

Muon-Catalyzed Fusion Effects in the Precision Measurement of Muon
Capture on the Deuteron

Michael H. Murray

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

Doctor of Philosophy

University of Washington

2017

Reading Committee:

Peter Kammel, Chair

David Hertzog

Martin Savage

Program Authorized to Offer Degree:
Physics

©Copyright 2017

Michael H. Murray

University of Washington

Abstract

Muon-Catalyzed Fusion Effects in the Precision Measurement of Muon Capture on the Deuteron

Michael H. Murray

Chair of the Supervisory Committee:
Professor Peter Kammel
Physics

The MuSun experiment will measure the rate of muon capture on the deuteron to 1.5 % precision. This reaction is related to solar pp fusion and the νd scattering reaction at the Sudbury Neutrino Observatory through effective field theories like chiral perturbation theory. The Gamow-Teller transition in all of these processes contains a single unknown low-energy constant that determines the strength of the axial coupling to the two-nucleon system. Muon capture on the deuteron provides a precise and theoretically clean determination of this low-energy constant. The capture rate is determined by comparing the free muon lifetime to a 10 ppm measurement of the lifetime of negative muons stopped in a deuterium target. MuSun achieves this precision by tracking muons with a cryogenic time-projection chamber to ensure they stop in deuterium. This dissertation characterizes and quantifies a systematic measurement error known as fusion interference, which is a class of tracking error caused by muon-catalyzed fusion reactions following the muon stop. An efficiency difference between events with and without fusion leads to a non-exponential decay time distribution and causes a shift in the measured lifetime. The design and operation of the MuSun experiment and the data analysis procedures are summarized. A formalism is developed to describe the parameters of fusion interference and a correction procedure based on a specialized muon tracking algorithm is presented.

ACKNOWLEDGMENTS

The work I describe in this thesis is only possible because of the contributions of many friends and colleagues, especially the MuSun collaboration, with whom I am very fortunate to have had the opportunity to work. My development as a student and researcher has been undeniably improved by the joint effort running the experiment and performing the scientific analysis, and on a personal level, I am grateful for the chance to know everyone from the several research groups in the collaboration. For advice and tutelage, I thank Fred Gray, Tim Gorringer, Volodya Tischenko, Rob Carey, Claude Petitjean, Alexey Vorobyov, Alexander Vasilyev, Marat Vznuzdaev, Peter Kravtsov, Nikolay Voropaev, Alexander Nadtochy, and Volodya Ganzha. The students and post-docs from the various MuSun institutions share with me the hectic experience of rapidly learning practical experimental physics while losing sleep running shifts for months at a time at PSI. To Justin Phillips, Luis Ibanez, Nandita Raha, Xiao Luo, Kres Neely, Wes Gohn, Emma Barnes, Ilya Ivanov, and Marco Mariti, I am grateful that we all learned together and found time for friendship amidst all the hard work. I want to thank Bernhard Lauss, Françoise Mulhauser, and Malte Hildebrandt, who graciously provided an outsized amount of their time and guidance.

I am grateful to everyone at the CENPA laboratory, including the students, professors, and lab staff who have all contributed, directly or indirectly, to my education and research. The muon group at the University of Washington has been a wonderful place to work and learn, and I owe thanks to the many people that make it great: Frederik Wauters, Rachel Ryan, Pete Alonzi, Duncan Prindle, Ethan Muldoon, Dan Salvat, Serdar Kizilgul, Brendan Kiburg, Peter Winter, Sara Knaack, David Webber, and Jason Crnkovic, and of course to my advisors, David Hertzog and Peter Kammel. Dave, thank you for bringing your lighthearted wit, piercing scientific insight, and drive for pragmatism to both the experiment and my tenure as a graduate student. Peter, I am deeply grateful for your commitment to

the pedagogical enrichment of the students of MuSun. Whether during late-night detector debugging, intense collaboration meetings, or weekend ski trips, I have always been learning from your patient instruction.

Thank you Mom, Dad, Scott, and John for instilling in me an interest in science and the natural world and encouraging and enabling me to pursue the study of physics. Alexis, your unfaltering support, patience, and advice has meant the world to me.

This material is based upon work supported by the U.S. Department of Energy Office of Science, Office of Nuclear Physics under Award Number DE-FG02-97ER41020. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1053575.

TABLE OF CONTENTS

	Page
Chapter 1: Introduction	1
Chapter 2: Theory and Motivation	5
2.1 Muon Decay	5
2.2 Weak Nuclear Processes	6
2.3 Physics of Light Nuclei	10
2.4 Chiral Symmetry	11
2.5 Chiral Perturbation Theory	12
2.5.1 Hybrid χ PT	13
2.5.2 Pionless Effective Field Theory	13
2.6 Few-Nucleon Observables	15
2.7 Calculation of Two- and Three-Nucleon Observables in HB χ PT	16
2.8 History of Λ_d Calculations	20
2.9 Summary	21
Chapter 3: Experimental Design	23
3.1 Lifetime Method	23
3.2 Deuterium Target and CryoTPC	24
3.3 Muon Kinetics	26
3.3.1 Formation of μd Atoms	26
3.3.2 Hyperfine Structure	26
3.3.3 Molecular Formation and Muon-Catalyzed Fusion	28
3.3.4 Target Gas Conditions	30
3.3.5 Numerical Analysis	31
3.3.6 Muon Capture on Nuclei Other than Deuterium	35
3.3.7 Muon-Catalyzed Fusion with Protium	36
3.4 History of Λ_d Measurements	38
3.4.1 Wang et al.	39
3.4.2 Bertin et al.	39

3.4.3	Bardin et al.	40
3.5	Summary	40
Chapter 4:	Experimental Apparatus	41
4.1	π E3 Muon Beamline	42
4.1.1	Kicker	44
4.2	Muon Entrance Detectors	45
4.3	CryoTPC	48
4.3.1	μ SR Magnet	51
4.4	Electron Detectors	51
4.4.1	Electron Multi-Wire Proportional Chambers	52
4.4.2	Electron Scintillators	54
4.5	Neutron Detectors	56
4.6	Deuterium Gas	56
4.7	Data Acquisition System	58
4.8	R2011 Beamtime Summary	60
Chapter 5:	Software and Analysis	64
5.1	Data Analysis	64
5.1.1	Data Organization	65
5.1.2	Analysis Stages	66
5.1.3	Diagnostics	67
5.1.4	Error Identification and Rejection	67
5.1.5	Standardized Analysis Interfaces	68
5.2	Muon Events	69
5.2.1	μ SC	69
5.2.2	μ PC	69
5.2.3	Kicker	70
5.2.4	Pile-up Protection	70
5.3	TPC	71
5.3.1	Waveforms and TPC Event Topologies	72
5.3.2	Pulse-Finding	77
5.3.3	Pulse Clustering	79
5.3.4	Muon Tracking	80
5.3.5	TPC Fiducial Volume	87
5.4	Decay Electron Tracking	88

5.4.1	eSC	88
5.4.2	ePC	90
5.4.3	Electron Track Objects	91
5.5	Muon Decay Events	91
5.6	Lifetime Fitting	92
5.7	Blinding in MuSun	94
Chapter 6:	Systematic Errors	96
6.1	Chemical Gas Impurities	96
6.2	Stops in Materials Other than Deuterium	97
6.3	Electron Interference	99
6.4	Deviation From Exponential Decay-Time Distribution	100
Chapter 7:	Fusion Interference	102
7.1	Distortion of the Observed Decay Time Distribution	103
7.2	Fusion Migration and Tracker Errors	106
7.2.1	Monte Carlo Simulation Isolating Fusion Migration	108
7.2.2	Properties of Tracker Errors	109
7.2.3	Fusion Time Dependence of $\eta(t_f)$	113
7.2.4	Separation of $\eta(t_f)$ Shape and Magnitude	119
7.2.5	Lifetime Shift from Net Migration	119
7.3	Migration Formalism	121
7.4	Correction Strategy	124
7.5	Modeling Fusion Migration	125
7.5.1	Fusion Event Tagging	125
7.5.2	Migration Probability Based on Muon Penetration Depth	128
7.5.3	Interpolated Muon Stop Density	133
7.5.4	Determining the Fusion Fraction	135
7.5.5	Migration in R2011 Data Set	137
7.6	Disappearance Rate Correction	140
7.7	Results	143
7.8	Outlook	145
7.8.1	Direct Measurement of $\eta(t_f)$	145
7.8.2	Fusion Tag	145
7.8.3	Fusion Fraction	145
7.8.4	Delayed Fusion Tag for Muon Stop	146

Chapter 8: Conclusion	148
8.1 Upgrade Program Informed by R2011 Analysis	148
8.1.1 Impurity Monitor	148
8.1.2 Wall Stop Monitor	149
8.1.3 Stop Distribution Balancing	150
8.2 Conclusion	150
Glossary	153
Bibliography	158

Chapter 1

INTRODUCTION

The muon is an unstable fundamental particle with similar properties to the electron, but a mass of $106 \text{ MeV}/c^2$. It is part of the second generation of leptons, which include the electron and tau as well as the three neutrinos. The primary decay mode is to an electron and two neutrinos,

$$\mu^- \rightarrow e^- + \nu_\mu + \bar{\nu}_e \quad (1.1)$$

with a lifetime of $\tau = 2.1969811(22) \text{ } \mu\text{s}$ [1][2].

Because this short lifetime would classically limit the range of muons, even moving at the speed of light, their observation on the surface of the earth provided an early test of time dilation in Einstein's theory of special relativity [3]. In modern physics, muons continue to find use as probes of fundamental physics through their decay properties and interaction with matter. They are used in precision tests and searches for new physics such as the Muon $g - 2$ measurement of the anomalous magnetic moment [4] and experiments attempting to measure charged lepton flavor violation [5][6]. Negative muons form hydrogen-like atoms with ordinary nuclei and undergo capture with protons in these nuclei in a process that is analogous to β -decay. Spectroscopy of muonic atoms and measurement of capture rates continues to determine basic properties of light nuclei.

The MuSun experiment is a measurement of nuclear capture of negative muons on the deuteron (the nucleus of a ^2H atom) from the doublet hyperfine level of the ground state muonic atom:

$$(\mu^- d)_{F=1/2} \rightarrow \nu_\mu + n + n. \quad (1.2)$$

This reaction is sensitive to the two-nucleon interaction between the proton and neutron. It proceeds at a rate of $\Lambda_d \approx 400 \text{ s}^{-1}$, which is substantially smaller than the free muon decay rate $\lambda_0 = 1/\tau = 455170 \text{ s}^{-1}$. The rate Λ_d will be measured to 1.5 % by determining the difference in lifetime between muons stopped in deuterium gas and muons decaying via

reaction (1.2).

Nuclear physics in the low-energy regime cannot be described by perturbative QCD because renormalization of the coupling constant leads to a strong coupling at low energy ($\lesssim 1$ GeV). Instead, an effective field theory (EFT) can be formulated that considers only low-energy degrees of freedom relevant for these processes. High-energy (short-range) particle interactions are integrated into simplified effective interactions. In nuclear physics, low-energy EFTs treat the short-range behavior of particles heavier than pions (or even pions themselves) as contact interactions with coefficients that must be determined empirically for the theory to become predictive.

The weak nuclear form factors of the proton and neutron, g_V , g_A , g_M , and g_P , are the coefficients of the allowed Lorentz-covariant components of the generalized nuclear current that describes single-nucleon weak interactions like neutron β -decay. Application of the single-nucleon form factors to the two-body system while assuming no proton-neutron interaction is called the impulse approximation (IA), and accounts for only ~ 90 % of the total doublet capture rate Λ_d [7]. Historically, contributions to the remainder of this rate have been modeled with an exchange of light mesons like the pion and η , but modern calculations are performed with EFTs such as chiral-perturbation theory (χ PT) and pionless-EFT [8]. In these theories, the two-nucleon interaction is parameterized by a single low-energy constant (LEC) called d_R or L_{1A} that must be determined by experiment or calculated with a first-principles method like lattice QCD. This constant appears not only in two-nucleon observables, but contributes to multi-nucleon systems. MuSun will reduce the uncertainty in this LEC from 100 % to ~ 20 % in the two-nucleon system.

The physics of two-body nuclear interactions has connections with many areas of nuclear physics. The most concrete is the connection to the other two-nucleon weak processes, pp -fusion and charged-current neutrino scattering on deuterons,

$$\begin{aligned} p + p &\rightarrow d + e^+ + \nu_e \\ \nu_e + d &\rightarrow p + p + e^-, \end{aligned}$$

which share the same weak hadronic vertex.

A precise measurement of muon capture on the deuteron determines the rate and cross-




Figure 1.1: Theoretical calculations and previous measurements of the doublet capture rate. The projected MuSun precision is the red error bar, with an arbitrary central value. The Bertin 1973 work is omitted because the interpretation of the result assumed a statistical mixture of μd hyperfine states, which is now known to be incorrect by up to a factor of 3.

section of these reactions directly. MuSun takes its name from this connection to the pp -fusion reaction that powers the sun.

Perhaps less obvious is the relation to transition matrix elements between heavy nuclei that are candidates for neutrino-less double β -decay ($0\nu\beta\beta$). Recent calculations of these matrix elements have included the two-body chiral interactions that are determined by studies of light nuclei. When compared with transitions involving only one-body interactions, inclusion of two-body forces can change the magnitude of the matrix element by 20 % or more [9][10].

The history of experimental attempts to determine d_R/L_{1A} is outlined in Fig. 1.1 and described in detail in Chapters 2 and 3. Experiments measuring reaction (1.2) do not achieve the precision on Λ_d needed to determine L_{1A} with an error better than 100 %, but theoretical

calculations have reached this precision due to substantial advancement in the last twenty years. MuSun is the first experiment capable of measuring Λ_d precisely enough to fix the two-body axial LEC entirely within a clean two-nucleon system. The MuSun collaboration has members from several institutions in Switzerland, Russia, and the United States. The experiment is operated at the Paul Scherrer Institute (PSI) in Villigen, Switzerland, with the most recent data run in autumn of 2016.

MuSun will achieve a 1.5 % precision measurement of Λ_d by careful consideration of muonic atomic and molecular kinetics, event selection with an active stopping target, and the use of a high-flux muon beam at PSI. The target gas density and temperature strongly affect the relative populations of the hyperfine states of the μd atom, so these are selected to optimize the time the average muon spends in the doublet state. The use of active ultrapure gas targets has been a crucial development for measurements of muon capture and muon-catalyzed fusion [11][12][13]. Operating the target as a time-projection chamber (TPC), individual muons are tracked to stringently require stops in the pure gas, far from the vessel walls, substantially reducing backgrounds. In this thesis, I will discuss the experimental design and gas conditions in Chapter 3 and the detector systems and data analysis in Chapters 4 and 5.

The advantages of an active target are not without cost, however, as event selection cuts in a lifetime experiment can be coupled to the muon decay time in a subtle fashion, distorting the otherwise exponential decay time distribution. In the MuCap experiment [13], the decay electron caused small signals in the TPC that biased event selection in favor of prompt muon decays. In MuSun, the dominant event selection systematic error is fusion interference, where events containing muon-catalyzed fusion are selected with different efficiency, biasing towards later decay times. In Chapter 7 of this thesis, I will describe this systematic error, a formalism to describe the errors induced by fusion, and a method to correct for the dominant error contribution.

This work is based on the analysis of the first large production dataset, R2011, collected in summer 2011. The analysis of this data influenced the optimization of the experiment leading up to the final production runs in R2014 and R2015 that collected the full statistics needed for the lifetime measurement.

Chapter 2

THEORY AND MOTIVATION

Low-energy weak interactions like muon decay and nuclear physics processes such as β -decay and muon capture take the form of a point-like current-current interaction with amplitude

$$\mathcal{M} = \frac{4G_F}{\sqrt{2}} J^\alpha J_\alpha, \quad (2.1)$$

where J^α represents a generic charge-raising or lowering weak current between fermions. The mass of the W boson mediating charged weak interactions is much larger than the energies and momenta involved in muon decay or nuclear processes, $q^2 \ll M_W^2 \approx (80 \text{ GeV})^2$, so in these cases, the propagator reduces to $1/M_W^2$ and the effective coupling constant is

$$\frac{G_F}{\sqrt{2}} = \frac{g^2}{8M_W^2} \quad (2.2)$$

where g is the dimensionless weak coupling constant.

The charged current $J_{ff'}^\alpha$ of two fundamental fermions f and f' is the matrix element of a current operator

$$\langle f' | \hat{J}^\alpha | f \rangle = \bar{\psi}_{f'} \gamma^\alpha \frac{1}{2} (1 - \gamma^5) \psi_f, \quad (2.3)$$

where $\bar{\psi}_{f'}$ and ψ_f are Dirac spinors, γ^α are the gamma matrices, and $\gamma^5 \equiv i\gamma^0\gamma^1\gamma^2\gamma^3$. The projection operator $(1 - \gamma^5)/2$ selects left-handed particles or right-handed anti-particles, reflecting the parity violation of the weak interaction. In the above, I use the convention of Halzen and Martin [14] to define the projection operators normalized with a factor of 2.

2.1 Muon Decay

Muons are unstable fundamental leptons whose primary decay mode is to an electron and neutrinos.

$$\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu \quad (2.4)$$

This is a leptonic weak process governed by the amplitude in Eq. 2.1 with currents defined by Eq. 2.3 [14]:

$$\mathcal{M} = \frac{G_F}{\sqrt{2}} \bar{u}_{\nu\mu} \gamma^\alpha (1 - \gamma^5) u_\mu \bar{u}_e \gamma_\alpha (1 - \gamma^5) v_{\nu e}. \quad (2.5)$$

Here, u_μ , \bar{u}_e , $v_{\nu e}$, and $\bar{u}_{\nu\mu}$ are Dirac spinors for the incoming muon and outgoing electron, neutrino, and anti-neutrino, respectively. The decay rate is obtained via integration over phase space Q

$$d\Gamma = \frac{1}{2E} |\overline{\mathcal{M}}|^2 dQ, \quad (2.6)$$

and an average is performed over all spin states to obtain

$$\Gamma = \frac{1}{\tau} = \frac{G_F^2 m^5}{192\pi^3} \quad (2.7)$$

where m is the mass of the muon and τ is the mean lifetime. This is a tree-level result that is modified by higher-order and radiative corrections in the full calculation.

The constant G_F is fixed by measurement of the muon lifetime, most recently by the 1 ppm measurement of the MuLan collaboration [1]. The world average for the lifetime (decay rate) is [2]

$$\tau = 2.1969811(22) \text{ } \mu\text{s} \quad (2.8)$$

$$\lambda_{\mu^+} = \frac{1}{\tau} = 455170.1 \pm 0.5 \text{ s}^{-1}. \quad (2.9)$$

When radiative corrections are included, the world-average value $G_F = 1.1663787(6) \times 10^{-5} \text{ GeV}^{-1}$ is obtained.

2.2 Weak Nuclear Processes

The current-current interaction of Eqs. 2.1 and 2.3 is written in terms of spinors for bare quarks and leptons, but in contrast to the muon decay example (Eq. 2.5), the quark is part of a bound state of a neutron or proton.

The first complication is that unlike the leptons, the left-handed propagating quark states are not the same as the states that participate in the weak interaction. This quark-sector mixing is not a consequence of any nuclear physics, but the effect appears in nuclear transitions through the elements of the unitary matrix that specifies the mixing. A heuristic

introduction to the quark mixing matrix is presented here, but a more thorough derivation can be found in Ref. [2].

The simplest evidence for mixing between the left-handed quark mass eigenstates and the eigenstates of the weak interaction is the decay of the charged kaons, K^\pm , which are bound states of a strange quark and an up or down quark [14]. To account for this, consider the left-handed lepton doublets that participate in the weak interaction:

$$\begin{pmatrix} e \\ \nu_e \end{pmatrix}_L, \begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix}_L, \begin{pmatrix} \tau \\ \nu_\tau \end{pmatrix}_L. \quad (2.10)$$

Instead of a doublet consisting of the propagating u and d quark states, use a modified quark doublet,

$$\begin{pmatrix} u \\ d' \end{pmatrix}_L. \quad (2.11)$$

where the state d' is a mixture of s_L and d_L states with mixing angle $\theta_C = 13^\circ$, known as the Cabibbo angle,

$$d'_L = \cos(\theta_C)d_L + \sin(\theta_C)s_L. \quad (2.12)$$

To include the third generation of quarks, this concept is generalized to the Cabibbo-Kobayashi-Maskawa (CKM) matrix,

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix}_L = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix}_L. \quad (2.13)$$

Alternatively, the CKM matrix can be defined in a manner that is symmetric with respect to up-type and down-type quarks using two unitary matrices V_L^u and V_L^d that account for mixing in the left-handed up-type and down-type quarks, respectively. The product of these matrices is the CKM matrix, $V_L^u V_L^{d\dagger} \equiv V_{\text{CKM}}$ [2]. The nuclear transition between p and n in β -decay or muon capture involves the CKM element $V_{ud} = 0.97427(15)$.

