doublet Capture Rate	Λ_d	400	
Sticking Probability	$\bar{\omega}$ (unitless)	0.1206(6)	

Table A.1: Rates and branching ratios for Muon kinetics.

A.1 Quartet Population

Let us first focus on the population of the F=3/2 quartet hyperfine state. The number of muonic atoms in the quartet state is decreased when the muon decays via ordinary decay with a rate λ_{μ^+} , nuclear capture occurs with a rate Λ_q , the atom transitions to the doublet state, in a collisional de-excitation dependent on the density, with a rate $\Phi\lambda_{qd}$, or a dd μ molecule is formed in a mostly resonant formation process also proportional to the density with a rate $\Phi\lambda_q$.

A muon can enter the quartet from the doublet state via the hyperfine transition with rate $\Phi\lambda_{dq}$. A muon can also be recycled into the quartet state after catalyzing either a $p + t$

or ${}^3\text{He}$ fusion process. The branching ratio of a muon remaining after a muon catalyzed fusions one minus the probability of the muon sticking to the ${}^3\text{He}$, $(1 - \beta\bar{\omega})$, such that the rate of a muon originally in the quartet state undergoing fusion and re-entering the quartet state is given by $\Phi\lambda_q q(1 - \beta\bar{\omega})$, and the rate of a muon starting in the doublet state being recycled into the quartet state after fusion is given by $\Phi\lambda_d q(1 - \beta\bar{\omega})$

The time dependent population of muonic atoms within the quartet hyperfine state is then given by

$$\frac{dN_q(t)}{dt} = N_q(t) \left[-\Lambda_q - \lambda_{\mu^+} - \Phi\lambda_{qd} - \Phi\lambda_q \left(1 - q(1 - \beta\bar{\omega}) \right) \right] + N_d(t) \left[\Phi\lambda_{dq} + \Phi\lambda_d q \left(1 - \beta\bar{\omega} \right) \right] \quad (\text{A.1})$$

A.2 Doublet Population

Next, we determine the population of the $F=1/2$ doublet state. Similar to the quartet state, there are loss terms due to nominal muon decay with rate λ_{μ^+} , nuclear capture from the doublet state with rate Λ_d , a hyperfine transition into the quartet state with rate $\Phi\lambda_{dq}$, and resonant molecular formation with rate $\Phi\lambda_d$.

Just as with the quartet state, a muon can enter the quartet state via hyperfine transition from the quartet state with rate $\Phi\lambda_{qd}$, or a muon can be recycled back into the doublet state after catalyzing a fusion process.

The time dependent population of muonic atoms within the quartet hyperfine state is then given by

$$\frac{dN_d(t)}{dt} = N_d(t) \left[-\Lambda_d - \lambda_{\mu^+} - \Phi\lambda_{dq} - \Phi\lambda_d \left(1 - (1-q)(1 - \beta\bar{\omega}) \right) \right] + N_q(t) \left[\Phi\lambda_{qd} + \Phi\lambda_q (1-q)(1 - \beta\bar{\omega}) \right] \quad (\text{A.2})$$

A.3 μ^3He Population

The population of μ^3He molecules is decreased by either free muon decay or nuclear capture with rate Λ_{3He} . New μ^3He molecules are created when a molecular formation from either hyperfine state leads to a 3He fusion reaction in which the muon sticks to the 3He , with the overall branching ratio $\beta\bar{\omega}$. The time dependent μ^3He population can then be written

$$\frac{dN_{\mu^3He}(t)}{dt} = N_{\mu^3He}(t) \left[-\Lambda_{3He} - \lambda_{\mu^+} \right] + N_d(t) \left[\Phi\lambda_d\beta\bar{\omega} \right] + N_q(t) \left[\Phi\lambda_q\beta\bar{\omega} \right] \quad (\text{A.3})$$