The nuclear processes of β -decay and muon-capture are semi-leptonic, meaning that the the current-current interaction of Eq. 2.1 is the product of a leptonic and hadronic current, $\mathcal{M} \propto L_\alpha J^\alpha$. For this section, I will consider muon capture on the proton,

$$\mu^- p \rightarrow n + \nu_\mu, \quad (2.14)$$

but similar equations apply to neutron β -decay. The leptonic current is given by Eq. 2.3,

$$\langle \nu_\mu | \hat{L}_\alpha | \mu \rangle = \bar{u}_{\nu_\mu} \gamma^\alpha \frac{1}{2} (1 - \gamma^5) u_\mu, \quad (2.15)$$

but a complete description of the interaction with the nucleon is complicated by the non-perturbative nature of QCD at the relevant energy scale. Instead, a phenomenological current is written down including all terms compatible with Lorentz covariance, split into vector terms V^α and axial terms A^α :

$$\begin{aligned} \langle n | \hat{J}^\alpha | p \rangle &= \bar{u}_n [V^\alpha - A^\alpha] u_p \\ V^\alpha &= g_V \gamma^\alpha + i \frac{g_M}{2m_N} \sigma^{\alpha\beta} q_\beta + \frac{g_S}{m_\mu} q^\alpha \\ A^\alpha &= g_A \gamma^\alpha \gamma^5 + i \frac{g_T}{2m_N} \sigma^{\alpha\beta} q_\beta \gamma^5 + \frac{g_P}{m_\mu} q^\alpha \gamma^5. \end{aligned} \quad (2.16)$$

The coefficients of these terms, g_V , g_M , g_S , g_A , g_T , g_P , sometimes called *nuclear form factors*, reflect that the nucleon is composed of quark and gluon constituents that are governed by fundamental QCD. The coefficients have a small dependence on the energy scale of the process being described, $g \rightarrow g(q^2)$. The coefficients g_S and g_T of the so-called *second-class* scalar and pseudo-tensor currents are expected to be zero due to their G-parity transformation properties [15]. G-parity is a generalization of charge conjugation symmetry to multiplets of particles like the proton and neutron or the pions. The strong interaction is invariant under G-parity, to the extent that it is isospin invariant, so the vector currents of Eqs. 2.16 with couplings induced by the strong interaction will retain the same G-parity transformation properties of the bare vector current in Eq. 2.3, and likewise for the axial terms [16]. However, scalar and pseudo-tensor currents transform with the opposite G-parity from the vector and axial vector terms, respectively. Specifically, the vector and magnetic terms transform with positive G-parity, but the scalar term has negative G-parity:

$$\begin{aligned} \bar{u}_p \gamma^\alpha u_n &\rightarrow +\bar{u}_n \gamma^\alpha u_p \\ \bar{u}_p \sigma^{\alpha\beta} u_n &\rightarrow +\bar{u}_n \sigma^{\alpha\beta} u_p \\ \bar{u}_p u_n &\rightarrow -\bar{u}_n u_p. \end{aligned}$$

Likewise, the axial vector and pseudoscalar terms transform with negative G-parity but the

pseudo-tensor term has positive G-parity:

$$\begin{aligned}\bar{u}_p \gamma^\alpha \gamma^5 u_n &\rightarrow -\bar{u}_n \gamma^\alpha \gamma^5 u_p \\ \bar{u}_p \gamma^5 u_n &\rightarrow -\bar{u}_n \gamma^5 u_p \\ \bar{u}_p \sigma^{\alpha\beta} \gamma^5 u_n &\rightarrow +\bar{u}_n \sigma^{\alpha\beta} \gamma^5 u_p.\end{aligned}$$

There is currently no experimental evidence for the existence of second-class currents.

With this parameterization of the hadronic weak current, the form factors are determined empirically. The vector coefficients g_V and g_M are determined from the electromagnetic form factors f_1 and f_2 by the conserved vector current hypothesis (CVC):

$$\begin{aligned}g_V(q^2 = 0) &= 1 \\ g_M(q^2 = 0) &= \frac{1}{2}(\mu_p - \mu_n - 1),\end{aligned}\tag{2.17}$$

where $\mu_p = 2.79$, $\mu_n = -1.91$ are the magnetic moments in units of the nuclear Bohr magneton [17]. The momentum dependence of the weak form factors is the same as the electromagnetic analogues.

The CVC hypothesis essentially states that the hadronic electromagnetic current and the weak charged current between neutron and proton states are components of a single isovector of current operators. The conservation of the electromagnetic current,

$$\frac{\partial j_{EM}^\alpha(x)}{\partial x^\alpha} = 0,\tag{2.18}$$

implies the conservation of the weak vector current $\bar{u}_p V^\alpha u_n$, and the weak form factors of the nucleon are directly related to the hadronic electromagnetic form factors (Eq. 2.17). In addition, the CVC hypothesis implies that the scalar current is zero separately from its status as a second-class current. This analogue allows the vector current form factors to be determined quite precisely via electromagnetic experiments such as electron scattering, magnetic resonance, and the spectrum of a hydrogen maser [18].

Muon capture on the proton (Eq. 2.14) and neutron β -decay,

$$n \rightarrow p + e^- + \bar{\nu}_e\tag{2.19}$$

are the most experimentally accessible processes related to the induced nuclear axial current in Eq. 2.16. The axial vector form factor $g_A = -1.2723 \pm 0.0023$ is determined via the

asymmetry in the β -decay of cold neutrons [2][19][20]. The contribution of the pseudoscalar coupling g_P is small compared to the axial current in experiments such as β -decay because of the factor of the momentum, q^α , in the pseudoscalar current. For the same reason, muon capture on the proton is uniquely placed to determine g_P due to the momentum transfer of nearly the whole muon mass energy. The MuCap collaboration measured the rate of reaction (2.14) from the singlet state of the μ^-p atom, determining the pseudoscalar form factor [13]

$$g_P(q^2 = -0.88m_\mu^2) = 8.06 \pm 0.55. \quad (2.20)$$

With this measurement, all of the single-nucleon form factors are determined precisely enough that they are not the dominant source of uncertainty in multi-nucleon interactions.

2.3 *Physics of Light Nuclei*

The deuteron is the simplest nucleus heavier than a proton, consisting of a bound state of a proton and a neutron. In addition to the one-body parameterization of the weak nuclear current in Eq. 2.16, there are two-body currents involving both nucleons. Likewise, for the three-body systems of the triton and ^3He , two- and three-body currents contribute to observables.

Historically, observables in two- and three-nucleon systems, including the muon capture rate on the deuteron, have been successfully calculated using a phenomenological approach, sometimes called the standard nuclear physics approach (SNPA). Wavefunctions are derived from potentials that have been developed to match existing nuclear observables and reproduce scattering data. There are several of these potentials to choose from, reflecting a lack of constraints on the short-range behavior, but also highlighting that this method is explicitly model-dependent. Transition operators are a combination of the one-body nuclear current of Eq. 2.16 with two- and three-body terms derived from meson exchange currents (MECs), such as single-pion or two-pion exchange.

There are a number of deficiencies to this approach. The choice of potentials and method of accounting for MECs comprise a significant model-dependence. It is difficult to systematically calculate the largest-contributing terms, since there is no explicit expansion and

power-counting hierarchy. A perturbative theory such as high-energy QCD, for instance, relies on the next order of diagrams to constrain the magnitude of error in a calculation. Perhaps most distressing is that the calculations are not based in fundamental QCD, which is known to be the underlying description of the quarks and gluons that constitute nucleons. Modern calculations of nuclear observables use low-energy effective field theories (EFTs) to avoid these deficiencies. I will briefly discuss the chiral symmetry of the QCD Lagrangian and introduce the EFT known as heavy-baryon chiral perturbation theory as well as an elegant alternative called pionless-EFT that is applicable only at very low energy.

2.4 Chiral Symmetry

The QCD Lagrangian can be written in terms of right- and left-handed chirality states,

$$\mathcal{L} = \bar{q}_L i \not{D} q_L + \bar{q}_R i \not{D} q_R - m (\bar{q}_L q_R + \bar{q}_R q_L) - \frac{1}{4} G_{\mu\nu}^a G^{a\mu\nu}. \quad (2.21)$$

Here q is the doublet of up and down quark states,

$$q = \begin{pmatrix} u \\ d \end{pmatrix}, \quad (2.22)$$

and q_L and q_R are the left and right handed projections of these fields. The gluon field strength tensor contains the gluon dynamics and can be written in terms of the gluon fields \mathcal{A} and the structure constants of SU(3), f^{abc} ,

$$G_{\mu\nu}^a = \partial_\mu \mathcal{A}_\nu - \partial_\nu \mathcal{A}_\mu + g f^{abc} \mathcal{A}_\mu^b \mathcal{A}_\nu^c. \quad (2.23)$$

In the limit of massless quarks, the Lagrangian has a ‘‘flavor chiral symmetry’’ $U(2)_L \times U(2)_R$ in the rotation of q_L or q_R by a unitary matrix. This decomposes into the symmetry

$$SU(2)_V \times SU(2)_A \times U(1)_V \times U(1)_A \quad (2.24)$$

where $U(1)_V$ corresponds to the conservation of baryon number, and $U(1)_A$ is broken by the chiral anomaly. The remaining $SU(2)_V \times SU(2)_A$ symmetry is spontaneously broken to the $SU(2)$ isospin symmetry group. This spontaneously broken symmetry would result in massless, scalar Goldstone bosons corresponding to the three broken generators. However,

chiral symmetry is explicitly broken by the fact that the light quark masses are small but not zero, so the pions are in fact pseudo-Nambu-Goldstone bosons with non-zero masses that are still relatively light when compared with other hadrons.

If the strange quark is included in the chiral symmetry in addition to the u and d quarks, the arguments stated above are similar. The $SU(2)$ symmetries are replaced with $SU(3)$ and there are $N^2 - 1 = 8$ broken generators, corresponding to the octet of light pseudoscalar mesons of pions, kaons, and the η [21].

2.5 Chiral Perturbation Theory

An effective field theory (EFT) is a modification of a fully renormalizable quantum field theory where the high energy (short-range) behavior has been “integrated out”, for some chosen energy scale, Λ . This can be interpreted as replacing the non-local interactions governed by high-mass particle exchange with contact interactions of the appropriate strength [22]. A perturbative theory is generated, where processes are expanded in powers of the energy scale relative to the high-energy scale of the theory, Q/Λ . The non-local corrections to the effective contact interactions show up at higher powers of Q/Λ . The interactions representing unknown behavior at high energies are parameterized by a finite (but potentially large) number of low-energy constants (LECs) which are empirical inputs to the theory. The well-known pedagogical example of an effective theory is Fermi’s 4-fermion contact interaction describing weak decays with the effective coupling G_F .

Chiral perturbation theory (χ PT) is the effective field theory of hadronic interactions with energies small compared to the chiral symmetry breaking scale, $\Lambda_\chi \lesssim 1$ GeV [23]. The nucleon mass is not small compared to Λ_χ , so χ PT is useful primarily for calculating processes involving pions and light mesons. For processes involving nucleons, the baryons are treated non-relativistically, with corrections appearing at higher powers of Q/m_B . This extension is called heavy-baryon chiral perturbation theory (HB χ PT), and is one of the primary frameworks for calculating strong and weak observables in light nuclei [24][25].

Chiral perturbation theory is a systematic procedure for classifying terms into different orders of a perturbative parameter, and introduces LECs as parameters that must be fit to experiment or calculated using first-principles methods like lattice QCD. The claim that

HB χ PT is model-independent, in contrast to the SNPA, is essentially stating that HB χ PT is based on the fundamental symmetries of QCD and admits a systematic power expansion and order classification for diagrams in the calculation. This is not trivial since, in a calculation to a given order, the next higher order is used to estimate errors, and interactions parameterized by unknown LECs appear at higher orders. For multi-nucleon systems, however, there is a substantial advantage that LECs from few-nucleon interactions explicitly appear, allowing knowledge about light nuclei to contribute to calculations of heavier nuclei.

The 2-nucleon system contains a LEC called d_R that fixes the strength of the weak axial current interacting with a 2-nucleon contact vertex. This diagram, shown in Fig. 2.3, contributes to muon capture on the deuteron as well as tritium β -decay. Before the rest of the chapter in which I discuss the history and strategy of determining this LEC, I will discuss two alternatives to HB χ PT. The first is a hybrid effective field theory that uses calculation of tritium β -decay to fix the LEC. The other is pionless effective field theory, where the LEC is instead called L_{1A} , but still must be fixed experimentally.

2.5.1 Hybrid χ PT

Deriving the nuclear wavefunctions within HB χ PT to the same precision as the empirical wavefunctions is quite challenging, so a hybrid approach was used for a number of years. Transition operators are derived using HB χ PT and applied between phenomenological initial- and final-state wavefunctions. Since the HB χ PT nuclear potentials are not used, the only unknown LEC in this procedure is d_R , even in three-nucleon processes like muon capture on ^3He and triton β -decay, which have been measured precisely [11][26][27]. In a fully-consistent HB χ PT calculation at NNLO, the three-nucleon potential includes two LECs, as discussed later in Sect. 2.7, and therefore two observables are required to make the calculation predictive.

2.5.2 Pionless Effective Field Theory

The standard HB χ PT expansion is complicated by the existence of additional length scales, especially the large S-wave NN scattering length. Kaplan, Savage, and Wise (KSW) argue

that this indicates that the short-range physics is non-perturbative and “finely-tuned” [28]. In particular, the original power counting scheme proposed by Weinberg leads to graphs at some orders being substantially cancelled by higher-order graphs, implying that the strength of contributions does not necessarily become smaller at higher orders. In response, KSW developed a different expansion to deal with this problem when applied to low-energy nuclear interactions [8]. In this theory, only nucleon fields are considered (no pion fields), so all interactions are contact interactions. The expansion scheme uses powers of q/m_π , the momentum divided by the pion mass, to perturbatively calculate two-nucleon interactions and electroweak currents. Similar to the parameter d_R in HB χ PT, the high-energy behavior of the weak axial current interaction with two nucleons up to NNLO involves a single LEC, called L_{1A} .

Expansion in the small parameter q/m_π may not seem justified, since the reaction (1.2) has a 3-body final state, with neutron energies potentially as large as the muon mass, $m_\mu \approx 0.75m_\pi$. Pionless-EFT is certainly not applicable in this regime, but the high neutron energy ($\gtrsim 30$ MeV) region of phase space contributes a relatively small amount (~ 10 %) to the total capture rate. To complete the theoretical calculation at high energies, one could match the results of HB χ PT at lower energies to the pionless-EFT result and extend the partial capture rate $d\Lambda_d(E)/dE$ to high energy. While this would not be a “pure” pionless-EFT result, any uncertainty from HB χ PT would be suppressed by its limited domain.

Alternatively, an experimental approach would be to measure the partial capture rate at energies where pionless-EFT is applicable. This could be accomplished by analyzing only neutrons with energies below 30 MeV, for instance, which would not represent a significant loss of statistics. However, this requires counting neutrons, rather than the lifetime method (described in the next chapter) used by MuSun that detects only the decay electrons.

Chen et al. calculated the μd capture rate in terms of L_{1A} up to NNLO in pionless EFT [29]. They report that a precision of ± 2 fm³ is possible for a 2 % measurement of Λ_d . They do not predict a central value because the three-nucleon system had not been treated consistently with pionless-EFT, as mentioned above. There is recent progress in this area, however, with an effort to calculate three-nucleon observables, including the tritium β -decay rate, in pionless-EFT [30].

2.6 Few-Nucleon Observables

It is desirable to determine 2-nucleon LECs using observables of two nucleon weak interactions or from lattice QCD. One reason is that the 3-nucleon system is considerably more difficult to calculate to the same order as the 2-nucleon potentials and transition operators. Moreover, as the precision of calculations rises to further orders or as more nucleons are considered, additional LECs are introduced, parameterizing different forms of short length scale physics. To maintain a predictive theory, more clean experimental observables will be needed at higher orders. For these reasons, one of the two-nucleon weak interactions, (1.2) or

$$p + p \rightarrow d + e^+ + \nu_e \quad (2.25a)$$

$$\nu_e + d \rightarrow p + p + e^- \quad (2.25b)$$

should be measured precisely to fix the value of the LEC d_R or, equivalently, L_{1A} in pionless effective field theory.

All three of these processes share the same weak hadronic vertex, as illustrated in Fig. 2.1, and the contribution of the two-body weak axial current is the dominant uncertainty in each case. Several efforts have been made to determine the weak axial LEC d_R using these two-nucleon interactions and tritium beta-decay, $t \rightarrow {}^3\text{He} + e^- + \bar{\nu}_e$, summarized in Table 2.1. These are reported in terms of L_{1A} , but in principle could be converted to a measurement of d_R using a calculation within HB χ PT.

Solar pp -fusion is suppressed by the low probability of tunneling through the large Coulomb barrier between protons. Reproducing such a fusion process on Earth is infeasible with current technology. Through the standard solar model and helioseismology, the pp -fusion cross-section could be constrained in principle, but this effort has proven to be difficult [32]. Neutrino break up of the deuteron is limited in statistics because of the low probability of neutrino interaction with even massive detectors like the heavy water detector at the Sudbury Neutrino Observatory (SNO). Nonetheless, a low-precision determination of the LEC for the two-body weak axial current has been performed in the context of pionless EFT. In SNO and Super-Kamiokande, charged-current and neutral-current channels

Process	L_{1A} (fm ³)
$\nu - d$ NC, CC, ES	4.0 ± 6.3
Reactor $\bar{\nu} - d$ breakup	4.0 ± 6.3
Tritium β -decay	4.2 ± 0.1

Table 2.1: Summary of experimental extractions of the strength of two-body weak axial current coupling to two nucleons. Note that tritium β -decay cannot be calculated within pionless-EFT, and a bridge via HB χ PT is necessary. See text for details. Table reproduced from Ref. [31] and references within.

are sensitive to L_{1A} , while the elastic electron scattering channel is used for calibration. Chen et. al. showed that L_{1A} can be self-calibrated using the neutrino data, obtaining $L_{1A} = 4.0 \pm 4.7(\text{stat}) \pm 4.5(\text{sys}) \text{ fm}^3$ [31]. This uncertainty is large, but does not change the conclusions about neutrino flavor change.

Muon capture differs from the pp -fusion and neutrino reactions in that the momentum transfer is much larger, comparable to the muon mass, and experimentally, the process is much easier to produce terrestrially with high statistics. Experimental design will be considered in detail in the next chapter, but the primary advantages are that negative muons can be created in large numbers at pion production facilities and that muons form atoms, enhancing the wavefunction overlap with the nucleus.

Although it is beyond the scope of this thesis, the first lattice calculation of the two-nucleon axial current matrix element was performed recently at a pion mass of 805 MeV [33].

2.7 Calculation of Two- and Three-Nucleon Observables in HB χ PT

One of the attractive features of HB χ PT is the emergence of a hierarchy of nuclear forces. That is, with an appropriate power-counting scheme, the lowest-order diagrams are interactions between two nucleons, and interactions between three or more nucleons appear at successively higher orders of Q/Λ_χ (Fig. 2.2). This explains the empirical fact that three-nucleon forces are weaker than two-nucleon forces and in general, N-body interactions are




Figure 2.1: All three processes (2.25a), (2.25b), and (1.2) involve the same weak hadronic vertex. A high-precision measurement of any of these fixes the LEC in effective field theories that can be used to predict the other two processes.

weaker as N rises. For instance, it is expected that the 4-nucleon contribution to the helium-4 nuclear binding energy ($\lesssim 1$ MeV) is small compared to the total ($BE_4 = 28.3$ MeV) [34]. In this section, the calculation of the nuclear potentials and the weak axial and vector currents using HB χ PT is summarized following closely the work of Ref. [34] on nuclear potentials and Ref. [7] on weak observables.

The state of the art calculation of the two- and three-nucleon (NN and NNN) potentials includes terms up to order $(Q/\Lambda_\chi)^4$. The leading order (LO) at $(Q/\Lambda_\chi)^0$ contains a contact term with two couplings C_S and C_T and a one-pion exchange (OPE) diagram, which are determined through a fit to multiple two-nucleon observables. The pion exchange is necessary to explain some interesting phenomena such as large angular momentum partial wave contributions to NN scattering and the quadrupole moment of the deuteron [35].

As a consequence of time-reversal and parity invariance, there are no interactions at $(Q/\Lambda_\chi)^1$, and the next-to-leading order (NLO) is $(Q/\Lambda_\chi)^2$, where two-pion exchange (TPE) diagrams enter (see Fig. 2.2). These diagrams involve three LECs, c_1 , c_3 , and c_4 , parameterizing the NN π vertex. Two-pion exchange necessarily involves loop diagrams and the consequent divergent integrals are handled with dimensional regularization and a momentum cutoff function. Marcucci et. al. use the cutoff values $\Lambda = 500$ MeV and 600 MeV [7].

At NNLO, $(Q/\Lambda_\chi)^3$, additional TPE diagrams contribute to the NN potential, and the

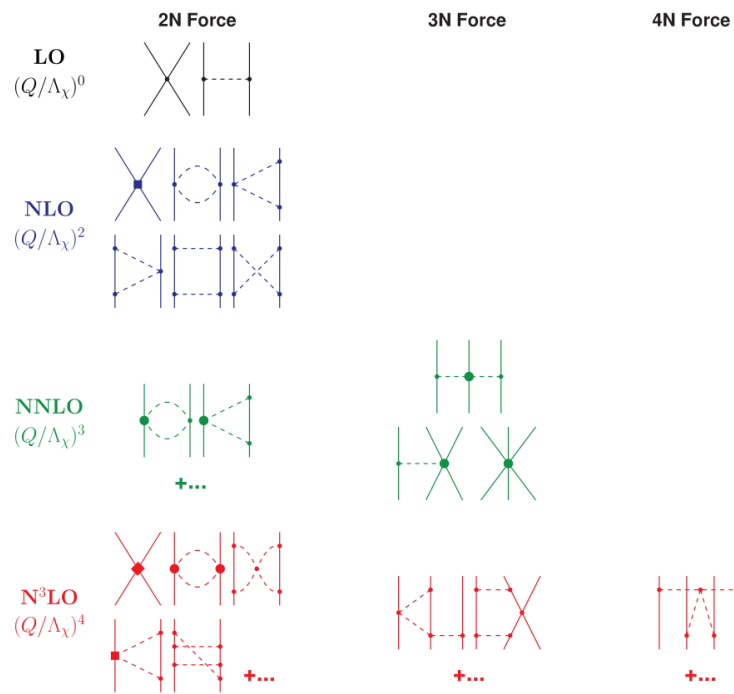


Figure 2.2: Diagrams contributing to the NN and NNN potentials at various orders in the chiral expansion. Figure from Ref. [34].

first NNN interactions arise. There are 3 diagrams that contribute to the NNN force, shown in Fig. 2.2: a three-nucleon contact term parameterized by LEC c_E , a TPE diagram that contains no additional parameters beyond the NN TPE LECs, and a pion exchange with a two-nucleon contact vertex parameterized by LEC c_D . This last diagram is especially important because it is closely related to the weak axial current term parameterized by the LEC d_R measured by the MuSun experiment (see Fig. 2.3):

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

This relation between the weak axial LEC and a purely nuclear diagram is interesting and also helpful in that only one independent measurement determines both.

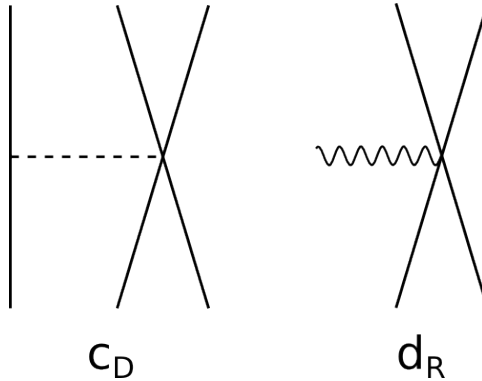


Figure 2.3: The diagrams that involve the LECs c_D (left) in the three-nucleon potential and d_R (right) in the two-nucleon weak axial coupling. Both c_D and d_R parameterize the axial coupling to a two-nucleon contact vertex and are related through Eq. 2.26. In three-nucleon weak interactions like tritium β -decay, the diagram on the right appears with an additional spectator nucleon.

Several LECs have been introduced that must be determined empirically. The c_i constants parameterize TPE in the dimension-two chiral Lagrangian and the d_i for TPE in the dimension-three Lagrangian. These and LO contact couplings C_S and C_T are all fit to pion-nucleon scattering experiments [36][37] and NN observables. The LECs associated with the NNN potential, c_D and c_E must be fixed with two observables involving at least

3 nucleons. For some time, the triton binding energy and the $n - d$ scattering length or the ${}^4\text{He}$ binding energy were used, but due to the substantial correlation between these, the uncertainty in individual LECs was large. A suggestion by Gardsteig and Phillips to fix the LEC c_D with the β -decay rate of the triton was carried out by Gazit et al. [38]. The β -decay rate, Γ_β^t , is strongly affected by the weak coupling involving d_R , but relatively insensitive to the contributions to the nuclear potentials from c_D and c_E . This revelation allows a much more precise determination of d_R , which also fixes c_D through Eq. 2.26 and then c_E with the triton binding energy. In practice, allowed trajectories for a given cutoff energy are created in the $c_D - c_E$ plane to reproduce the ${}^3\text{He}$ or triton binding energies, then Γ_β^t is calculated along this trajectory and then compared to the experimental value. The insensitivity of Γ_β^t to the nuclear multi-body potential terms also explains the success of the hybrid χPT approach to measuring d_R , since three-nucleon forces are not so important, and the dominant contribution is through the weak axial current that is fully treated with a systematic EFT.

2.8 History of Λ_d Calculations

Calculations of the rate of μd capture span decades. Table 2.2 and Fig. 1.1 summarize results from several groups. It is only relatively recently that the calculation has been performed fully within the context of HB χPT theory. The theoretical precision, even if inflated to account for the small difference between different groups, greatly exceeds the existing experimental precision. However, the discrepancy between results [39] and [7] is comparable to the target precision for MuSun.

In Ref. [7], the μd capture rate is broken down into the partial rate for the largest-contribution partial-wave components, and additionally the impulse approximation (IA), which ignores two-body contributions, is calculated separately. Just over half of the total capture rate proceeds through the S-wave channel, with the remainder being largely P-wave. The two-body contribution of 18 s^{-1} is largely S-wave.

Λ_d (s^{-1})	Method	Year	Reference
386 ± 4	hybrid-EFT	2002	[40]
± 10	pionless-EFT	2005	[29]
416 ± 6	Phenom+MEC	2010	[41]
392.0 ± 2.3	hybrid-EFT	2011	[42]
399 ± 3	HB χ PT	2012	[7]
$383.8 - 392.4$	HB χ PT	2012	[39]
409 ± 40	experiment	1989	[43]
470 ± 29	experiment	1986	[44]
± 6	MuSun target		

Table 2.2: Theoretical calculations of Λ_d . Phenom+MEC means phenomenological nuclear potentials are used with transition operators derived from meson-exchanged currents (sometimes called SNPA). Hybrid-EFT means the transition operators are derived in the context of χ PT, but the nuclear wavefunction is determined from phenomenological potentials. The EFT-based calculations use the precise measurement of tritium β -decay to fix the 2-body LEC and produce a numerical result. In the case of pionless-EFT, the LEC is left unfixed, and only the estimated error from the next order of calculation is reported.

2.9 Summary

The theoretical treatment of the physics of small nuclei is undergoing a shift from phenomenological modeling of nucleon-nucleon potentials to formal EFTs based in the underlying symmetries of QCD. Experimentally measured NN potentials and meson-exchange currents are giving way to systematic expansion in powers of the momentum scale over a higher-energy scale. Physics beyond this scale is encapsulated in the coupling strength of specific contact diagrams, parameterized by LECs. The systematic expansion justifies the assumption that N-body interactions are weaker as N rises, and also makes clear the relation between different nuclear reactions, with contact diagrams and their LECs being shared

	1S_0	3P_0	3P_1	3P_2	1D_2	3F_2	Λ_d
IA($\Lambda = 500$ MeV)	238.8	21.1	44.0	72.4	4.5	0.9	381.7
IA($\Lambda = 600$ MeV)	238.7	20.9	43.8	72.0	4.5	0.9	380.8
FULL($\Lambda = 500$ MeV)	254.4 ± 0.9	20.5	46.8	72.1	4.5	0.9	399.2 ± 0.9
FULL($\Lambda = 600$ MeV)	255.2 ± 1.0	20.3	46.6	71.6	4.5	0.9	399.1 ± 1.0

Table 2.3: Partial μd capture rate (in s^{-1}), broken down by partial-wave contributions and momentum-scale cutoff Λ . The impulse approximation, including only one-body interactions, is reported separately from the full EFT calculation that includes two-body effects. Table reproduced from Marcucci et. al. [7].

among several processes. All of the LECs must be determined by independent observables, so any independent measure of LECs contributes to precision and predictivity.

There is a lack of a high-precision 2-nucleon experimental result to fix the two-body weak axial current LEC. Instead, we must rely on the 3-nucleon sector, where the theoretical extraction of this LEC is complicated by the less-well-known 3-body contributions. The MuSun experiment will measure this LEC using the 2-nucleon process of muon capture, providing a theoretically clean input to higher-order calculations and larger nuclei.

Chapter 3

EXPERIMENTAL DESIGN

The discovery of muon-catalyzed fusion (MCF) of the $p\mu d$ molecule in 1956 by Luis W. Alvarez et al. ignited a decades-long study of the kinetic properties of muonic atoms of the hydrogen isotopes [45]. In parallel, measurements of muon capture on these nuclei probed the properties of the proton and the two-nucleon interaction in the deuteron. Early measurements were made before the hyperfine structure of the μd atom was well-understood. The two neutrons in the final state have a spectrum of energies, and fusion reactions source an irreducible background of neutrons that is difficult to quantify experimentally.

This chapter discusses the technique used in MuSun to measure Λ_d by comparing the lifetime of negative muons in deuterium to the free muon decay lifetime. The complex atomic and molecular kinetics of muons in deuterium are described along with a numerical model of the MuSun gas conditions. The optimal choice of temperature and density of the deuterium target gas is explained in light of competing experimental requirements. Finally, I discuss three experiments that have measured Λ_d in the context of the muon atomic and molecular kinetics described in this chapter.

3.1 Lifetime Method

Previous experiments measured the capture rate by comparing the number of observed final-state neutrons to the number of incident muons. This relies on determining the absolute efficiency for neutron detection, which limits the precision of these experiments. Additionally, there are significant backgrounds due to bremsstrahlung photons and neutrons from muon-catalyzed fusion. Capture neutrons are continuously distributed with a peak at 1.5 MeV, and the monoenergetic 2.45 MeV fusion neutrons produce an observed spectrum with considerable overlap. In previous experiments, gaseous and liquid targets are composed of various concentrations of protium and deuterium gas to balance the population of muons

capturing on deuterium with the rate of fusion.

MuSun employs the lifetime method, measuring the disappearance rate of muons via the decay channel (2.4). The capture rate on the deuteron is a small fraction of the free decay rate, so most stopped muons produce observable electrons. The total disappearance rate is equal to the sum of the capture and decay rates, up to small distortions from the atomic and molecular kinetics (discussed later in Sect. 6.4):

$$\lambda_{\mu^-} = \lambda_0 + \Lambda_d. \quad (3.1)$$

The free muon decay rate, $\lambda_0 \equiv \lambda_{\mu^+}$, given in (2.8), has been measured to part-per-million precision, so this contributes negligible additional uncertainty to the extraction of Λ_d .

The primary advantage of this method over measuring the final state neutrons is that the charged decay electrons are easy to detect. The measured disappearance rate does not depend on the absolute electron detection efficiency, but the disadvantage of the lifetime method is that significantly more observed decays are required. The number of neutrons observed from a population of N μd atoms follows a binomial distribution $B(N, p)$ with

$$p = \frac{\Lambda_d}{\lambda_0 + \Lambda_d}. \quad (3.2)$$

This is well-approximated by a normal distribution with mean Np and variance $Np(1-p)$, and achieving a 1 % uncertainty requires 10^4 observed capture events. Since $p \approx 10^{-3}$, the number of stopped muons is 10^8 , assuming 10 % acceptance of the neutron detectors. However, the muon disappearance rate must be measured to 10 parts-per-million, requiring the observation of 10^{10} muon decays.

3.2 Deuterium Target and CryoTPC

Muons are produced at the PiE3 beamline at the Paul Scherrer Institute (PSI) via charged pion decay

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu. \quad (3.3)$$

The beamline is tuned to select muons with a momentum of ~ 40 MeV/c, corresponding to a kinetic energy ~ 5 MeV. Muons pass through several detectors and beam windows before entering a gaseous deuterium target, where they lose the remainder of their energy




Figure 3.1: Drawing of the MuSun CryoTPC inside the deuterium pressure vessel. The muon beam enters through a hemispherical beryllium window and muons stop in the central volume of deuterium, ionizing the gas as they stop. Ionization electrons drift toward the anode plane, and the 3-D stop position is reconstructed from the anode pad grid and the time it takes for the charge to drift. Note that this drawing inaccurately displays more pads than the 6×8 pad plane of the real detector.

via ionization of the gas as they come to rest. The combination of range straggling and the width of the beam momentum selection results in a stopping distribution with full-width at half-maximum of 10 cm.

The target is instrumented with a time-projection chamber (TPC) (Fig. 3.1) that ensures a stop in the deuterium gas by reconstructing the three-dimensional muon stop position. Chapter 4 describes the experimental apparatus and detector systems in detail, but a short introduction to the TPC is provided here since it is central to achieving the precision goal of MuSun. The TPC is an ionization chamber operated in pure deuterium with an active

volume measuring $96 \times 71 \times 120 \text{ mm}^3$. A uniform electric field is established between the cathode plane at the top of the detector and anode pad plane at the bottom by biasing the cathode to -80 kV . The electrons from the muon track ionization drift toward the anode plane, where they are collected. The induced signals on a 6×8 segmented grid of anode pads provide the x- and z-coordinates of the stop, and the y-coordinate is derived from the time taken for the charge to drift. The energy of particles in the TPC is determined by the amplitude of the charge signal collected on the anode pads, although very dense energy deposition allows electron-ion pair recombination before the electrons can drift to the anode, reducing the observed energy.

3.3 Muon Kinetics

3.3.1 Formation of μd Atoms

A negative muon forms a muonic atom with a deuterium nucleus, which is analogous to the electronic version, but the larger muon mass ($m_\mu/m_e = 207$) results in a similarly larger reduced mass, smaller Bohr radius, and larger ground state binding energy. Muonic atoms form in one of two possible hyperfine states, with molecular formation at different rates from each and subsequent muon-catalyzed fusion from the molecular state. A summary of this complex behavior is diagrammed in Fig. 3.2 and the kinetic rates are given in Table 3.2.

Stopped muons initially displace an electron, forming a highly excited state ($\langle n \rangle = 14 \approx (m_\mu/m_e)^{1/2}$) that overlaps the electronic ground state wavefunction. The muon rapidly falls into the 1S state following a cascade of processes including radiative decay, Auger emission, and Coulombic de-excitation through interactions with nearby atoms. Any residual average polarization coming from the muon beam is expected to be substantially de-polarized by the cascade, but in the worst-case estimations could contribute a $\sim 5 \text{ s}^{-1}$ shift to the doublet capture rate (see Ref. [46], Appendix A, and references therein).

3.3.2 Hyperfine Structure

The 1S state of the μd atom has two hyperfine levels, a lower $F = 1/2$ doublet state and an upper $F = 3/2$ quartet state, with energy separation $\Delta_{hf} = 0.048 \text{ eV}$. Atomic muon

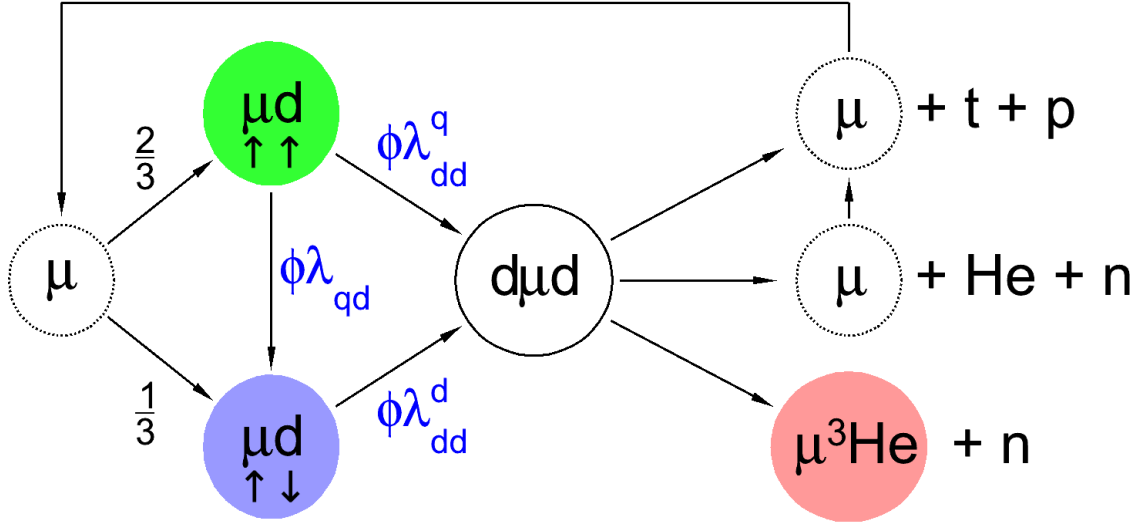


Figure 3.2: Atomic and molecular kinetics following a muon stop in deuterium gas. The muon is recycled in two of the three fusion branches, reforming a μd atom, and forms a bound state, or “sticks” with the ${}^3\text{He}$ nucleus in the third. Muon decay occurs from all states with rate $\lambda_0 = 455170 \text{ s}^{-1}$. MuSun aims to measure Λ_d , the rate of nuclear capture from the doublet hyperfine state.

capture and the subsequent cascade are processes with energy scales much larger than the hyperfine splitting, so the two states are populated according to their multiplicities, with $\nu = 2/3$ of μd atoms initially in the quartet state.

The nuclear capture rate from the quartet state, Λ_q is highly suppressed compared to the doublet capture rate, Λ_d . In order to clearly interpret the nuclear capture rate measured by MuSun it is necessary to precisely understand the time-dependent populations of μd atoms in the two hyperfine states. MuSun uses a high-density gaseous target to increase the hyperfine transition rate, $\Lambda_{qd} = \phi\lambda_{qd}$, and ensure that the muons spend most of their time in the doublet atomic state. At experimental conditions ($\phi = 0.06$), this rate is $\Lambda_{qd} = 2.2 \times 10^{-6} \text{ s}^{-1}$, which is only a few times larger than the free muon decay rate.

3.3.3 Molecular Formation and Muon-Catalyzed Fusion

A muonic deuterium atom can interact with a D_2 molecule to form a $d\mu d$ molecule, compensating for the change in binding energy by ejecting an electron [12]



The deuterons bound in this muonic molecule are a factor of $(m_\mu/m_e) \approx 200$ closer to each other than in the electronic ion, D_2^+ . This proximity of ~ 500 fm substantially reduces the Coulomb energy barrier, and muon-catalyzed fusion can proceed rapidly via one of the three fusion channels



In the first two reactions, the muon remains in the gas, and rapidly forms another μd atom, starting the kinetic evolution again. In the third, the muon “sticks”, forming a bound state with the ${}^3\text{He}$ nucleus, eventually disappearing via muon decay or nuclear capture on ${}^3\text{He}$. The bound $\mu^3\text{He}$ behaves in some ways like a singly-charged particle due to the screening by the muon, which is much more tightly bound than electronic atoms. In particular, the ionization density of the $\mu^3\text{He}$ ion in the MuSun target gas is lower than that for the bare ${}^3\text{He}$, which reduces the recombination factor and increases the observed energy. (The spectrum of these two particles can be found later in Fig. 5.4.) The fraction of $n+{}^3\text{He}$ fusion reactions in which the muon sticks to the helium nucleus, $\bar{\omega} = 0.1206(6)$, is a precisely measured quantity due to its role in limiting the energy yield per muon from muon-catalyzed fusion. Especially in deuterium-tritium mixtures, the analogous sticking probability limits the potential of this process for energy generation.

Resonant Molecular Formation

In the late 1960s through the 1970s, a high-rate resonant $d\mu d$ molecular formation process was discovered and studied extensively through the 1990s. Instead of the process in

	E(MeV)	E_{obs} (MeV)	R(mm)
${}^3\text{He}$	0.82	0.32	0.3
n	2.45	0-2.5	(long)
$\mu^3\text{He}$	0.80	0.50	0.6
t	1.01	~ 0.8	0.9
p	3.02	~ 3	12.9

Table 3.1: Properties of fusion products in the MuSun gas target and TPC detector at operating temperature and density. Neutrons interact rarely in the TPC, but can deposit substantial energy by scattering from a deuterium nucleus in the TPC volume. The observed particle energies are reduced due to recombination of electrons and ions for high dE/dx particles.