A.4 Combined equation

The system of three coupled equations written above can be written in matrix notation as

$$\frac{d\vec{N}(t)}{dt} = k\vec{N}(t) \quad (\text{A.4})$$

with a population vector defined as

$$\vec{N}(t) = \begin{pmatrix} N_q(t) \\ N_d(t) \\ N_{\mu^3He}(t) \end{pmatrix} \quad (\text{A.5})$$

and the matrix k defined as

$$\begin{pmatrix} -\Lambda_q - \lambda_{\mu^+} - \Phi\lambda_{qd} - \Phi\lambda_q(1 - q(1 - \beta\bar{\omega})) & \Phi\lambda_{dq} + \Phi\lambda_dq(1 - \beta\bar{\omega}) & 0 \\ \Phi\lambda_{qd} + \Phi\lambda_q(1 - q)(1 - \beta\bar{\omega}) & -\Lambda_d - \lambda_{\mu^+} - \Phi\lambda_{dq} - \Phi\lambda_d(1 - (1 - q)(1 - \beta\bar{\omega})) & 0 \\ \Phi\lambda_q\beta\bar{\omega} & \Phi\lambda_d\beta\bar{\omega} & -\Lambda_{3He} - \lambda_{\mu^+} \end{pmatrix} \quad (\text{A.6})$$

As none of the coefficients in A.4 are time dependent, this system of linear differential equations has an analytical solution of the form

$$\vec{N}(t) = \sum_i c_i e^{-\Gamma_i t} \eta_i \quad (\text{A.7})$$

where Γ_i and η_i are the eigenvalues and eigenvectors of the system, and c_i are constants determined by the initial values. Alternatively, the solutions can be obtained numerically.

A.5 Solutions for MuSun Conditions

The conditions of the MuSun target gas are carefully selected to mitigate several systematic sources of error, as discussed in section 3.2.3. The solutions for the relative populations of muonic atoms in the doublet and quartet state, as well as the population of μ^3He for the standard operating temperature of 31K and density of $0.65 L_{N_2}$ are displayed in figure A.2. The observable time distribution of neutrons and electrons can be written in terms of the




Figure A.2: Solution to the kinematic system of equations (A.4) using the standard operating conditions of the target gas, at 31K and .065 LN density.

solutions for each population above. The electron time distribution is given by

$$el(t) = \lambda_{\mu+} \sum_i N_i(t) \quad (\text{A.8})$$

and the observation of neutrons from 3He fusion and $\mu + d$ capture will follow the following time distributions

$$\begin{aligned} fus(t) &= \beta\phi(\lambda_q N_q(t) + \lambda_d N_d(t)) \\ cap(t) &= 2(\Lambda_q N_q(t) + \Lambda_d N_d(t)). \end{aligned} \quad (\text{A.9})$$

These distributions can be seen in figure A.3.




Figure A.3: Observed decay distributions obtained from solutions of the kinematic system of equations at standard operating conditions.

Appendix B

RESOLVING THE KINETIC EQUATIONS WITH IMPURITY CAPTURES

In order to assess the effect of impurities within the target gas on the disappearance rate, the kinematic equations must be written to include additional capture terms. The population vector of equation A.5 must now include an additional term

$$N(t) = \begin{pmatrix} N_q(t) \\ N_d(t) \\ N_{\mu^3He}(t) \\ N_z(t) \end{pmatrix} \quad (\text{B.1})$$

where $N_z(t)$ denotes the population of muons which have atomically captured on an impurity. The differential equations governing the populations of the quartet and doublet state obtain an additional disappearance term proportional to the impurity concentration, c_z and the transfer rate of muons from deuterium to the impurity nucleus, Λ_{tr} . The matrix k is then given by