Figure 3.3: (Left) Molecular formation rates from the quartet (λ_q) and doublet (λ_d) hyperfine states. At the operating temperature of 34 K, the quartet formation rate is dominant. (Right) Ground-state μd hyperfine transition rate, λ_{qd} . Figure reproduced from Ref. [12] with modifications.

reaction 3.4, the $d\mu d$ molecule is formed in a weakly-bound state that is rotationally and vibrationally excited ($J = 1, \nu = 1$) and the binding energy of this state is transferred to rotational and vibrational excitations of the whole molecular complex,



In contrast to the Auger process, the energy levels in the final state are discrete, so the μd kinetic energy, and therefore the gas temperature, greatly impacts the molecular formation rate. That is, this is a resonant process which relies on the condition that the μd kinetic energy distribution significantly overlaps the formation energy. At low temperatures, the molecular formation rate from the higher-energy quartet state dominates, as the resonance condition is more easily fulfilled. At higher temperatures the kinetic energy of the gas molecules satisfies the resonance condition and the hyperfine effects become less important [12].

3.3.4 Target Gas Conditions

The target gas conditions for MuSun are selected to fulfill the two goals of quickly depopulating the quartet hyperfine state, and minimizing the total number of muon-catalyzed fusion reactions. The gas density directly affects the stopping power of the target for beam muons. The specific hyperfine transition rate and molecular formation rate from the quartet state do not depend strongly on temperature, but the molecular formation rate from the doublet state falls by nearly two orders of magnitude from 300 K to 30 K (see Fig. 3.3). Additionally, the cryogenic temperature allows for a much higher gas density given the mechanical constraints on a pressure vessel with a low-mass beam window.

At a temperature of 34 K and pressure of 5.6 bar, deuterium is gaseous with a molecular number density of 2.25 mol L^{-1} . Typically, the density is reported as a fraction of the density of liquid H_2 at its boiling point at 1 bar, 20.32 K, which is $35.170 \text{ mol L}^{-1}$. With this normalization, the experimental gas density is $\phi = 6.40 \%$.

3.3.5 Numerical Analysis

The muon population in the two hyperfine states of muonic deuterium, $N_d(t)$ and $N_q(t)$, and $\mu^3\text{He}$, $N_{\text{He}}(t)$, evolve over time due to the kinetic processes described above. A numerical solution for these populations over time is constructed using the evolution matrix

$$M = \begin{pmatrix} -\lambda_0 - \Lambda_d - \phi\lambda_{qd} - \phi\lambda_q(1 - \rho_q) & \phi\lambda_{dq} + \phi\lambda_d\rho_q & 0 \\ \phi\lambda_{qd} + \phi\lambda_q\rho_d & -\lambda_0 - \Lambda_d - \phi\lambda_{dq} - \phi\lambda_d(1 - \rho_d) & 0 \\ \phi\lambda_q\omega & \phi\lambda_d\omega & -\lambda_0 - \Lambda_{\text{He}} \end{pmatrix} \quad (3.7)$$

where the muon state populations are given by the vector

$$\bar{N}(t) = \begin{pmatrix} N_d(t) \\ N_q(t) \\ N_{\text{He}}(t) \end{pmatrix} \quad (3.8)$$

with an initial state defined by statistical population of the deuterium hyperfine levels ($\nu = 2/3$ is the fraction initially in the quartet state)

$$\bar{N}(t=0) = \begin{pmatrix} 1 - \nu \\ \nu \\ 0 \end{pmatrix}. \quad (3.9)$$

These satisfy the equation

$$\frac{d\bar{N}(t)}{dt} = M\bar{N}(t). \quad (3.10)$$

With the parameters in Table 3.2, and a target gas density of 6.4 % of LHD, the solutions are shown in Fig. 3.4a [47]. The $F = 3/2$ hyperfine state depopulation occurs faster than the muon disappearance rate, but the time the muon spends in this state is not negligible. The residual quartet population at late time is defined by the balance between hyperfine depopulation and recycling into the quartet state after fusion

$$\lambda_{qd}N_q(t \rightarrow \infty) = \lambda_d\rho_qN_d(t \rightarrow \infty). \quad (3.11)$$

Here the left side is the rate of collisional hyperfine depopulation and the right side is the rate of MCF reactions from the doublet state that end up with a μd atom in the quartet

state. Rearranging, the late-time ratio of the quartet and doublet populations is

$$\left[\frac{N_q}{N_d}\right]_{t \rightarrow \infty} = \frac{\lambda_d \rho_q}{\lambda_{qd}}. \quad (3.12)$$

With the parameters in Table 3.2, this ratio is $\sim 10^{-3}$.

In Fig. 3.4, two example kinetic solutions show the effect of low gas density (left figure) and high temperature (right figure). The quartet state depopulates much slower at lower density, while at higher temperature, the significantly larger molecular formation rate from the doublet state results in more bound $\mu^3\text{He}$ atoms. The error contribution from uncertainties in other muon kinetic parameters (Table 3.2) is shown in Fig. 3.5.

Fusion Fraction

The fraction of events, ϵ_f , in which the muon undergoes at least one p+t fusion reaction is an important quantity for the systematic effect discussed in Chapter 7. We can estimate the ratio of rates of fusion from the quartet state to the rate of μd atoms leaving the quartet state:

$$\begin{aligned} \epsilon_f &\approx \frac{\nu(1-\beta)\lambda_q}{\lambda_{qd} + \lambda_q[1 - \nu\beta(1-\omega)] + \lambda_{\mu^-}/\phi} \\ &\approx \frac{0.033}{1 + 0.011/\phi}. \end{aligned} \quad (3.13)$$

This estimate gives $\epsilon_f = 2.8\%$ for the R2011 gas density. The density dependence of the fusion fraction is somewhat weak because there is a balance between the depopulation of the quartet state, from which most fusion reactions occur, and the molecular formation rate, both of which are proportional to density. The sensitivity of ϵ_f to changes in gas density is given by

$$\frac{d\epsilon_f}{\epsilon_f} = 0.15 \times \frac{d\phi}{\phi}. \quad (3.14)$$


The full numerical model can be adjusted to measure the fusion fraction by adding states for p+t fusions that do not recycle the muon. The population of these states is the cumulative fraction of events that undergo the p+t fusion reaction, which asymptotically approaches ϵ_f (See Fig. 3.6). The uncertainty in this quantity due to the kinetic parameters is large. The fusion fraction is required for the systematic correction in Chapter 7, but the

Quantity	Symbol	300 K	34 K
initial quartet fraction	ν	2/3	
hf energy splitting	Δ_{qd} (meV)	48.5	
hf transition q \rightarrow d	λ_{qd} (μs^{-1})	35(5)	37.0(4)
hf transition d \rightarrow q	λ_{dq} (μs^{-1})	$\sim (0.30)\lambda_{qd}$	0*
dd μ form. from q	λ_q (μs^{-1})	$\sim 3.60^\dagger$	3.98(5)
dd μ form. from d	λ_d (μs^{-1})	2.549(23)	0.053(3)
effective fusion fraction	β	0.590(6)	0.517(15)
sticking probability	$\bar{\omega}$	0.1206(6)	
eff. sticking fraction	$\omega = \beta\bar{\omega}$	0.0711(7)	0.0623(19)
recycling fraction into d	$\rho_d = (1 - \nu)(1 - \omega)$	0.310	0.312
recycling fraction into q	$\rho_q = \nu(1 - \omega)$	0.619	0.625
^3He total capture rate	Λ_{He} (s^{-1})	2216(70)	
μd quartet capture rate	Λ_q (s^{-1})	~ 10	
μd quartet capture rate	Λ_d (s^{-1})	~ 400	

Table 3.2: Kinetic parameters of muons in deuterium at 300 K and 34 K [12].

* Detailed balance gives $\lambda_{dq}/\lambda_{qd} = \nu/(1 - \nu) \exp(-\Delta_{qd}/kT)$, which is negligible at 34 K.

† This is a theoretical interpolation of two separate measurements at 125 K and 400 K.



(a) μd hyperfine state populations over time for $T = 30$ K and $P = 5$ bar, which are similar to the experimental conditions of 34 K and 5.6 bar. The relatively high gas density, $\phi = 6.4$ %, increases the hyperfine transfer rate, ensuring muons spend most of their time in the doublet state.

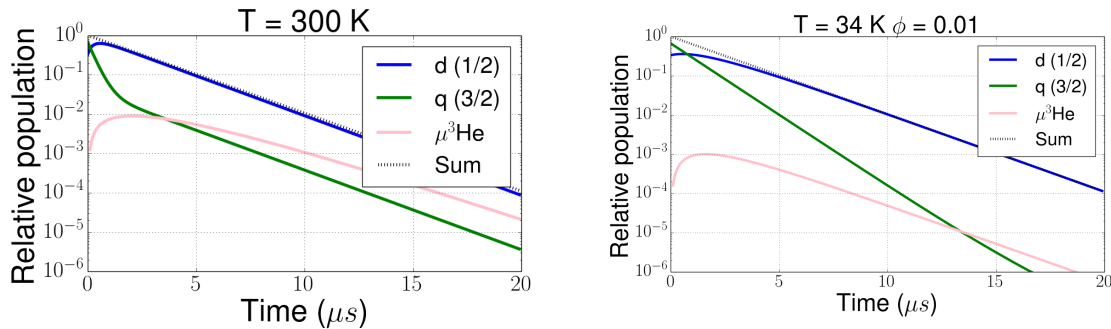


Figure 3.4: Muonic atom populations with the experimental density, but at 300 K (left) and at the experimental temperature, but with a density of $\phi = 0.01$ (right). Higher temperature results in many more fusion reactions, which populates the $\mu^3\text{He}$ state, while reducing the density retains a significant population in the upper hyperfine state.




Figure 3.5: Effect on the measured doublet capture rate, Λ_d , due to $\pm 2\sigma$ variation of the kinetic parameters from Table 3.2. Figure reproduced from Ref. [46].

precision in the kinetic parameters is not sufficient to simply calculate this quantity; it must be measured using MuSun data.

3.3.6 Muon Capture on Nuclei Other than Deuterium

Muonic atoms can form with any nucleus, and the rate of muon capture scales as $\Lambda_Z \propto Z^4$, where Z is the atomic number of the nucleus. Because of the high capture rate, buildup of even small populations of μZ atoms must be avoided to unambiguously interpret the muon disappearance rate. In MuSun, there are two major sources of μZ atoms: direct muon stops in the non-deuterium materials that make up the beam windows, pressure vessels, and the TPC detector, and chemical impurities in the target gas.

Prompt stops in a non-deuterium material contribute an additional exponential component to the decay time distribution with a disappearance rate $\lambda_0 + \Lambda_Z$. The systematic error in the lifetime measurement due to these events is discussed in Sect. 6.2.

A proportion of muon stops will initially form atoms with chemical gas impurities such as O_2 and N_2 , but this fraction is directly proportional to the concentration c_Z , which is




Figure 3.6: Cumulative fraction of events with p+t fusions. There is significant uncertainty in this quantity from the kinetic parameters.

\sim ppb, and therefore this contribution is negligible. However, μd atoms are small neutral systems that easily penetrate the heavier atoms' electron clouds, so the transfer process



has a very large rate Λ_{dZ} due to the high density of the target gas (see Table 3.3). This process leads to a population buildup of μZ atoms that is much larger than the relative impurity concentration. This systematic effect is discussed in Sect. 6.1.

The temperature of MuSun target gas is 34 K, which is well below the freezing points of oxygen (54 K) and nitrogen (63 K), but the vapor pressure is sufficient that a ppb-level impurity of O_2 and several tens of ppb of N_2 would exist if a trace amount of those molecules was frozen inside the deuterium vessel (see Fig. 3.7). To remove chemical impurities, the target gas is continuously circulated through a purification system that is described in Sect. 4.6.

3.3.7 Muon-Catalyzed Fusion with Protium

The deuterium gas must be purified to exclude 1H , which is called *protium* in this text to distinguish it from an isotopic mixture of hydrogen. The isotopic purity requirement for the

Element	$\Lambda_Z(\text{s}^{-1})$	$\lambda_{dZ}(\text{s}^{-1})$	$\Lambda_{dZ}(\text{s}^{-1})$	$\Delta\lambda_{\mu^-}(\text{s}^{-1})$
N	62×10^3	1.45×10^{11}	$9.3 \frac{c_Z}{\text{ppb}}$	$2.0 \frac{c_Z}{\text{ppb}}$
O	102×10^3	0.63×10^{11}	$4.0 \frac{c_Z}{\text{ppb}}$	$1.2 \frac{c_Z}{\text{ppb}}$

Table 3.3: Properties of muon transfer to chemical gas impurities and subsequent muon capture. The capture rate, Λ_Z is much larger than Λ_d . The effective transfer rate at the density of the MuSun target gas from μd to μZ for an impurity concentration c_Z is $\Lambda_{dZ} = \phi\lambda_{dZ}c_Z$. Finally, $\Delta\lambda_{\mu^-}$ is the expected lifetime shift for a given impurity concentration. The experimental results in this table are collected in Ref. [48]. Preliminary analysis of nitrogen-doped deuterium gas in MuSun indicate that the rate shift from to muon transfer to nitrogen is $(3 \text{ s}^{-1}) \frac{c_Z}{\text{ppb}}$.





Figure 3.7: Calculated N₂ and O₂ concentrations based on partial pressure at 34 K and the MuSun gas density.

target gas in MuSun is set by the effect of the radiative and conversion fusion reactions from the $p\mu d$ molecular state:



The less-common reaction (3.16a) ejects the muon with enough energy to reach non-deuterium materials, where the capture rate is much larger than Λ_d .

To eliminate the systematic error from this process, the protium concentration is reduced to a level 100 ppm with a cryogenic distillation column before each operation period. Before the R2011 beam period the gas was measured with gas chromatography to have 10 ppm of protium, and 100 ppm after the run. Although the concentration remained below the required limit, the source of additional protium is unknown.

The Bardin et. al. measurement of the muon capture rate with low protium concentration in a liquid deuterium target determined a correction to the capture rate of 43 s^{-1} for $c_p = 0.17 \%$ [44]. The MuSun gas density is more than a factor of 15 smaller than the gas target in [44] and the concentration of protium is smaller by a factor of 17. The rate shift scaled by these two factors is negligible, but the lower target gas density in MuSun means the 5.5 MeV muons travel farther and are more likely to reach non-deuterium materials. Additional study of this effect is necessary.

3.4 History of Λ_d Measurements

Previous experiments measuring the muon capture rate on the deuteron took a variety of approaches to deal with the complex muon atomic and molecular kinetics in the extraction of Λ_d . The most important differences are whether they measure the final state neutrons or employ the lifetime method, and the density and protium concentration of the target. The results are plotted in Fig. 1.1 and three of these experiments are discussed in the context of the experimental design principles described in this chapter.

3.4.1 Wang et al.

This pioneering experiment done by a group from Columbia University first measured the muon capture rate to be $\Lambda_d = 365 \pm 96 \text{ s}^{-1}$ via the rate of neutron emission [49]. The stopping target was composed of mostly protium with a deuterium concentration of 0.32 %. At this concentration, the high transfer rate from μp to μd is sufficient that muons are quickly transferred to the deuteron. However, the high density led to large molecular formation rates such that the capture process takes place largely from the $p\mu d$ state and has a significant additional component from the $\mu^3\text{He}$ atom formed following $p\mu d$ fusion. These channels were disentangled using the time distributions of the hyperfine populations of these bound states. Additional theoretical work was employed to convert the rate of capture from the molecular state to that of the μd atom.

3.4.2 Bertin et al.

Bertin et. al. reported a result $\Lambda_d = 445 \pm 65 \text{ s}^{-1}$ using a low-density gaseous hydrogen target containing 5 % deuterium [50]. The spin flip reaction $\mu d(\uparrow\uparrow) + p = \mu d(\uparrow\downarrow)$ proceeds at a rate that is $\sim 10\times$ slower than the corresponding pure deuterium reaction. Therefore the quartet depopulation is very slow, and the muon effectively spends its entire lifetime in a statistical mixture of the two hyperfine states. The authors acknowledge the lack of information about the μd spin state at the time of this experiment, but they erroneously conclude that consistency of their result with the theoretical calculations is evidence that the muons depopulate the upper hyperfine state rapidly due to interactions with the protium. With modern knowledge of muon kinetics in deuterium and protium, it is known that this experiment was measuring a nearly-statistical population of the hyperfine states. A factor of ~ 3 is necessary to correct for this, leading to complete contradiction with theory and other experiments. Subsequent experiments employed liquid targets to ensure the spin state of μd atoms is well known, but this exacerbates the formation of molecules and rate of fusion reactions.

3.4.3 *Bardin et al.*

Bardin et al. obtained a result $\Lambda_d = 470 \pm 29 \text{ s}^{-1}$ using the lifetime method and a liquid deuterium target with a small protium impurity $c_p = 0.13 - 0.18 \%$ [44]. The high density of the liquid phase rapidly depopulates the quartet state, but also greatly increases the buildup of the $\mu^3\text{He}$ state from both the doublet and the short-lived quartet states, where muon capture proceeds with a 5 times higher rate than that of deuterium. Muon sticking via both fusion processes 3.16 and 3.5c contribute. The shift in the measured capture rate due to $p\mu d$ fusion was corrected with a zero-extrapolation on the concentration of protium. A $60 \pm 16 \text{ s}^{-1}$ correction was applied to account for sticking due to $d\mu d$ fusion.

3.5 *Summary*

Measurement of the rate of reaction (1.2) is especially difficult due to the neutral final state and the complex atomic and molecular kinetics of negative muons in deuterium, so MuSun measures the disappearance rate via electrons from the decay channel. A high-density cryogenic gaseous deuterium target is used to achieve the optimal balance between rapid hyperfine depopulation and the prevalence of MCF, both of which complicate the interpretation of the disappearance rate.

Muons capture on heavier nuclei occurs at significantly higher rates, so stops in non-deuterium materials and trace chemical gas impurities must be avoided. A small isotopic gas impurity allows a specific type of fusion reaction to eject the muon far from the stopping point, potentially leading to additional disappearance channels. This error can only be avoided by careful target gas preparation and verification of purity throughout the data run.

Experiments in this field must be designed to deal with the kinetics of muons in deuterium, and previous results have included significant uncertainties due to muon-catalyzed fusion and the population of hyperfine levels. The MuSun experiment is designed to minimize these uncertainties through the choice of thermodynamic gas properties and by tracking muon stops in an active target in order to measure Λ_d with the precision needed to constrain the two-body weak axial current.

Chapter 4

EXPERIMENTAL APPARATUS




Figure 4.1: Diagram of MuSun detector systems. The muons enter from the left through the entrance detectors μSC , μSCa (not pictured) and μPC that register the start time of the event. The muon passes through beam windows and comes to rest in a cryogenic deuterium vessel equipped with a TPC to measure the stop position. The muon eventually decays to an electron that is measured by the two multi-wire proportional chambers ePC1 and ePC2, and the barrel scintillator detector, eSC, which registers the decay time for the event.

MuSun is an ongoing experiment at the Paul Scherrer Institute (PSI) in Villigen, Switzerland. The data studied in this thesis was collected between June and October of 2011 at

the $\pi E3$ secondary beamline of the High Intensity Proton Accelerator facility, although the experiment has since moved to the $\pi E1.1$ and $\pi E1.2$ beam areas. This chapter will discuss the physical experimental setup, including the beamline and detector systems, as well as the electronics, digitization, and data-acquisition systems.

The experiment is an observation of Michel decay of muons in the bound atomic state μd to determine the disappearance rate. A clean decay event begins with a muon guided through a package of *entrance detectors* into a stopping target instrumented with a time-projection chamber (TPC) that measures the 3-D stop position of the muon. The muon forms an atom and eventually decays to an electron with energy between 0 and 53 MeV that is tracked through several detectors. The decay time is registered by a scintillator detector with large solid-angle coverage, and the disappearance rate is extracted from the full distribution of these decay times. See Fig. 4.1 for an overview of the detector systems. The primary challenges in this experiment are observing a large number of clean decay events and ensuring that these events are not polluted with background signals or other physical processes such as muons capturing on materials other than deuterium. The several detector systems described here are used to address these challenges and permit a data analysis that can use event selection to measure and constrain any systematic sources of error.

4.1 $\pi E3$ Muon Beamline

The primary high-energy proton beam is generated via a Cockroft-Walton accelerator and two cyclotron acceleration stages reaching 590 MeV with a beam current of 2.1 mA. The primary beam impacts on a graphite target, producing many particles including charged pions that decay to positive and negative muons with a branching ratio $> 99.99\%$. The parity-violating structure of the weak interaction in pion decay produces muons with their spin aligned along their momentum direction in the pion rest frame, and therefore the muon beam supplied to the MuSun target is polarized. This average polarization is discussed in Sect. 3.3.1 for negative muons. Positive muons will retain their polarization even after coming to rest and then precess in the ambient magnetic field. An external field is applied to raise the frequency and make the precession observable (see Sect. 4.3.1).

The MuSun beamline was tuned to select muons with a momentum of $40 \text{ MeV } c^{-1}$ (kinetic




Figure 4.2: The π E3 beamline extending from the primary proton target, under the concrete block at the top of the image, and through the beamline elements to the electron detector in the foreground. The 4 blue cabinets near the top are the beam kicker, followed by the electron separator and two red quadrupole magnet triplets. A beam slit is installed between the two triplets. This is actually the MuCap experiment installed in the π E3 beamline. For the MuSun installation, see Fig. 4.3.

energy 7 MeV) such that the muons penetrate the various entrance detectors and thin windows to come to a stop in the middle of the TPC in the deuterium vessel. The accepted range of momenta is adjusted to maximize the number of muon stops in the target using a slit system between dipole bending magnets. The significant contamination of electrons in the muon beam with the same momentum is removed using an $\vec{E} \times \vec{B}$ velocity-selecting




Figure 4.3: The MuSun experiment in the $\pi E3$ beam area. CHUPS and the deuterium gas handling system are in the foreground, the TPC and electron detector are in the background. The muon beamline from the primary proton beam target extends behind the electron detector.

separator. Quadrupole beam elements guide the beam and focus it to a $2 \times 2 \text{ cm}^2$ spot at the muon entrance detectors. The raw rate of muons (without the kicker) measured by the entrance scintillator was 67 kHz.

4.1.1 Kicker

A beam kicker originally built for the MuLan experiment and used in the MuCap experiment allows a time structure to be introduced into the DC muon beam. In a section of beam pipe, four parallel plates are charged to a potential difference of 25 kV and the resulting electric field deflects the beam into a set of collimating slits. The beam is deflected following the detection of a muon in the entrance scintillator to avoid the pile-up of multiple muons in

the target at the same time. The beam is reverted after 25 μs (~ 11 muon lifetimes). The charging of the plates is driven by 2 cabinets that each contain a stack of 17 MOSFET cards to drive the switching on and off of the potential in 60 ns. The fast switching time combined with electronic signal processing delays and transmission times from the entrance detectors to the kicker and the delay as muons evacuate the beam pipe results in a total 600 ns delay for the muon rate seen by the entrance scintillator to fall. The extinction of the beam rate is not complete because the transverse beam profile has tails that are not perfectly deflected. The ratio of the unkicked average beam rate to the kicked rate is called the *kicker extinction*, and was typically 70-80 when the beamline was tuned for μ^- and only ~ 30 for μ^+ . While not all the beam elements are symmetrical with respect to polarity inversion of the muon beam, the cause of the discrepancy in kicker extinction is unknown. The autocorrelation of muon signals in the entrance scintillator can be seen in Fig. 4.4. The shoulder starting at 0 ns shows the unkicked rate of muons, modulated by the proton cyclotron RF, and the extinction to the kicked rate after 600 ns.

Without the kicker, only ~ 30 % of events are free of additional muons entering the detector system within the 25 μs pile-up protection window (when the beam rate is optimized for the kicker-less condition). Increasing the raw beam rate and introducing the kicker boosts this fraction to about 90 %, which is critical to accumulate the needed 10^{10} pile-up-free decay pairs in a beam period lasting a few months. The ideal raw beam rate that maximizes the pile-up protected event rate is a balance between the number of pile-up muons entering following the trigger muon despite the kicker extinction and the time between the end of the pile-up protection window and the subsequent trigger muon entering the target.

4.2 Muon Entrance Detectors

After exiting the beam pipe through a 75 μm Mylar window, the muon beam passes through an entrance scintillator (the μSC) and wire-chamber (the μPC). It then travels through a small air gap into the insulation vacuum vessel containing thin layers of super-insulation. Finally, it passes through a beryllium beam window into the cryogenic D_2 vessel and stops in the active volume of the TPC.

The entrance detectors begin with the μSC , a 500 μm thick scintillator covering an area

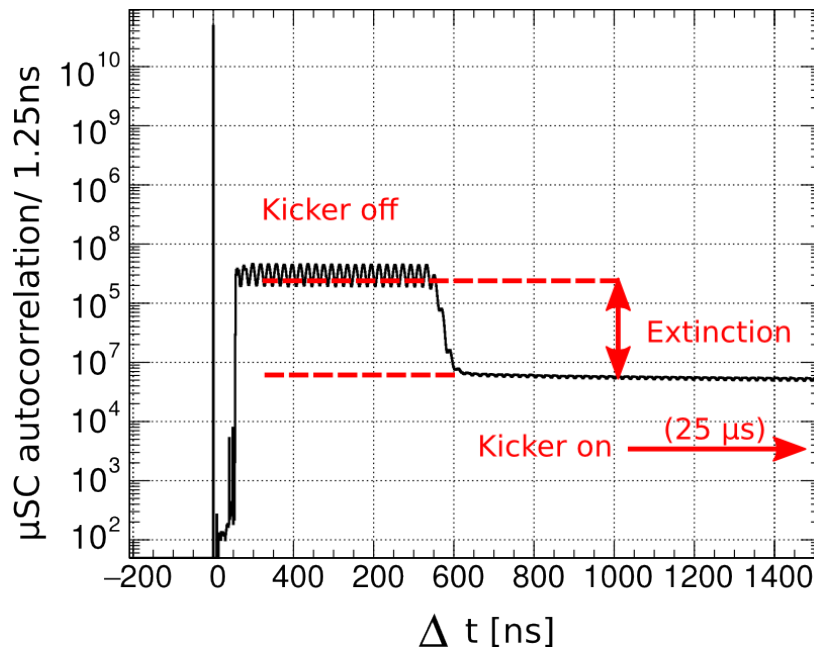


Figure 4.4: Autocorrelation of μ SC muon signal, showing the deadtime of the scintillator followed by the full unkicked beam rate, modulated by the 50 MHz cyclotron RF. After nearly 600 ns, the rate of incident muons is reduced by a factor of ~ 70 .

of $70 \times 70 \text{ mm}^2$ providing the fast timing that serves as the start signal for the lifetime measurement. The signal from the photomultiplier tube (PMT) is fed to a linear splitter, creating two copies of the analog waveform. One copy of the analog waveform is recorded by a 500 MHz, threshold-triggered waveform digitizer and the other is converted to a logic signal via a threshold discriminator. This signal is recorded by two separate CAEN time-to-digital converter (TDC) modules and is also used to drive the fast logic associated with the beam kicker. Beam electrons having the same momentum are distinguished from muons by their smaller energy deposition using a threshold on the μ SC signal (see Fig. 4.6). Electron signals do not trigger the kicker or pile-up protection in the data analysis. Digitized waveforms allow more fine-tuned separation of electron pulses, as well as improved pulse-pair resolution when compared to the TDC signals.

Following the scintillator, a set of two multi-wire proportional chambers collectively called




Figure 4.5: Detector package mounted on the end of the muon beam pipe, including the μ SC and μ SCa (not shown, behind the beam pipe) scintillators followed by the μ PC wire chamber. During operation, the electron detector is moved to the left in this image so there is only a small air gap between the μ PC and the entrance window to the insulation vacuum and TPC cryo-vessel.

the μ PC provides the x- and y-coordinates of the muon as it passes through the detectors, as well as redundancy with the μ SC. Each of the x- and y-wire planes of the μ PC consists of an array of sensitive anode wire separated by thin cathode foils. Each anode plane comprises 24 wires at a pitch of 2 mm, for a total sensitive area of $48 \times 48 \text{ mm}^2$. A muon passing through the chamber will ionize gas near an x-wire and a y-wire, and this coincidence is used to identify the position of the particle. With the μ SC scintillator alone, two muons that arrive at nearly the same time are not resolved separately due to the limited pulse-pair resolution. However, the additional x/y coincidence with the μ PC reduces this occurrence to only the cases where the muons also fire the same x- and y-wires. Additionally, by histogramming the location of coincident x/y hits, an image of the transverse beam distribution is made, which is used in beam tuning and as an input to the GEANT simulation of the experiment. The beam is focussed to about a $2 \times 2 \text{ cm}$ spot on the μ PC (see Fig. 4.6) and the entrance

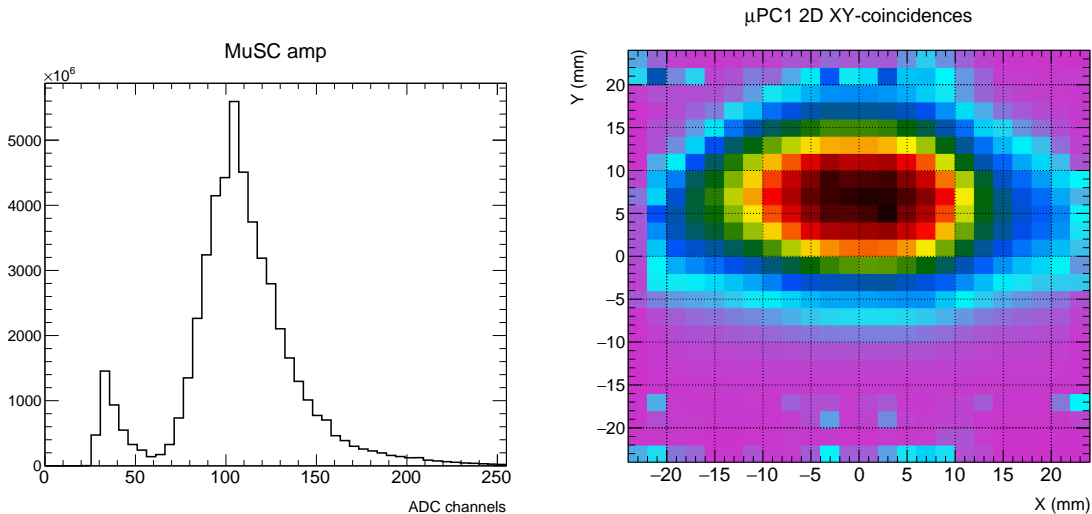


Figure 4.6: (Left) Amplitude spectrum of μ SC waveforms. The small signals just above threshold are beam electrons and the large peak is the muon energy deposition. (Right) Histogram of xy-coincidences in the μ PC wire chamber.

to the target vessel.

4.3 CryoTPC

The cryogenic deuterium vessel is pressurized to 6 bar and cooled to 34 K and enclosed in an insulation vacuum vessel. Muons enter the deuterium vessel through a 0.4 mm hemispherical beryllium window, designed to minimize multiple scattering of the muon beam. The momentum of the muon beam is selected so that the muons come to rest in the deuterium gas with the distribution of stopping muons centered in the sensitive volume of the TPC detector.

The rectangular prism TPC is unconventional in a number of ways in addition to the cryogenic operating temperature. The charge collection anodes are $16 \times 18 \text{ mm}^2$ gold-plated (thickness $2 \text{ }\mu\text{m}$) copper pads on a MACOR substrate. Gas amplification with thin wires is not possible due to the high deuterium density, so the TPC operates as a simple ionization chamber and not as a proportional counter. The drift electric field is provided by an aluminum cathode plate separated from the Frisch grid by 71 mm and biased to




Figure 4.7: Photograph of the TPC detector removed from the deuterium vessel, with the coordinate system used in this thesis labeled (left) and close-up of the pad plane (right). The muon beam enters from the left, traveling in the positive z -direction. The cathode plate is mounted on top of the 4 slotted MACOR posts, which stand on the frame supporting the Frisch grid. Ionization electrons drift from the interaction site downward. Readout cables and the high-voltage bias supply connect via feedthroughs in the back flange. These pictures are taken after several upgrades were made following the collection of the R2011 data, but the geometry and general features remain the same. The R2011 pad plane was plated with a thin layer of gold, not silver.

-80 kV. The field is made uniform by seven stainless steel field-shaping wires wrapped around MACOR support posts that are biased with a resistive voltage-dividing ladder from the cathode to the grounded frame.

The MuSun TPC Frisch grid is a plane of wires oriented parallel to the anode pad plane and separated by 1.5 mm. In an ionization chamber without a Frisch grid, the induced current on a collection anode is a complicated function of the geometry of the electrodes, and will typically include a collection signal from drifting electrons that accumulate on that anode and a bipolar induction signal from drifting charge collected on neighboring electrodes. More importantly, the total induced charge will depend on the height (distance from the anode plane) of the ionization event and the ions drifting toward the cathode will shield

a part of the electric field. A Frisch grid shields the anodes from charge in the main drift volume while allowing a well-defined signal from charge drifting in the small gap between the grid and anode plane. The bias, spacing, wire pitch, and wire diameter of the Frisch grid are selected to optimize isolation of the drift field and transparency to drifting electrons [51]. In the R2011 TPC, the wire diameter is 25 μm , the wire pitch is 250 μm , and the ideal bias voltage is -3.6 kV. However, due to high-voltage sparking, we were unable to maintain the grid bias and instead used -3.0 kV.

The drift velocity of ionization electrons, which generally depends on the ratio of the drift field to the gas density, is $5.3 \frac{\text{mm}}{\mu\text{s}}$ at R2011 conditions, leading to a maximum time difference of 14 μs between the μSC signal and the arrival of the ionization signal from the stopped muon. This velocity is measured simply by dividing the height of the TPC from the Frisch grid to the cathode plane by the difference in times of the earliest and latest signals observed from muon tracks. The final analysis cuts define a *fiducial volume* that only includes the central gas volume. Only decays from stops within this volume are used to ensure that the muon did not leave the TPC. In the y -direction, along the drift field, this cut is $15 \text{ mm} < y < 55 \text{ mm}$, while in the x - and z -directions, a single row of pads along the border of the TPC is used to veto muon stops.

Signals in the form of current pulses on the anode pads are transformed via a custom-built charge-integrating pre-amplifier and several shaping amplification stages, all designed by the PNPI group. These voltage pulses are recorded by a threshold-triggered waveform digitizer (WFD) operating at 25 MHz, originally designed and built by the Boston University group for the MuLan experiment.

Because the pre-amplifiers were not designed to operate in the insulation vacuum of the CryoTPC, there was a $\sim 1 \text{ m}$ cable run between the anode pad and the pre-amplifier inputs. The signals are transmitted on unshielded traces on flexible Kapton cables to reduce the input capacitance at the pre-amplifiers. These signal cables were effectively shielded from the environment by the vacuum vessel, but not from the high-voltage electrodes that were also inside the vacuum. This resulted in microphonics, sensitivity to high-voltage sparking in the vacuum volume, and other sources of radio-frequency pickup before the initial gain stages. In the runs following R2011, this system was substantially upgraded with in-vacuum,

cooled preamplifiers located just outside the cryogenic vessel, resulting in improved resolution as well as reduced sensitivity to the above-mentioned pick-up problems [52].

4.3.1 μ SR Magnet

Muons produced from pion decay, $\pi^- \rightarrow \mu^- + \bar{\nu}_\mu$, are polarized with their spin direction aligned parallel to their momentum, due to the structure of the weak interaction. Negative muons that form muonic atoms will lose this polarization during the atomic cascade to the 1S state through large spin-orbit couplings producing high-speed spin precession for random durations (See Sect 3.3). Positive muons stopping in the TPC will not form atoms and will retain their initial polarization. Precession in ambient magnetic fields and depolarization of these muons during the measurement period can result in a systematic error in the lifetime measurement. This arises through the imperfect cancelation of detectors on opposite sides of the target due to gain and threshold variations as well as a muon stopping distribution that is not perfectly centered. In order to control the precession, a ~ 50 gauss magnetic field was applied to the fiducial volume, transverse to the beam axis using a saddle-coil magnet wrapped around the insulation vacuum vessel. While the magnetic field covers the active TPC volume, it is not strong enough to bend particles for momentum selection, as is common with TPCs in collider experiments. The period of precession of a free muon in a field of 0.5 G (approximately Earth's magnetic field) is $T_\mu(0.5 \text{ G}) = 150 \text{ } \mu\text{s}$ while in an applied field of 50 G, this period is 1.5 μs . In a fit of the decay time distribution, a precession signal with the larger period is highly correlated with the muon lifetime, but the shorter period is much easier to separate.

4.4 *Electron Detectors*

The electrons from muon decay are tracked with two detector systems - a pair of cylindrical multi-wire proportional chambers, ePC1 and ePC2, and a barrel of double-layer plastic scintillator panels, called the eSC. The wire chambers allow measurement of the electron track azimuthal angle and the angle to the beam axis, further reducing backgrounds and allowing systematic studies based on the electron track direction. The scintillator panels provide the fast timing signal used to measure the decay time of the muon and requiring

a coincidence between PMT signals reduces non-particle backgrounds. In coincidence, the solid angle covered by the electron detectors is 3π and the average efficiency of detection for these angles is 80 %.

4.4.1 Electron Multi-Wire Proportional Chambers




Figure 4.8: Photograph of wire chambers ePC1 (right) and ePC2 (left) being serviced. The amplification and digitization electronics mounted on small cards around both ends of the chambers can be seen in more detail in Fig. 4.9. Cylindrical copper mesh sheets are wrapped around the inside of ePC1 and the inside and outside of ePC2. These provide additional RF shielding for the long wires of each of the chambers, but the added mass is minimized to avoid scattering and degrading electron tracking.

The two cylindrical wire chambers ePC1 and ePC2 (Fig. 4.8) each comprise longitudinally oriented anode wires sandwiched between two layers of aluminized mylar cathode strips, which together constitute three sensing electrodes for ionization events. The smaller chamber, ePC1, with a diameter of 384 mm is mounted just outside the target insulation vacuum vessel. It has an active length of 580 mm spanned by 512 anode wires and 192

cathode strips in each plane. The larger chamber, ePC2, with a diameter of 640 mm, is mounted inside the eSC scintillators. It has an active length of 800 mm, 1024 anode wires, and 320 cathode strips in each plane. The cathode strips wrap around the cylinder with a screw angle of $\sim 45^\circ$, but opposite helicities for each plane. Both cathode planes and the anode wires are instrumented, so the coincidence allows for a 2-dimensional reconstruction of the point of intersection with the electron track, although only one cathode plane and an anode wire is required for the coincidence. With both wire chambers, a 3-dimensional track is reconstructed.

The high voltage bias on the anode wires was 2630 V for ePC1 and 2800 V for the larger chamber. The HV bias stability for both chambers was good over the run, though for $\sim 40\%$ of the data, the bias of ePC2 was elevated to 2850 V to improve efficiency as long as excessive sparking could be avoided.

The planes of each chamber are divided into sectors of 16 wires each. Cathode sectors are read out via an integrated circuit card containing a charge-integrating stage as well as shaping stages. On the same circuit board, mounted directly on the chamber, threshold discriminators identify these pulses and they are converted to digital signals. The anode readout of the chamber is isolated from the bias voltage via coupling capacitors and each card handles two sectors, but is otherwise similar. The digitized signals are passed to an electronics rack containing custom-built TDCs that time-stamp the signals and record time and channel number words into a buffer. The time resolution of this system (> 100 ns) is much poorer than that of the eSC detector system, owing to the required amplification and shaping time, so the decay time of the electron is always derived from the eSC. Conversely, the spatial/angular resolution of the eSC is very limited in comparison to that of the wire chambers.

The anode readout electronics were very stable throughout the R2011 run, but the cathodes had a tendency to pick up stray RF signals (from e.g., switching power supplies) and the LVDS drivers would oscillate if the old output pins made a poor connection to the readout wires. To aid both problems, the cards were redesigned using a 4-layer board and a mounting system that provided more complete shielding of the input pins (see Fig. 4.9) The new output connectors solved many cases of the oscillating LVDS drivers, but a few sectors




Figure 4.9: Amplification and digitization cards mounted on the cathodes of the ePC1 and ePC2 chambers. The older cards (left) had connector pins that were aged and faulty, causing oscillation of the LVDS drivers, and the shielding of the inputs from stray RF signals was insufficient. The new design (right) uses a 4-layer board with ground planes and a more robust aluminum shield covering (not pictured) to avoid RF pickup.

were still problematic.

At the beginning of the R2011 beam period, it was discovered that one of the anode wires in ePC2 had broken, probably during crane movement of the detector. The chamber electronics are designed to be able to operate with a broken wire by ramping the bias down to zero for one or more sectors. Operating in this mode is non-trivial, since the leakage current of the HV bias must be monitored, and this must be separated from the current drained by the resistor ladder ramping down the bias across the broken wire sector. The sector containing the broken wire as well as the sectors with partial HV bias were not sensitive to decay electrons, representing an overall loss of about 10 % efficiency in this detector system.