$$\begin{pmatrix} -\Lambda_q - \lambda_{\mu^+} - \Phi\lambda_{qd} - \Phi\lambda_q(1 - q(1 - \beta\bar{\omega})) - c_z\Lambda_{tr} & \Phi\lambda_{dq} + \Phi\lambda_d(1 - \beta\bar{\omega}) & 0 & 0 \\ \Phi\lambda_{qd} + \Phi\lambda_q(1 - q)(1 - \beta\bar{\omega}) & -\Lambda_d - \lambda_{\mu^+} - \Phi\lambda_{dq} - \Phi\lambda_d(1 - (1 - q)(1 - \beta\bar{\omega})) - c_z\Lambda_{tr} & 0 & 0 \\ \Phi\lambda_q\beta\bar{\omega} & \Phi\lambda_d\beta\bar{\omega} & -\Lambda_{3He} - \lambda_{\mu^+} & 0 \\ c_z\Lambda_{tr} & c_z\Lambda_{tr} & 0 & -\Lambda_z - \lambda_{\mu^+} \end{pmatrix} \quad (\text{B.2})$$

where Λ_z is the capture rate on the impurity.

As discussed in section 3.3.2, the most concerning impurity for the MuSun target gas is nitrogen. The solutions to the kinematic equations with 1 ppb of nitrogen at the standard running conditions can be seen in figure B.1.




Figure B.1: Solution to the kinematic system of equations using the standard operating conditions of the target gas with 1 ppb of nitrogen impurity.

To approximate the effect of a given concentration of nitrogen impurity on the observed disappearance rate, the kinematic equations can be solved with and without the additional capture term. The difference between the nominal electron distribution and the electron distribution observed with 1ppb of nitrogen impurity at the standard MuSun temperature and pressure can be seen in figure B.2. An exponential fit to each distribution yields a shift of 3.3 s^{-1} in the rate.




Figure B.2: The difference between the solution to the kinematic system of equations with no impurity, $el(t)$, and with 1ppb of nitrogen, $el_N(t)$, using the standard operating conditions of the target gas.

Appendix C

RESOLVING THE KINETIC EQUATIONS WITH PROTIUM IMPURITY

In order to assess the effect of a small concentration of protium within the target gas on the disappearance rate, the kinematic equations must be written to include additional terms for the disappearance of muons due to $pd\mu$ fusion. The high energy muons ejected following the fusion reaction are likely to leave the gaseous target and capture on wall materials. In order to estimate the worst case senario, it is assumed that all muons resulting from a $pd\mu$ fusion are lost.

The vector of equation A.5 must now include an additional term for the $pd\mu$ population

$$N(t) = \begin{pmatrix} N_q(t) \\ N_d(t) \\ N_{\mu^3He}(t) \\ N_{pd\mu}(t) \end{pmatrix} \quad (C.1)$$

where $N_{pd\mu}(t)$ denotes the population of $pd\mu$ molecules. The differential equations governing the populations of the quartet and doublet state obtain an additional disappearance term proportional to the protium concentration, c_p and the rate of molecular formation, $\lambda_{pd\mu}$ defined as $\Lambda_{pd\mu} = c_p\phi\lambda_{pd\mu}$. The matrix k is then given by

$$\begin{pmatrix} -\Lambda_q - \lambda_{\mu^+} - \Phi\lambda_{qd} - \Phi\lambda_q(1 - q(1 - \beta\bar{\omega})) - \Lambda_{pd\mu} & \Phi\lambda_{dq} + \Phi\lambda_dq(1 - \beta\bar{\omega}) & 0 & 0 \\ \Phi\lambda_{qd} + \Phi\lambda_q(1 - q)(1 - \beta\bar{\omega}) & -\Lambda_d - \lambda_{\mu^+} - \Phi\lambda_{dq} - \Phi\lambda_d(1 - (1 - q)(1 - \beta\bar{\omega})) - \Lambda_{pd\mu} & 0 & 0 \\ \Phi\lambda_q\beta\bar{\omega} & \Phi\lambda_d\beta\bar{\omega} & -\Lambda_{\mu^3He} - \lambda_{\mu^+} & \lambda_{f\gamma} * \lambda_f \\ \Lambda_{pd\mu} & \Lambda_{pd\mu} & 0 & -\lambda_f - \lambda_{\mu^+} \end{pmatrix} \quad (C.2)$$

where λ_f is the rate for the $pd\mu$ to undergo fusion.

These equations are assumed numerically, assuming a protium concentration of 20ppm, at the MuSun target conditions, and the relative populations can be seen in figure C.1.




Figure C.1: Relative populations of

Appendix D

QUADRAPOLE CURRENT ERROR AND EFFECT ON BEAM TUNING

A series of of upstream beam elements, described in section 4.2, work in unison to achieve the following

1. Produce a well focused muon beam at the target entrance window
2. Separate the muons from beam electrons
3. Divert the beam once a muon is detected in the target.

The various elements of the beamline, and their effects on the beam envelope in X and Y can be seen in figure D.1. All bending magnets are shown as blue boxes, while the quadrapole magnets are shown in red. The top half of the figure shows the beam envelope in Y, while the bottom half shows the beam envelope in X. The solid red curve shows the beam envelope as a function of distance along the beamline for nominal magnet currents. For the entirety of the R2014 dataset, upstream quadrapole magnet QSL2, labeled in the figure below, was running at half the nominal current value due to an error in the PSI database. The resulting beam envelope is shown as a solid black line.

A few key beamline elements are labeled in figure D.1. The first element specific to the MuSun beamline is the MORE kicker, consisting of two parallel plates which are rapidly charged to divert the beam by 13 mRad once a muon is registered in the entrance counters. Immediately following the kicker is the separator, an ExB field designed divert beam electrons without influencing the trajectory of muons. Both the kicked muons and beam electrons are diverted into the horizontal collimator located 3 m downstream in the bending magnet. In




Figure D.1: Beam profile in X and Y as a function of the distance along the beampipe. The beam profile with nominal QSL2 current is shown in red, and the beam profile with half the intended value of the QSL2 current is displayed in black. The red boxes indicate quadrapole magnets and bending magnets are shown in blue. The dashed lines indicate beam dispersion, while the solid lines show the beam width.

order to effectively divert muons and obtain optimal muon-electron separation, the divergence of the beam in the region of the kicker and separator should be small. This is clearly achieved in with the nominal current settings, but the mis-tuning of QSL2 results in a much larger divergence in the kicker and separator regions. In addition, the focus points at the collimator and downstream slit labeled are slightly displaced. Despite numerous efforts to improve the beam tune throughout the R2014 collection period, the reduced QSL2 current prevented adequate muon-electron separation, precipitating a large beam electron background.

Appendix E

ERROR INFLATION FOR ELECTRON DOUBLE COUNTING

E.1 Sources of doubly counted electron events

There are several cases in which a single event can be double counted within the electron detectors. Cosmic muon events, in which a cosmic muon passes directly through the electron hodoscope depositing energy on opposing pairs of scintillators. Additionally, afterpulsing and scattering within the scintillators can cause events in which a single electron is counted twice.

E.2 Error correction for doubly counted electrons

In the creation of the lifetime histograms, the error associated with each bin is given a default value of

$$\sigma = \sqrt{N} \tag{E.1}$$

where N is the total number of counts in the bin. This is incorrect for the treatment of events in which an electron is double counted. In this case, the events are completely correlated, and the error is given by

$$\sigma_2 = 2\sqrt{N_2} \tag{E.2}$$

where N_2 is the number of events in which a single electron is double counted. The correct error of each bin is then given by the single and double counted errors added in quadrature

$$\sigma = \sqrt{N_1 + 4N_2} \tag{E.3}$$

Written in terms of the observed double counted electrons, $\tilde{N}_2 = 2N_2$, we then have

$$\sigma = \sqrt{N_1 + 2\tilde{N}_2} = \sqrt{\tilde{N} + \tilde{N}_2} \tag{E.4}$$