4.4.2 Electron Scintillators

The eSC detector system comprises 16 double-layer plastic scintillator panels arranged in a barrel of diameter 78 cm and length 90 cm around the TPC vessel. The panels are each 5 mm thick and are stacked directly on top of each other. Each of the 16 segments has a PMT on

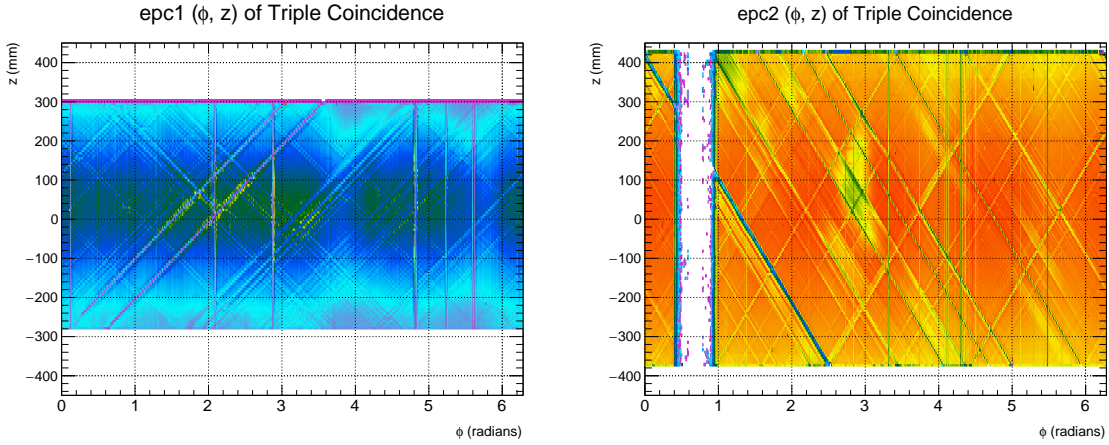


Figure 4.10: ϕ - Z distribution of triple coincidences of two cathode planes and anode wires for the ePC1 (left, linear z -scale) and ePC2 (right, log scale) wire chambers. The sector of ePC2 with the broken wire is clearly visible, as well as several individual wires that are either inefficient or too noisy and masked out of the data collection.

each end of both layers, for a total of 64 PMT signals. These signals are split and digitized with a threshold discriminator chained to a TDC and separately with a threshold-triggered waveform digitizer (WFD). The CAEN multi-hit TDCs used to record the time stamps of the discriminated PMT signals subdivide the external 25 MHz clock with an internal 32-step interpolator for an ultimate timing resolution of 1.25 ns. However, non-linearity in the interpolator forces almost all analysis applications to re-bin the signal to the external clock period of 40 ns. The WFDs recording the full waveform use the same digitizer boards as the TPC, but operated at 500 MHz. A 4-fold coincidence of all PMTs associated with a given eSC segment is required to reconstruct a muon-electron decay pair. The relative timing and scintillation light division between the two ends of each scintillator panel can be used to determine the angle between the electron track vector and the beam axis, but this has much poorer resolution ($\sim 25^\circ$) than the wire chambers. The timing resolution achieved by the eSC is 1 ns.

4.5 Neutron Detectors

Outside the eSC, eight 1.2 L liquid-scintillator neutron detectors were mounted around the target at 45° to the vertical at a distance 42 cm from the center of the TPC. Combining the overall solid angle coverage of 4.5 % with detector efficiencies including discrimination of neutrons from photons gave an overall acceptance of about 1 %.

Neutrons in MuSun are produced primarily through 3 reactions: muon capture on deuterons (Eq. 1.2), muon-catalyzed fusion (Eq. 3.5b), and muon capture on other nuclei, mainly TPC support materials and beam windows such as steel, beryllium, Mylar, and MACOR. By monitoring the production rates of fusion neutrons, the average population in the two hyperfine states of the muonic atoms can be determined. This is critical for fitting the decay time distribution to the full kinetics of muons in deuterium. Neutrons from muon capture on high-Z wall materials provide a strong check on mis-identification of good muon stops in the fiducial volume of the TPC. Muon capture on deuterium yields two neutrons with a distribution of energies because of the three-body final state. In addition to providing a check on the rate of this reaction with less precision than the lifetime method, these detectors allow the measurement of this spectrum.

For more detail on the neutron detectors and analysis of neutrons in MuSun, see the thesis of graduate student Nandita Raha [53].

4.6 Deuterium Gas

The deuterium gas target used in MuSun must satisfy several purity and thermodynamic requirements for the experiment to yield interpretable results. Chemical impurities such as oxygen and nitrogen must be present at levels less than 1 part-per-billion (ppb) due to the higher rate of muon capture on heavier nuclei. The isotopic impurity of ^1H must be under 100 parts-per-million due to the muon-catalyzed fusion process $p\mu d \rightarrow ^3\text{He} + \mu + 5.5 \text{ MeV}$ with the potential to eject the muon from the deuterium volume. The selection of thermodynamic conditions for the target gas was covered in detail in Chapter 3 leading to the optimized choice of 6.4 % density, corresponding to 34 K and 6.0 bar.

The deuterium gas is prepared using a cryogenic distillation column that separates pro-

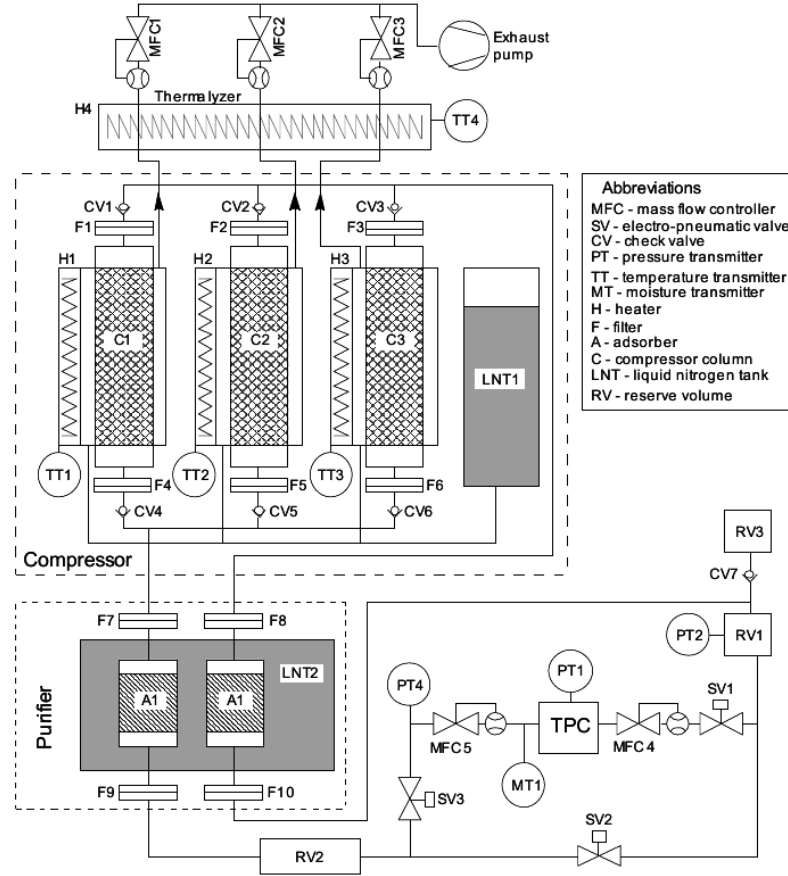


Figure 4.11: Simplified diagram of the CHUPS gas handling and purification system. Figure from Ref. [55].

tium (^1H) gas with high efficiency [54]. The system is portable so the gas is prepared on-site at PSI and stored just prior to beam periods. The quantity of HD impurity in the D_2 product is under 100 ppb and the rate of gas production can be as high as 1500 L d^{-1} . During the R2011 beam period, the isotopic purity was measured to be 10 ppm at the beginning of the run and 100 ppm at the end. The change is probably due to exchange with ordinary hydrogen and water in the filter or chamber walls over the course of the run, but both measurements are within requirements.

Chemical impurities of N_2 and O_2 will distort the measured lifetime because of the additional disappearance channel of muon capture on nuclei, $\mu Z \rightarrow (Z - 1)^* + \bar{\nu}_\mu$. This

capture rate scales as Z^4 and the high rate of transfer $\mu d + Z \rightarrow \mu Z + d$ means the gas must be clean of chemical impurities to <1 ppb. The TPC is operated below the freezing points of N_2 (63 K) and O_2 (54 K), but the vapor pressure is still appreciable at 34 K, yielding concentrations of 200 ppb (N_2) and 2 ppb (O_2) in undisturbed equilibrium (see Fig. 3.7). The gas circulation and filtration system employed by MuSun is called the Circulating Hydrogen Ultra-high Purification System (CHUPS) which maintains the gas at the required temperature, pressure, and purity.

The heart of the CHUPS system is a group of three 1-liter compressor columns connected in parallel and operated out of phase. Each column, containing a hydrogen adsorbent, fills with gas during the cooling cycle using LN_2 . As the column is heated and pressure rises, a check valve on the input closes and the output valve opens so that the deuterium gas flows through the output line. The gas passes through a zeolite adsorber in another volume maintained at 77 K to remove nitrogen and oxygen impurities. This design minimizes the number of moving parts, allowing for high-purity operation, and is capable of pumping at a rate of 3 standard liters per minute.

4.7 Data Acquisition System

Up to this point, I have discussed detector systems from incoming particles to when the data is digitized, usually as a single time stamp record or a full waveform trace. The data-acquisition system (DAQ) is the accumulation of these many digital data streams for permanent recording as well as online analysis. MuSun uses a customized version of the Maximum Integrated Data Acquisition System (MIDAS) developed at the PSI and TRIUMF laboratories [56]. The data-acquisition cycle is mediated by central FPGA logic that integrates the status of all readout module buffers to synchronize the periods when they are active and recording data. The readout block cycle is a series of ~ 140 ms-long segments where all detector systems are active at once separated by ~ 25 ms periods where the module buffers are recorded and emptied. When a module buffer is nearly full it signals the central logic and requests that all other modules discontinue data collection. This cycle of alternating deadtime-free data collection blocks and simultaneous readout results in an overall DAQ live-time of 80 %.




Figure 4.12: Block diagram of data acquisition system. The frontend electronics are controlled and read by drivers on frontend computers following each MIDAS block of data collection. The data streams from these many frontends are combined by an Event Builder on the backend computer into a MIDAS data file, which is written to storage. The data files are transferred to backup storage and to a small local cluster for near real-time analysis of the full dataset.

MIDAS divides the tasks of operating detector readout modules, reading and writing to storage, and performing online data analysis into programs called *frontends* (Fig. 4.12). Each detector system frontend will interface with one or more electronic modules to read data in the readout period of the DAQ cycle, as well as handle configuration and report errors to the user. The detector system frontend programs in MuSun run on separate computers connected by a local network. The data from the front ends is accumulated into a single stream on the main DAQ computer by a program called the Event Builder and is written to disks and archives by Loggers in the MIDAS file format. This computer also runs many slow control frontends that monitor environmental data, detector settings such as PMT bias voltages, beamline element settings, and data from the CHUPS system, recording directly into the data stream. The MIDAS system includes a program, MU, that analyzes a handful of data blocks per file in real time as they are collected. This program is nearly identical to the offline data analysis, so this provides an essential near-real-time first look at a fraction of the incoming data stream.

One data file collected by the DAQ is truncated at either 1.61 GB or 3000 data collection blocks, but in practice most data runs hit the file size limit. The average data file size is 1.60 GB, containing 4.6×10^6 raw muon entrances as defined by a hit in the μ SC, before any analysis cuts.

MIDAS files are initially recorded to a disk buffer on the main DAQ computer. The Logger programs copy this data to permanent storage on 2 TB hard disk drives that are ultimately shipped to the University of Washington. Additionally, the data is transferred over the internet using Globus Online [57] to the Ranch mass-storage system at the Texas Advanced Computing Center where the main data analysis is performed on the Stampede computing cluster [58]. A third copy of the data is stored on the FTP Archive system at PSI. The rate of data transfer to the Ranch and PSI FTP Archive systems was usually sufficient to outpace the rate of data collection from the experiment, but system outages or network congestion often led to a backup of data. An additional RAID array of 2 TB served as a buffer, allowing enough time for the networked data transfer to catch up during beam outages or other experimental shutdowns.

In addition to the MIDAS online analyzer program, a small cluster of three data analysis computers would each run 3-4 MU jobs on complete data files as they were collected. This cluster provided an analysis of the data with much higher statistics than the online analysis built into the DAQ, and subtler problems and features could be detected in daily or weekly sums of these files.

4.8 *R2011 Beamtime Summary*

MuSun was granted 12 weeks of time on the π E3 beamline at PSI, from 22 June through 11 September 2011, divided into three 4-week operation periods. Preparation for the run began several weeks before the allocated beam time with the PNPI group purifying deuterium gas and preparing the CHUPS system. UW post-doc Frederik Wauters and I spent 3 weeks improving the noise and RF pickup difficulties with ePC2 experienced during previous beam periods. This included re-wiring detector and electronics grounding, isolating interfering RF signals, and replacing many of the chamber-mounted amplifier cards, as well as the previously discussed bypass of the broken sector. To setup in π E3, the electron detector was craned

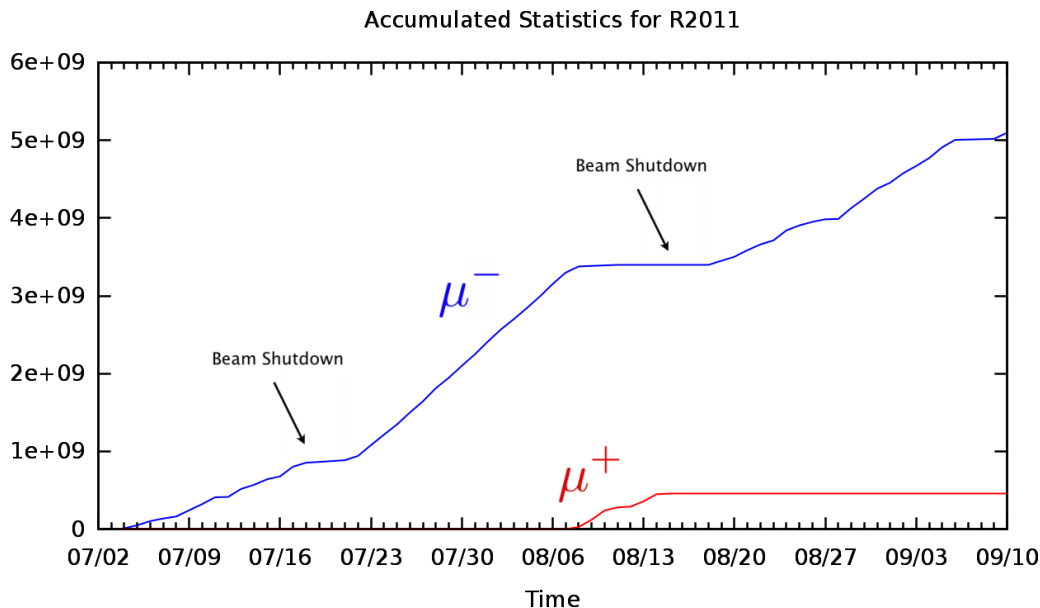


Figure 4.13: Accumulated statistics of muon-electron decay pairs that pass all analysis cuts over the R2011 beam period. Flat sections correspond to shutdown periods of the primary proton accelerator.

into place, cables were routed to the counting room barracks, and a air-conditioned enclosure was constructed around the whole beam area to provide a low-humidity environment for the high-voltage electronics. Significant efforts were made to reduce microphonics in the TPC readout electronics by anchoring the TPC to the floor and adding a base-line restorer circuit to the amplifiers before digitization.

The 3.5-week operation periods of the primary proton beam separated by maintenance shutdowns roughly divide the R2011 data into three major μ^- datasets and one μ^+ dataset (Table 4.1). Before the first shutdown, setup of the experiment was being completed, so there were small changes such as beam tuning and the powering of the μ SR magnet midway through this period. The field from the magnet was strong enough to distort the pulse shape and reduce gain in the large photomultiplier tubes used for the neutron detectors, so during the first shutdown the detectors were shielded with mu-metal sheaths that eliminated the effects of the μ SR magnet. Most of the second beam operation period comprises the

Data set	Start-End Date	$\mu - e$ decay pairs
Pre-neutron shielding	8 Jul - 22 Jul	1.4×10^9
Golden	22 Jul - 15 Aug	2.1×10^9
Wire chamber noise	18 Aug - 10 Sep	1.8×10^9
μ^+	7 Aug - 14 Aug	0.5×10^9

Table 4.1: Summary of major data-taking periods in the 2011 run separated by data quality. All listed data is viable for the lifetime measurement. The number of $\mu - e$ decay pairs includes all analysis cuts.

golden dataset where all detector systems were operating as expected. The third period had problems with noise in the μ PC and the ePCs as well as a kicker malfunction requiring emergency repair. These problems were not expected to substantially change the quality of the data, but were separated from the golden data for ease of comparison.

After the final production data collection, several days of systematic tests were performed, including variations in the target gas temperature and density. We discovered an unexplained signal loss in the TPC for events with long drift times that became more severe with rising temperature. The best explanation for this effect is an oxygen impurity that is frozen on the walls of the vessel during normal operation. At elevated temperature, the oxygen vapor pressure is higher and drifting ionization electrons attach to the electronegative impurity before collection on the anode pads. This discovery led to the suspicion of chemical impurities in the prepared gas, and motivated a campaign to measure the N_2 and O_2 concentration *in-situ* in the following runs, R2014 and R2015.

To summarize the recorded data, 14804 MIDAS files of good quality were collected with total size of 23.6 TB during R2011. Of these, 13254 files and 21.2 TB used a negative muon beam with the remainder running positive muons for systematic checks and controls. The number of muon-electron decay pairs that pass the standard analysis cuts is 5.2×10^9 for μ^- and 0.5×10^9 for μ^+ .

The remaining chapters will discuss the analysis of this data, including integration of basic

detector system data, description of the analysis software and procedures, and a systematic error caused by interference of muon-catalyzed fusion reactions with TPC tracking.

Chapter 5

SOFTWARE AND ANALYSIS

The MuSun experiment aims to measure the lifetime of muons in deuterium gas. Muon stops and decay electrons are identified by separate detector systems and assigned a time stamp. The difference between electron and muon times is histogrammed and the lifetime extracted by fitting this time distribution. Analysis cuts are imposed on events to ensure data quality, select for good muons stops in deuterium, and reduce backgrounds.

Understanding systematic errors is at least as important as achieving statistical precision in MuSun, where the muon lifetime is measured to 10 ppm. Several potential errors arise from event selection criteria that favor certain decay times over others, distorting the final distribution. These distortions will vary in strength when measured against an event parameter that would otherwise be irrelevant, such as electron track geometry or muon stop position within the gas. Therefore, at every analysis stage, as higher-level particle objects are considered, they are flagged with additional parameters and algorithm choices so that the final decay time distributions can be assembled with analysis cut variations.

5.1 Data Analysis

The main framework for the analysis of the MuSun data is based on the MIDAS analyzing software allowing developers to write user-defined analysis modules. For this purpose, most of the software is written in C++ as well as some modules based on the special-purpose high-level Muon Query Language that was specially developed to efficiently find coincidences between detectors. The raw data files are processed and elementary consistency checks disregard data sections that do not pass basic quality cuts. Hereafter, the raw detector signals are combined to reconstruct the physical objects of the incoming muon, its track within the TPC and the outgoing decay electron. These objects are then stored in a ROOT tree, which is a special container for C++ objects. The tree stores all the physically relevant

information, such as particle energy, detection times, and track geometry.

5.1.1 Data Organization

The data-acquisition system (DAQ) records data in the MIDAS file format, where detector data is written to a formatted binary buffer called a MIDAS bank, and banks are combined into MIDAS events. MIDAS events are concatenated to form the MIDAS files (1.6 GB) that are written to disk and serve as the input to the analysis software. Each MIDAS file containing $\sim 4 \times 10^6$ muons is processed independently through each analysis stage and the final histograms are combined. The 32 TB R2011 dataset, of which 20 TB is analyzable production data, is stored on a tape archive system. The first analysis stage, called MU, analyzes the MIDAS events sequentially and produces higher-level C++ objects for each detector system. The output of this stage is a ROOT `TTree` with entries based on the class `TMusunEvent` and a ROOT file containing histograms used for diagnostics and preliminary decay time distributions. The `TMusunEvent` class, built around a muon signal in the μ PC and μ SC entrance detectors, includes the C++ objects from all detector systems if their time stamp falls within ± 35 μ s of the entrance detector time.

The second stage of analysis, called MTA, takes the ROOT tree as input and processes `TMusunEvent` objects, ultimately filling decay time histograms for many analysis variations. The primary output of this stage is a ROOT file containing diagnostic histograms and decay time histograms.