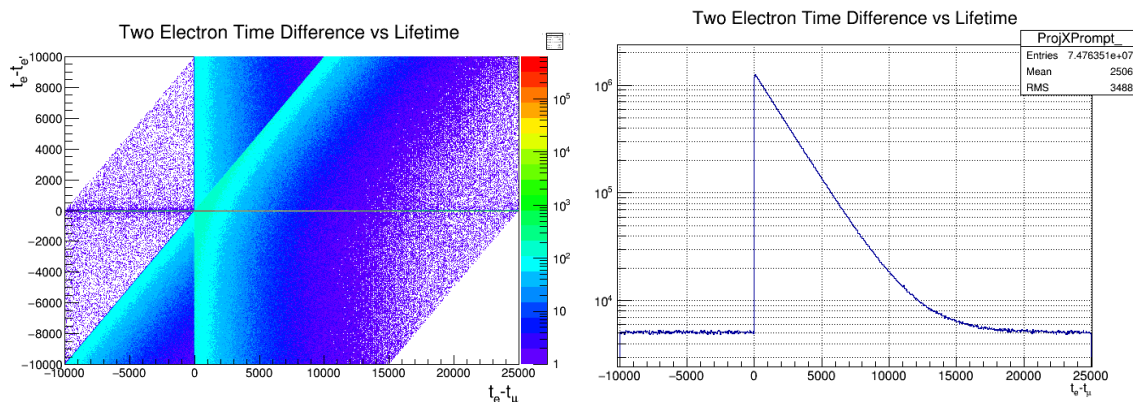
where $\tilde{N} = N_1 + \tilde{N}_2$ is the total number of observed events. For each lifetime fit, the fit is first performed to obtain the background rate. The error of each bin is then corrected to

$$\sigma = \sqrt{\tilde{N}(1 + n)} \tag{E.5}$$

where $n = \tilde{N}_2/\tilde{N}$, and the fit is repeated with the corrected bin errors.

E.3 Determination of double counted fraction

In order to correct the error in each bin, we must determine the fraction of double counted events in each bin. A Lifetime histogram of the two electron events versus the time difference between the two electrons can be seen in figure E.1, along with an x projection showing the prompt coincidences within +/- 10ns. The lifetime component comes from either decay events with an accidental electron within 10ns, or decay events with afterpulsing or branching in the scintillator such that two electrons are created.

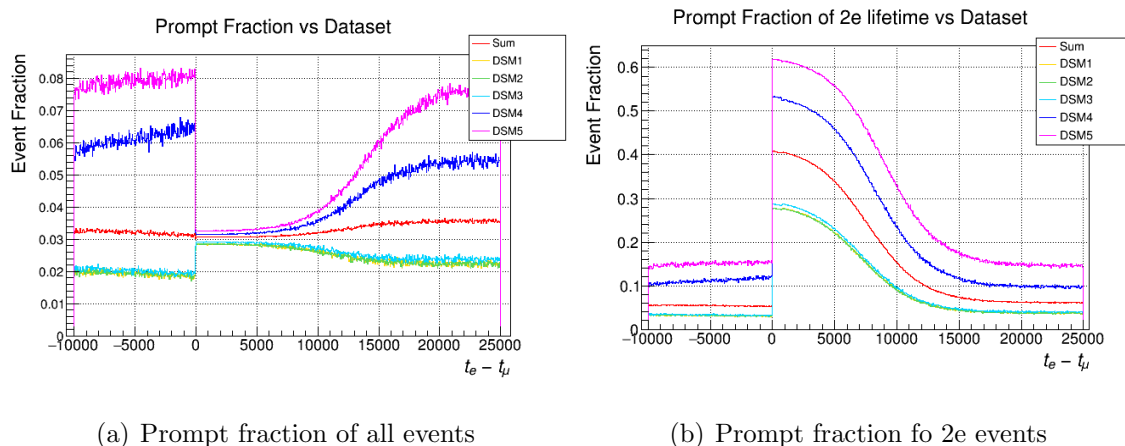


(a) Two electron time difference vs lifetime (b) Lifetime of electrons in prompt coincidence

Figure E.1: (a) Time difference between two electrons versus the time of each electron relative to the muon entrance time and (b) X project of the prompt electrons within a coincidence window of +/-10ns.