The MU and MTA computation is performed in parallel batches of files on the Stampede high-performance computing cluster at the Texas Advanced Computing Center [58]. The processing time per MIDAS file for the MU stage is 80 min and the output file sizes are 50 MB for histograms and 2.5 GB for the `TTree`. The processing time per tree file for the MTA stage is 50 min and the output file size is 50 MB. The full R2011 dataset is divided into batches of 1000 MIDAS files containing (before any analysis cuts) 4×10^9 muon events that are processed in parallel as a single job on Stampede with one file per core. The total number of files in all batches is 13234. Each Stampede node has 16 cores, so 63 nodes are reserved at once to process a single batch. A full-dataset pass of the MU analysis takes

18,000 cpu-hours and the MTA analysis takes 11,000 cpu-hours. Typically, multiple MTA passes are performed following a single MU pass as the analysis strategy evolves.

The remaining analysis, including fits of the decay time distributions, is performed on the smaller output files on local computing facilities. The procedure to fit the decay time distributions for the muon lifetime is handled by an additional library, LtFit, that implements the standard fit functions and variations for systematic studies.

5.1.2 Analysis Stages

The analysis software package for the MuSun experiment is executed in two stages in which raw detector information is progressively structured into objects representing physical particles, from which correlations can be formed.

The first stage, MU, performs low-level data interpretation, beginning with decoding the data format as recorded by the DAQ modules. Objects are created that represent waveforms from individual triggers of digitizers for the TPC, eSC, μ SC, and neutron detectors. TDC data for the μ PC and ePC wire chambers is reduced into clusters, and clusters from the different planes of each detector are combined into hit objects. Timing offsets between detectors are corrected and diagnostic histograms are produced so these coincidences can be examined, both in real time during data acquisition, and post-analysis. The TPC waveforms are processed by pulse-finding algorithms. The muon entrance detectors are time-correlated to produce **TMuEntrance** objects, which have flags for possible pileup and coincidence conditions of the μ SC and μ PC. These entrance detector objects are the basis for analysis at the next stage.

The second analysis stage, MTA, is based on muon events, described above. Pulse finding is performed on eSC and μ SC waveforms and the eSC and ePC hit objects are combined into electron tracks with several variations. TPC pulses are analyzed by track-finding algorithms to determine whether the event has a muon stop in the TPC and determine the properties of that muon stop. Finally, the data from particle objects is combined to determine the decay time of the muon, and whether the event passes the analysis cuts required in order to be filled into the decay time histograms.

5.1.3 Diagnostics

The MU and MTA programs produce many histograms as output that are used for diagnostics. Examples include timing differences between detectors, properties of clusters in wire chambers, and raw count rates of detector channels. A display program is used to format, draw, and examine these histograms for anomalous conditions that could indicate poor data quality or other errors.

5.1.4 Error Identification and Rejection

The preferred strategy is to identify sections of the data with errors or inconsistencies, flag them, and remove them at the last stages before making lifetime histograms. This allows for the comparison of error-free and error-prone data to identify the source of error. See Table 5.1 for a summary of error conditions that are identified at various stages in the analysis, but only cut just before filling decay time histograms.

TPCWFDBankError	TPC waveform digitizer error
ESCWFDBankError	eSC waveform digitizer error
MuSCWFDBankError	μ SC waveform digitizer error
MuSCMatchError	Too many μ SChits without copy signal
CAENErrorDAQ	Data-acquisition error in CAEN TDC module
CAENErrorEOB	End-of-block flag set in CAEN TDC module
CAENErrorTrailingEdge	Trailing-edges when only leading-edges requested
RolloverMissingError	Rollover sync. not present
RolloverMismatchError	Rollover sync. signal at the wrong time
RolloverMultHitError	Extra rollover sync.
CompErrorEOB	Compressor end-of-block error recorded
RoutedSequenceError	Missing or extra copies of μ SCsignals

Table 5.1: MIDAS block level data quality cuts. All of these are imposed in the final analysis except CAENErrorTrailingEdge.

5.1.5 Standardized Analysis Interfaces

Quantification of different sources of systematic error requires the ability to substitute any of several different algorithms in the final analysis cuts. For instance, in the analysis of R2011 data, three separate TPC pulse-finding algorithms and at least six muon tracking algorithms are employed, including the primary Road tracker described in Chapter 7. The analysis software must be highly modular to allow for parallel analysis with these variations. I constructed a uniform interface to use TPC pulses, through the class `TTPCGenericPulse` so that pulse-finding algorithms could be substituted for one another without the need to change the subsequent track-finding algorithm, preventing otherwise time-consuming and error-prone customization efforts. The `TTPCGenericPulse` class was extended using subclasses `TTPCGaussPulse`, `TTPCMiniPulse`, and `TTPCTOTPulse` before eventually determining that the TOT algorithm was the most generally useful, and the Mini algorithm was useful in special cases.

In a similar vein, we use several track-finding algorithms, each of which can be implemented using any pulses that conform to `TTPCGenericPulse`. The code that builds decay time distributions requires a standard set of TPC track properties, including stop position, energy deposition, and whether the stop is inside the fiducial volume. The tracking algorithms have a standard interface through `TTPCGenericMuonTrack` that provides these properties and a class `TTPCEventProperties` that accumulates data about the tracks into easily-queryable properties, such as whether the decay event has exactly one muon stop in the fiducial volume of the TPC.

Diagnostics for the tracking algorithms include many standard properties that are required regardless of the tracking algorithm used. The class `TTPCMuonTrackHistograms` encapsulates all of these. Similarly, decay time histograms also often need to be produced for different combinations of the TPC and electron tracking algorithms. For each combination, we make systematic checks of the lifetime as a function of quantities like stop position in the TPC, impact parameter of the electron vector, or angle of the electron track. All of these standard decay time distributions are accumulated in the class `TLifetimeHistograms` to ensure consistency of the cuts applied for the different combinations of algorithms.

5.2 Muon Events

Muons entering the detector system are identified by the entrance detectors and tracked as they come to rest by the TPC. The data analysis is ultimately centered around *muon events* with start time defined by a muon signal in the μ SC and coincidence and pile-up protection (PP) flags from the other entrance detectors. The set of flags identifying a good muon entrance is called `HasBestEntrance` in the analysis, and I will use this term in the text. A summary of these conditions can be found in Table 5.2.

5.2.1 μ SC

The entrance scintillator is sensitive to electrons as well as muons, so muons are identified with a threshold discriminator (see Fig. 4.6) producing a TDC hit. The time of this signal, $t_{\mu\text{SC}}$, is the start time for the event. An artificial deadtime (ADT) can be applied to μ SC signals to suppress afterpulsing in the PMT, but in the R2011 data, the default for this parameter is shorter than the 54 ns width of the logic signal recorded by the μ SC TDC and therefore has no effect. The PMT waveform of the μ SC is also digitized, and a pulse-shape template fit is applied to this data. The pulse-pair resolution with this method is only a few nanoseconds and arbitrary deadtimes and muon amplitude thresholds can be applied for systematic analysis. Effects due to pulse-pair resolution and deadtime are expected to be negligible for the analysis in this thesis, and only the μ SC TDC signal is used as the requirement for a good muon event.

5.2.2 μ PC

The entrance wire chamber, μ PC, separately records hits on x-oriented and y-oriented planes of anode wires. In each plane, hits on neighboring wires within a time difference of 230 ns are clustered together, forming x- and y-clusters. A deadtime of 300 ns is imposed on individual wires to limit re-triggering of the same anode. The x- and y-clusters are combined in coincidence within 190 ns to form xy-clusters. Good events require a μ PC-xy cluster in time coincidence with the μ SC ($-60 \text{ ns} < t_{\mu\text{PC}} - t_{\mu\text{SC}} < 260 \text{ ns}$).

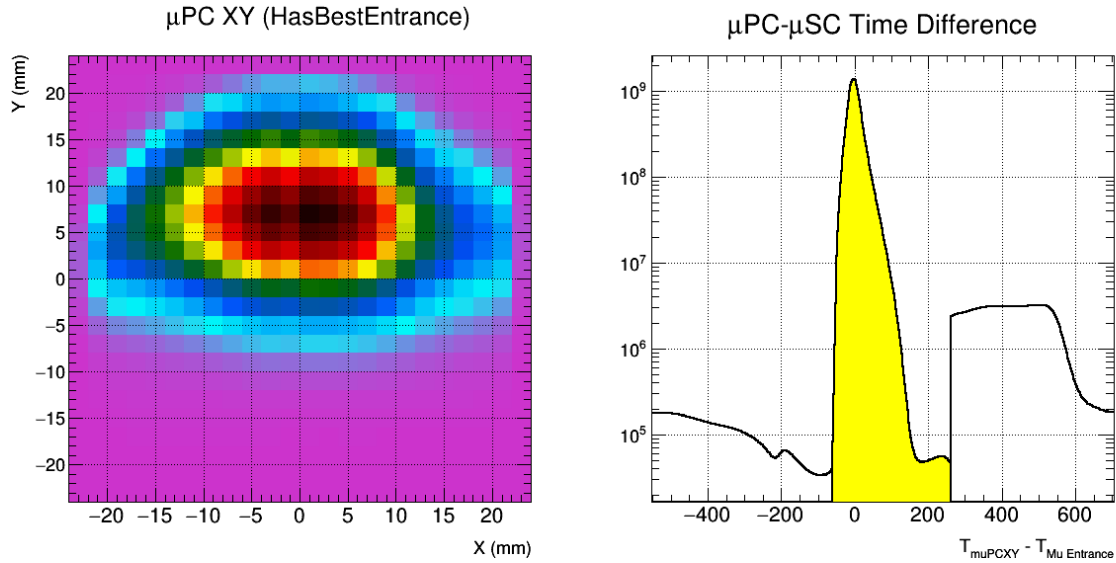


Figure 5.1: (Left) The x-y distribution of muons as seen passing through the μ PC detector for `HasBestEntrance` muon events. (Right) Time difference between μ SC signals and μ PC x-y clusters. The yellow fill indicates events within the coincidence window.

5.2.3 Kicker

The beam kicker is triggered by a muon signal in the μ SC, and the activation of the kicker is recorded as a separate TDC signal. Muons that enter after the kicker has fired or in a small protection window after the kicker beam deflection ends will not re-trigger the kicker. To be considered a good muon entrance, the kicker TDC signal must be in coincidence with the μ SC signal to be sure that the kicker was active during the event.

5.2.4 Pile-up Protection

MuSun analyzes only pile-up protected (PP) events, where a single muon has entered the detector system unaccompanied by additional muons within a $[-25 \mu\text{s}, +25 \mu\text{s}]$ interval. The beam kicker increases the fraction of PP events (see Sect. 4.1.1) to $\sim 90\%$ of muon entrances, and the remaining pile-up events must be rejected using event selection cuts. The signals available for pile-up include the μ SC, μ PC-x or -y clusters, μ PC-xy clusters, and the μ SCa

veto scintillator. The good event definition `HasBestEntrance` includes all of these except the single-plane $\mu\text{PC-x}$ or $\mu\text{PC-y}$ clusters. This is because during the R2011 run, the μPC system suffered from substantial RF pickup producing many non-physical hits in the data stream. Vetoing the single-anode-plane clusters would amount to a $\sim 10\%$ data loss, so these were left out of the pile-up definition, pending further systematic study of the PP efficiency. The $\mu\text{PC-xy}$ coincidence suppresses this RF pickup sufficiently that the reduction in data is minimal, so these are included in the PP definition.

Events near the beginning and end of MIDAS blocks warrant special consideration for PP. Recall that the MIDAS block is an uninterrupted period of data acquisition. A muon arriving at the beginning of the block may have been preceded by a pile-up muon that arrived before the MIDAS block data collection began. There is no way to know if muons near the end or beginning of the MIDAS block are pile-up protected, so they are not included in the data analysis; however, these detector signals are used for PP in neighboring events. This process is known as *bookending*, and 30 μs time intervals on either end of each MIDAS block is subject to this process. In principle, any section of data that needs to be vetoed (e.g. due to a detector system malfunction) undergoes a similar treatment, bookending around the missing data section. However, in the R2011 data, these kinds of errors were rare enough that whole MIDAS blocks were vetoed instead, with a minimal loss of statistics.

5.3 TPC

The MuSun CryoTPC is in many ways the heart of the experiment. Tracking the muon as it comes to rest is critical to ensure that only clean stops in D_2 are considered for the lifetime measurement. Particular attention must be paid to the interaction of muon tracking with additional signals from the muon-catalyzed fusion (MCF) reactions (3.5) because these interfere with the reconstructed stop position in a way that can lead to decay-time-dependent event selection. This source of systematic error is called *fusion interference* and is the main topic of this thesis, addressed in Chapter 7.

This section will first describe the pulse shape and event topologies that can arise from the interactions that can occur in the TPC, followed by the pulse-finding algorithm used to identify and characterize these ionization events. Next, I introduce three tracking algorithms

HasBestEntrance condition	Comment
μ SC TDC signal	Defines start time, $t_{\mu\text{SC}}$
μ PC xy-cluster	Coincidence $[-60 \text{ ns}, +260 \text{ ns}]$
Kicker fired TDC	Coincidence with μ SC
μ SCa TDC signal	Anti-coincidence with μ SC
μ SC PP	No addtl. μ SC signal in $[-25 \mu\text{s}, +25 \mu\text{s}]$
μ PC-xy PP	No addtl. μ PC-xy cluster in $[-25 \mu\text{s}, +25 \mu\text{s}]$
μ SCa PP	No μ SCa signal in $[-25 \mu\text{s}, +25 \mu\text{s}]$
Bookending	$t_{\mu\text{SC}}$ not within 30 μs of MIDAS block start/end

Table 5.2: Contributions to the definition of a good muon entrance event, `HasBestEntrance`.

that characterize events based on the pulse characteristics and topology. Finally, I will describe the TPC analysis cuts used for event selection in the lifetime measurement.

5.3.1 Waveforms and TPC Event Topologies

Particle signals in the TPC begin as clouds of electrons from gas molecules ionized by a charged particle traveling through or stopping in the TPC. The electron cloud drifts towards the anode pad plane and Frisch grid. In the final 1.5 mm between the grid and anode plane, the moving cloud of electrons induces an image charge on the anode. This induced charge is sensed by a charge-integrating pre-amplifier and several amplifier stages to produce a shaped pulse, which is digitized by waveform digitizers (WFDs).

Several sources of ionization signals in the TPC are enumerated here and will be discussed with examples.

1. muon tracks, spanning many pads
2. low- dE/dx particles like decay electrons and cosmic muons
3. electronic noise
4. ${}^3\text{He}$ and $\mu^3\text{He}$ from MCF, usually occurring on a single pad
5. neutrons from MCF single-site scattering from a deuteron

6. p+t from MCF, occurring on one pad or two adjacent pads

Sometimes these sources are distinguishable, but sometimes the overlap with other signals makes them difficult or impossible to disentangle.

1. Muons Beam muons stop with a distribution that spans the entire TPC due to the spread in momentum selection and longitudinal straggling and scattering on beam windows. However, tracks are generally straight within the TPC, aligned with the beam direction, and with several pads activated in a row leading up to the stop pad. The energy deposition per unit distance, dE/dx , increases as the muon comes to a stop, so that more energy is deposited at the end of the track. This peak in energy at the end of a track is referred to as the Bragg peak. A typical muon track is displayed in Figure. 5.2.




Figure 5.2: TPC event display showing the waveforms generated by a simple muon track, stopping in pad 22. The waveforms (left) are digitized in islands of 88 samples triggered by an electronic threshold. The pad display (right) shows the anodes that these waveforms correspond to. Darker colors indicate larger energy deposition. Signals arriving earlier in time originated lower in the TPC, and later signals came from ionization further from the anode plane. The signals from this track arrived at nearly the same time, indicating a relatively straight track, parallel to the z-axis in the TPC.

2+3. Decay electrons and Noise The electrons from muon decay deposit very little energy in the TPC, usually remaining below the detection threshold of the electronics. It is generally difficult to distinguish these signals from electronic noise.

4+5. ${}^3\text{He}$ from MCF The fusion reactions (3.5b) and (3.5c) yield a mono-energetic ${}^3\text{He}$ nucleus (or $\mu^3\text{He}$ ion) with energy 0.82 MeV (0.80 MeV) and a range of 0.3 mm (0.6 mm). The dense ionization leads to a large fraction of the electron-ion pairs recombining before they drift apart, reducing the observed energy of this particle to 0.32 MeV (0.50 MeV). The muon partially screens the charge of the $\mu^3\text{He}$, reducing the ionization density and charge recombination and increasing the range and observed energy. The short range of these particles means they are local to the pad on which the muon stopped. The signals appear as additional pulses overlapping the muon signal or delayed pulses appearing after the muon stop pulse. See Figure 5.3 for an example of a delayed ${}^3\text{He}$ event. The time delay between the muon stop pulse and the ${}^3\text{He}$ pulse is time of the fusion relative to the μSC time, t_f .

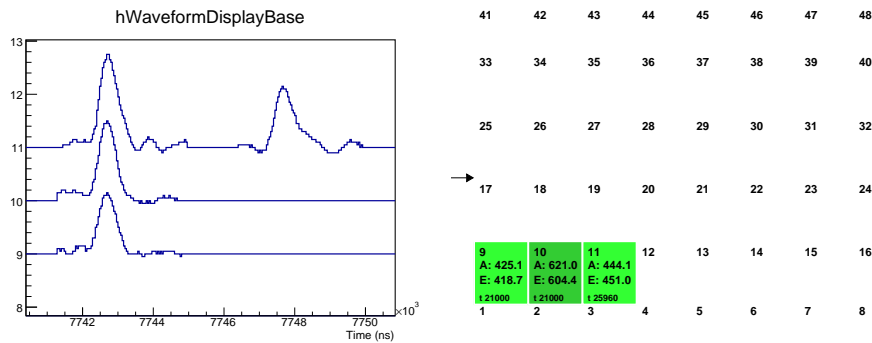


Figure 5.3: TPC event display of a delayed MCF reaction producing ${}^3\text{He}$. The mono-energetic pulse appears after the muon stop on pad 11. Depending on the time of the fusion, this additional pulse may partially or completely overlap the muon pulse. Pulses from prompt fusions are difficult to separate from a bare muon track.

The neutron from these fusion reactions interacts only very rarely in the TPC, but occasionally $n - d$ scattering events are observed with energy depositions of up to 2 MeV in a single pad. The neutron can interact anywhere in the TPC, so these pulses can appear far from the muon stop location.

Another source of events with this same topology is muon capture on nitrogen and oxygen nuclei from impurities in the deuterium gas. The nuclei break apart and deposit a range of




Figure 5.4: Spectrum of delayed pulses following a muon stop. The two largest peaks at 300 keV and 500 keV are ${}^3\text{He}$ and $\mu^3\text{He}$ from MCF. The red curve includes all signals, and the black curve uses a delayed electron requirement to suppress recoils from muon capture. Double-fusion peaks are barely visible at higher energies. Figure produced by F. Wauters.

energies in the TPC up to ~ 1 MeV. The difference between the red and black curves in Fig. 5.4 is due to these captures. The measurement of impurity capture recoils is discussed in Sect. 6.1.

6. p+t from MCF The p+t fusion branch (3.5a) yields a 3.02 MeV proton and a 1.01 MeV triton. The triton is local to the stop pad with a range <1 mm, while the proton range is 12.9 mm at the R2011 gas density, which is comparable to the $16 \times 18 \text{ mm}^2$ anode pad dimensions. Additional energy from the proton will appear on the stop pad and potentially the pads neighboring the stop pad, but the proton will never reach a pad that is further than 1 pad from the stop pad. Nevertheless, there is a large variety of waveform topologies from p+t fusion events. The simplest is a delayed fusion where the proton travels parallel to the anode (xz) plane so that all of the charge arrives at the anode at the same time as a single, large pulse. A delayed fusion event with this topology is shown in Fig. 5.5.

Unlike the muons that travel parallel to the z-axis and thus perpendicular to the TPC drift field, the protons from MCF are emitted isotropically and can travel vertically, parallel to the drift field. This spreads out the arrival of the proton's ionization cloud at the anode, and the Bragg peak of the proton becomes visible. This topology is shown in Fig. 5.6.

The p+t topologies shown so far have been from events with delayed fusions in which the

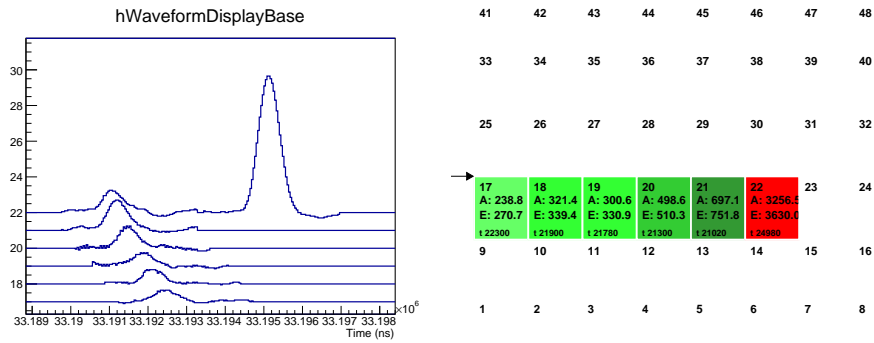


Figure 5.5: TPC event display of a delayed MCF reaction producing a proton and triton. The proton travels perpendicular to the TPC drift field, so that all of the charge arrives at once as a single pulse on the stop pad.