Todo: calculate the probability of getting an accidental within 10ns of a michel.

Let's first assume all prompt events are due to some detector effect and contribute to double counting, to see the maximal shift on the lifetime fit. By dividing the prompt two electron lifetime by the full electron lifetime, or the two electron lifetime, one can determine n , the fraction of events in each bin which are double counted. These can be seen for each dataset in figure E.2. The left distribution shows the fraction of doubly counted events relative to all electron events, and can be interpreted as follows. At early times there is some Michel related double counting which is seen nearly equally in all datasets. At later times, there is a background related double counting which is seen as a larger fraction in the later datasets, in which the overall background was reduced. In the earlier datasets, where the beam electron background was highest, this dominates the contribution at later times and a smaller relative fraction of doubly counted events is observed.



(a) Prompt fraction of all events

(b) Prompt fraction fo 2e events

Figure E.2: The fraction of double counted events in each time bin with respect to (a) all electron events and (b) two electron events, determined by dividing the lifetime distribution of events with two electrons in prompt coincidence by the nominal lifetime distributions.

E.4 Error correction for double counted events

Using the prompt two electron distribution as the number of doubly counted events, a correction can be made to each bin of the lifetime fit via equation 4, for each bin, the error is adjusted from $\sqrt{\tilde{N}}$ to $\sqrt{\tilde{N} + \tilde{N}_2}$. The lifetime fit parameters for both the full electron and two electron lifetime with and without the error inflation for each dataset can be seen in figure E.3. For the full lifetime distribution, the rate is unaffected by the error inflation, but

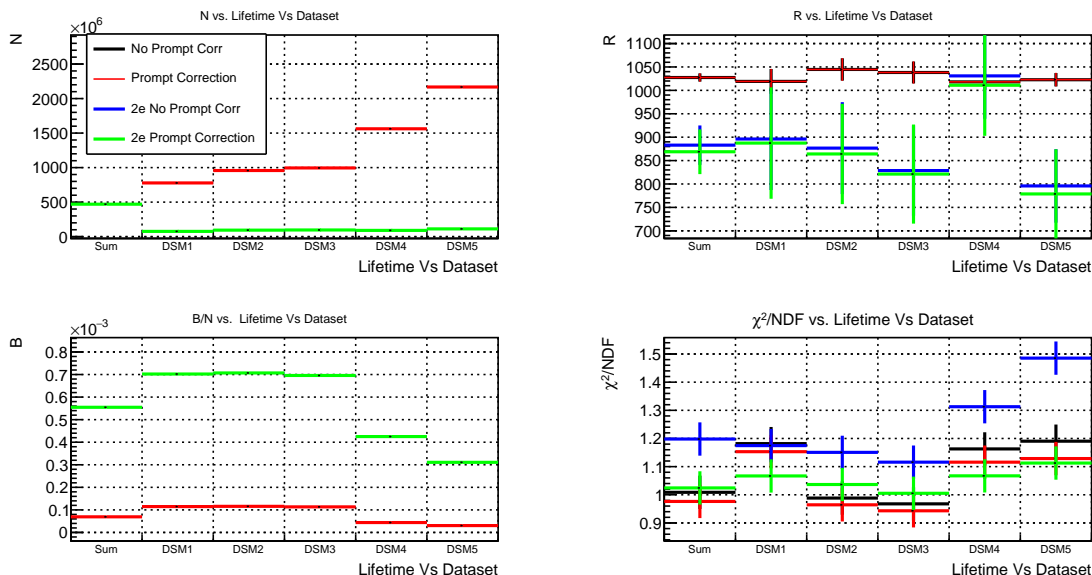


Figure E.3: Lifetime fit parameters for the full lifetime fit with and without (red and black) corrected bin errors to accounted for double counted events, in addition to the fit parameter for the two electron version of the lifetime histogram with and without (green and blue) the error inflation.

the chi squared of the fit for all datasets is reduced. For the two electron distributions, the lifetime is decreased by 15-20Hz and the chi squared is improved for all datasets.

E.5 Correction for non-correlated events

In principal, there are two electron events in which the two electrons are uncorrelated; one is a Michel electron and the other is an accidental that happened to come at the same time. To correct the prompt distribution used above, we can use the lifetime distribution from non-prompt events in an equivalent time window and subtract from the prompt distribution. Using an equivalently sized and symmetric off axis projection of figure E.1 from $[-30,-20]$ ns and $[20,30]$ ns, the non-prompt distribution is obtained. This is compare to the prompt distribution in figure E.4




Figure E.4: Comparison of the prompt distribution for two electron events within $[-10,10]$ ns (blue) to the non-prompt distribution created from to electron events within $[-30,-20]$ and $[20,30]$ ns time windows.

The lifetime fits with the error inflation was repeated with the versions of the prompt histograms in which the non-prompt contribution was subtracted. The obtained rates are unaffected, as shown in figure E.5.

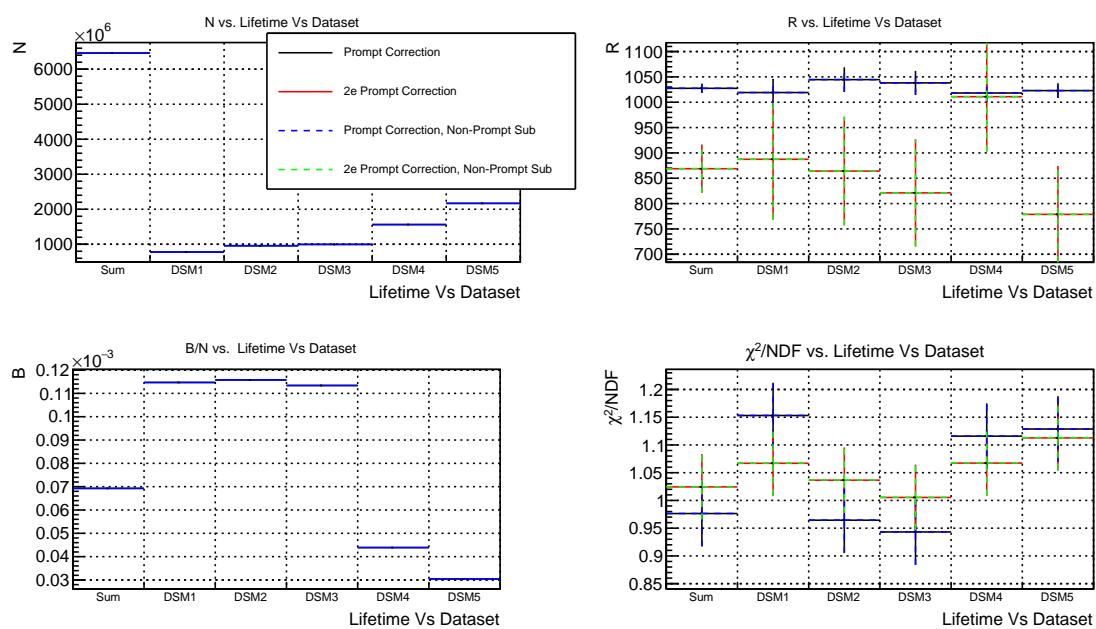


Figure E.5: Comparison of the fit results for the full and two electron lifetimes with error in inflation using both the full prompt distribution (solid) and the non-prompt subtracted prompt distribution (dashed).