Figure 5.6: TPC event display of a delayed MCF reaction producing a proton and triton. The triton pulse occurs just after the last pulse in the muon track, followed by the Bragg curve of the proton traveling vertically toward the cathode.




Figure 5.7: Events with p+t fusions overlapping the muon waveforms. (Left) The fusion occurs shortly after the stop, with the triton pulse overlapping the muon pulse and the proton traveling down towards the anode plane, causing the proton pulse to arrive earlier in time than the muon pulse. (Middle) The muon either stops on pad 28 and the proton travels downstream into pad 29 or vice versa. (Right) The muon stops in pad 21 and the proton travels downstream into pad 22 or vice versa.

p+t and muon waveforms are separated. The most common fusion events, however, involve waveforms that overlap the muon track waveforms in time, sometimes in ways that make the two impossible to disentangle. Some of these topologies are shown in Fig. 5.7. The middle figure shows an event where the true location of the muon stop is ambiguous. The proton could be traveling downstream from a muon stopping in pad 28 or upstream from a muon stopping in pad 29. This ambiguity is crucial for the fusion interference effect discussed in Chapter 7.

5.3.2 Pulse-Finding

The TPC waveforms are processed with a simple algorithm to find pulses before they are combined into clusters and evaluated as particle tracks. A pulse is defined as a stretch of contiguous digitizer samples above a threshold. A baseline pedestal is defined as the zero-amplitude point for a waveform trace, since the electronic offset from zero in the digitizer is arbitrary. The pedestal is determined pad-by-pad using forced digitization of baseline electronic noise. This was done periodically during TPC operation. A *search threshold* is defined as 5 ADC counts above this pedestal. The waveform samples, S_i , are traversed and

each continuous stretch satisfying

$$S_i \geq \text{pedestal} + \text{threshold} = \text{pedestal} + 5 \text{ ADC}$$

is collected into a *sample stretch*. Stretches are considered good pulses if they pass an additional amplitude cut. The maximum sample in the stretch must be larger than 8 ADC counts. Pulses have the following properties:

- Start and end times, t_{start} and t_{end} , are given by the first/last sample exceeding the search threshold.
- Pulse amplitude is given by the maximum sample, $A = \max(S_i)$.
- Center time is given by the time of the maximum sample.
- Pulse integral is given by the sum of all samples in the pulse $I = \sum S_i$

See Fig. 5.8 for an illustration of this algorithm.

The pulse amplitude and integral are both measures of the ionization energy of the pulse. The two measures, once calibrated, are only different in cases where a pulse is extended in time due to an ionization cloud that extends parallel to the TPC drift field, like in p+t fusions (See Fig. 5.6). In this thesis, the pulse *energy* will refer to the integral unless otherwise mentioned.

Coordinates are assigned to pulses in order to reconstruct the tracks of the particle that generated them. The x- and z-coordinates of pulses are simply given by the coordinates of the anode pad. The y-coordinate is reconstructed from the time between the μSC and the pulse arrival, called the drift time,

$$t_{drift} \equiv t_{pulse} - t_{\mu\text{SC}}. \quad (5.1)$$

The y-coordinate is then

$$y_{pulse} = t_{drift} \times v_{drift}, \quad (5.2)$$

where the drift velocity is determined by the earliest and latest pulse arriving after $t_{\mu\text{SC}}$ and the distance between the Frisch grid and cathode plate.

$$\begin{aligned} v_{drift} &\equiv 71 \text{ mm} / (t_{drift}^{max} - t_{drift}^{min}) \\ &= 5.3 \frac{\text{mm}}{\mu\text{s}}. \end{aligned}$$




Figure 5.8: Pulse properties determined by the time-over-threshold pulse finding algorithm for a p+t fusion waveform. The first sample that is *threshold* samples above the pedestal is the start time of the pulse, and the last sample above is the end time. The maximum sample is the amplitude, and the time of this maximum sample is the center time. The sum of all samples between t_{start} and t_{end} is the pulse integral, which is a useful measure for an irregularly-shaped pulse like this.

5.3.3 Pulse Clustering

Pulses are grouped into pulse *clusters*, which are then evaluated to determine whether or not they are muon tracks. A single pulse is grouped with a cluster if the x-coordinate is within one pad of another pulse in that cluster, $\Delta x \leq 1$, and if the z-coordinate is within two pads $\Delta z \leq 2$ (see Fig. 5.9). The gap in z is allowed because in some cases, the muon travels along the boundary between two pads in a way that the energy deposition is under threshold for both pads.

Additionally, the pulses must be clustered in y (or equivalently, time). The condition for matching pulses i and j is the following:

$$|t_i^{start} - t_j^{start}| \leq 1000 \text{ ns} \quad \text{OR} \quad (5.3)$$

$$|t_i^{start} - t_j^{end}| \leq 1000 \text{ ns} \quad \text{OR} \quad (5.4)$$

$$|t_i^{end} - t_j^{end}| \leq 1000 \text{ ns}. \quad (5.5)$$




Figure 5.9: Cartoon illustrating pulse clustering in x and z . Pulses are represented by the colored pads. In this figure, there are two pulse clusters, one of which is only a single pulse. The other consists of four pulses and has a gap in the z -direction.

That is, the start or end times of two pulses must be within 1000 ns of each other.

5.3.4 Muon Tracking

Pulse clusters are evaluated individually to determine if they can be classified as a muon track with three different algorithms, or *trackers*. Each tracker must determine the coordinates of muon stop location and some properties of muon tracks such as the muon stopping energy. First I will describe the “Basic” tracker, and use it to introduce a few concepts common to all trackers. The next algorithm is the Road tracker, which is the algorithm used when discussing fusion interference in Chapter 7. Finally, the PDir tracker performs better in distinguishing upstream-going from downstream-going protons (see Fig. 5.7), but makes the calculation of the lifetime shift from fusion interference more difficult.

As a prerequisite to describing the trackers, I will define the term *z-slice* and the concept of upstream energy, E_n . The tracking algorithms often make use of the projection of all track pulses onto the z -axis. The anode pad plane is a grid of 6×8 pads in the (x, z) plane, so

projecting out the x-dimension means summing the energy of all pulses that have the same z-coordinate. The y-coordinate is equivalent to time, so projecting out the y-dimension means summing the energies of all pulses in the cluster that appear on the same pad. I will use the term *z-slice* to describe the collection of pulses that have the same z-coordinate, where pulse energies are typically summed.

With this idea, I use the notation E_n to indicate the summed pulse energy on the z-slice that is n pads upstream from the muon stop location. Thus, E_0 is the energy deposited on the muon stop pad z-slice and E_1 is the next-upstream z-slice. In Fig. 5.10, the energies E_0 , E_1 , E_2 are plotted, showing how these energy distributions are shaped by the penetration depth of muons into the stop pad. For some trackers, the stop pad is not necessarily the most downstream pad, and one can also consider the energies downstream from the stop, E_{-1} , etc. It often makes more sense to discuss the sum of these energies, $\sum_{n<0} E_n$, which is called the *fusion energy*, because it is most often caused by downstream-going proton tracks.

All trackers discussed here have a minimum track length cut of 3 z-slices, to avoid confusion with delayed p+t fusions and local ionization events that cross pad boundaries but do not have the range to cross into 3 separate pads.

Basic Tracker The first tracker identifies the z-coordinate of the stop with the anode pad of the most downstream pulse in the cluster. On the stop pad z-slice, the x- and y-coordinates of the stop are given by the coordinates of the highest-energy pulse. For non-fusion muon tracks, this is usually the pulse created by the stopping muon. When fusion products are present, the highest-energy pulse is often part of the ^3He or p+t tracks, causing a tracker error.

In Figures 5.11 and 5.12, the distribution of E_0 and E_1 is shown for the basic tracker. The distribution in this parameter space is complicated, so some features are summarized in the figure captions. The Bragg peak of the muon is divided between the last two stopping z-slices producing the characteristic band in the bottom left corner of the figures. Additional energy from pile-up fusion products shifts this band by either the ^3He energy (300 keV, observed) or the p+t energy (4 MeV), which is also sometimes divided between the last two z-slices. Because the most downstream z-slice is always considered the muon stop for the Basic tracker, downstream- and upstream-going end up in the same region of this plot. This




Figure 5.10: Histogram of the sum of pulse energies on the last three z-slices in muon tracks. The variation in muon energy deposition on an anode pad is driven by the penetration depth into the stop pad. The energy on the stop pad, E_0 , varies from threshold to nearly 1200 keV for muons that stop on the downstream edge of the pad. The energy E_1 is highly correlated with E_0 , because the muon Bragg peak is being divided between the last two anode pads; on the other hand, E_2 is relatively well-peaked because the muon energy deposition is nearly constant far upstream of the Bragg peak.

is discussed further with the PDir tracker.

Examining Fig. 5.11 shows that in the absence of muon-catalyzed fusion products, the energies E_0 and E_1 of all tracks fall along a band of constant $E_0 + 2E_1$. This is a consequence of the division of the ionization electrons at the end of the Bragg curve between the final two anode pads, and is useful because it provides a well-defined energy peak corresponding to the muon stop. This linear combination of E_0 and E_1 is called the S-energy,

$$S = E_0 + 2E_1, \quad (5.6)$$

and is plotted for muon tracks in the fiducial volume and the whole TPC in Fig. 5.13. The




Figure 5.11: Histogram of E_0 and E_1 for the Basic tracker. See text for description and Fig. 5.12 for some annotations of the important regions in this plot.

Basic and PDir trackers impose a minimum S-energy of 430 keV which cuts out only 0.0075 % of muons stopping in the fiducial volume.

Road Tracker The Basic tracker has several deficiencies when applied to fusion events that I will discuss in the next chapter. An alternative algorithm, the Road tracker, improves on these by using less information from the final pads that may be affected by fusion products. The proton has a range of 12.9 mm in the R2011 TPC, which is smaller than the size of an anode pad, so it is not possible for more than the last two z-slices in a muon track to be affected by fusion.

The Road tracker examines the sum of pulse energies on individual z-slices, starting at the entrance to the TPC, and finds the first slice with an energy larger than a threshold energy, $E_{Road} = 1100$ keV, or the last z-slice with a pulse if none exceed the threshold. This z-slice is then declared to be the stop z-coordinate. The set of z-slices leading up to,




Figure 5.12: Annotated version of Fig. 5.11, with zoom to the low-energy region. The muon Bragg peak is divided between E_0 and E_1 in the bottom left corner, and the fusion events with ^3He are a shifted, overlapping copy of this distribution. Events with p+t fusion deposit significantly more energy, and this can be separated into events where all of this energy is contained in E_0 , like the event in Fig. 5.5, and events where the proton travels upstream or downstream in z, like the middle event in Fig. 5.7.

but not including, the stop is defined as the “road”, which is used to determine the x- and y-coordinates. The threshold energy is chosen to be larger than the maximum E_1 from typical muon tracks (see Fig. 5.10), so it is only ever exceeded by a muon stopping in a pad, or a fusion particle. The energy deposition by protons and tritons is very large, and they are emitted back to back, so there is no track geometry that allows these particles to enter a pad and not put the energy in that pad above the road threshold. On the other hand, a muon penetrating a very small distance into a pad followed by a ^3He particle could leave the track below E_{Road} .

The y-coordinate of the stop is determined by the pulses in the road by projecting the line defined by the first and last pulse downstream to the z-slice where the muon stops. I label the y-coordinates of the first and last road pulse as y_i and y_f and the z-coordinates




Figure 5.13: S-Energy ($S = E_0 + 2E_1$) for Basic muon tracks stopping in the TPC (black curve) and in the fiducial volume (gray, filled). Events with simple muon stops fall into the peak at 2000 keV. Pile-up $n\text{-}^3\text{He}$ fusions populate a small shoulder at 2300 keV. The distribution above 3000 keV is populated by $p+t$ fusion events.




Figure 5.14: Diagram showing the muon and $p+t$ fusion geometries that cause ambiguity in all three tracking algorithms in the text. Both of these tracks leave waveform traces similar to Fig. 5.7. The Basic tracker assumes that all tracks with this geometry are upstream-going protons, selecting the downstream pad as the muon stop pad. The Road tracker assumes the opposite, and selects the next-to-last pad as the stop pad. The PDir tracker partially distinguishes the two by examining the energies E_0 and E_1 .

(z-slice numbers) of these as z_i and z_f . From these, the stop coordinate y_{stop} is given by

$$y_{stop} = \frac{y_f - y_i}{z_f - z_i} + y_f. \quad (5.7)$$

If there are multiple pulses on these z-slices, the earliest pulse is used for y_f , while the average of the pulse times is used for y_i . If the road includes only one z-slice, the y-coordinate of the largest pulse on this z-slice is used.

The x-coordinate is determined by the highest energy pulse on the last z-slice in the road. Most muon tracks do not span pad boundaries in x, but for those that do, this typically has the effect of reconstructing an x-coordinate closer to the center of the TPC.

The major advantage of the Road tracker is in the projections of the x- and y-coordinates without reliance on pads that are “contaminated” by fusion products. As mentioned above, it is crucial that the projections are used even when no z-slice exceeds the road threshold. In Chapter 7, it will be shown that this tends to treat fusion and non-fusion events nearly identically in these coordinate directions, but the z-coordinate is incorrect for many p+t fusion events. Figure 5.14 is a diagram of the remaining ambiguity due to downstream- or upstream-going protons. The upstream-going proton almost always exceeds the road threshold on the next-to-last muon track z-slice, causing that pad to be called the muon stop pad.

PDir Tracker The third tracker uses the E_0 and E_1 energy depositions to distinguish upstream- from downstream-going protons. (The name comes from **Proton-Direction**.) Downstream-going protons travel from the muon stop pad to a pad with no other energy deposition, so the last two activated pads include the muon E_0 and the full proton and triton energies. Upstream-going protons, however, travel from the muon stop pad into the pad containing the muon E_1 energy deposition, so the sum of energies on the two pads is slightly larger. In Fig. 5.15, GEANT4 simulation data is used to select events based on extreme proton angles, showing that the two bands occupy different regions of the parameter space. Including all proton angles, finite detector resolution, and other ionization effects make the distinction between proton directions less clear, and there are events with opposite proton directions with identical E_0 and E_1 . The PDir tracker selects a region of this parameter space that is more likely to be downstream-going, and assigns the stop-z coordinate to the




Figure 5.15: Events with upstream-going (left) and downstream-going (right) protons can have similar but not identical distributions in the E_0/E_1 space. The energy difference is entirely due to the inclusion of the muon E_1 for upstream-going protons. This data is from simulation, where specific proton angles to the beam axis, θ , are compared.

second-to-last z-slice (see Fig. 5.16). For all non-downstream going protons, the Road tracker algorithm is used to determine z .

Once the z -coordinate is known, the same projection techniques used by the Road tracker are used to determine x and y .

The PDir tracker is ultimately found to make fewer mistakes in muon tracking in the presence of fusion products, but the energy cut is complex and the correction for fusion interference is more difficult to calculate. For this reason, Chapter 7 focuses on the Road tracker.

5.3.5 TPC Fiducial Volume

The primary purpose of tracking with the TPC is to veto muons that stop outside of pure deuterium, either upstream in the beam pipe or windows, or in the materials of the TPC. To this end, a *fiducial volume* cut is imposed on the stop position of the muon. Stops on the outer ring of pads in x and z and in the upper or lower 15 mm along y are vetoed to avoid potential wall stops, and there is a minimum track length of two pads. See Table 5.3 for a summary of these cuts and Figs.5.17 and 5.18 for projections of the stop distribution




Figure 5.16: Visualization of the cut imposed by the PDir tracker. Events cut out of the missing region are more likely to have a downstream-going proton with the muon actually stopping on the second-to-last z-slice.

into each dimension with the fiducial volume cut indicated.

5.4 Decay Electron Tracking

Electron detection and tracking begins with forming hit clusters in the individual planes of the two cylindrical MWPCs ePC1 and ePC2, and in the barrel of plastic scintillator panels, the eSC. The hit clusters are combined to reconstruct the path of the decay electron as it travels away from the muon stop point through the MuSun detector. Electron tracks can be defined with varying coincidence requirements between the eSC and ePCs, and these different electron definitions can be paired with muon entrances to form decay-time distributions with potentially different sensitivities to background and systematic effects.

5.4.1 eSC

The fast timing of the plastic scintillator and PMTs comprising the eSC provides the high-resolution time difference between muon entrance and decay. The PMT pulses are split into

Stop coordinate	TPC dimensions (mm)	Fiducial volume (mm)	Fiducial (pads)
X	-48.0 to 48.0	-36.0 to 36.0	1 to 4
Y	0.0 to 71.0	15.0 to 56.0	n/a
Z	-64.0 to 64.0	-64.0 to 64.0	2 to 7

Table 5.3: Summary of TPC dimensions and fiducial volume cuts for TPC tracks. The pads are counted from 0 on the most upstream edge in Z, and the right-most edge in X.

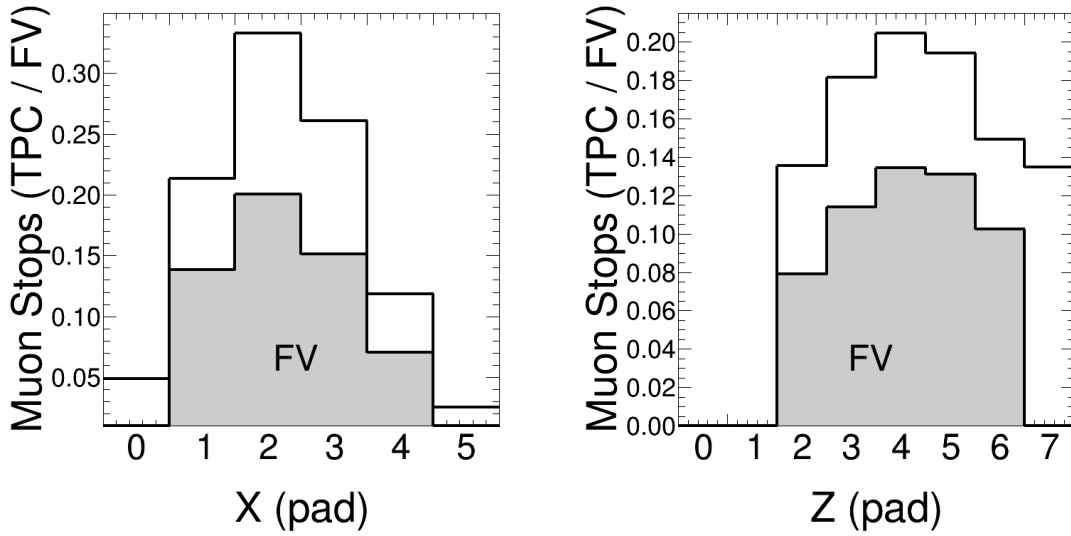


Figure 5.17: X (left) and Z (right) projection of muon stop distribution with the fiducial volume cuts applied in gray, using the Basic tracking algorithm. The population in the fiducial volume is 56 % of the total number of stops in the TPC.

a constant-threshold TDC and a waveform digitizer (WFD) with a much lower digitization threshold. From the digitized waveforms, the pulse amplitude and area are extracted using the maximum sample and integral of samples in the region around this maximum sample. The time stamp for the pulse is determined by interpolating between the two samples where the waveform goes above a threshold, giving sub-nanosecond precision.

Each of the 16 azimuthally-arranged segments in the eSC is made up of two panels, read




Figure 5.18: Y (left) and X-Z (right) projections of muon stop distribution using the Basic tracking algorithm.

out on both ends for a total of four PMT signals per segment. Pulses that are nearby in time are clustered into “N-fold” groups, where the time stamp of the group is the average of all pulses. Background particles and other non-electron sources contribute significantly to single-PMT and 2-fold hits, so typically a 3- or 4-fold coincidence is required to select predominantly through-going electron tracks.

5.4.2 *ePC*

The cylindrical wire chambers, *ePC1* and *ePC2*, provide the spatial information required to reconstruct the direction of the decay electron. Signals detected on individual anode and cathode planes of the chambers are clustered together in time and space. The average position of wires in the hit cluster and the time of the earliest hit are assigned as the time and position of this cluster. Clusters are combined into hits for each wire chamber using a time coincidence window of ± 200 ns. The anode planes of each wire chamber are not parallel to the two accompanying planes of spiral cathode strips, so an anode-cathode coincidence uniquely identifies a single point on the chamber.

5.4.3 *Electron Track Objects*

Time-coincident hit clusters from the ePC chamber planes and the eSC are combined into track objects representing an electron’s geometric path through the detector. The hits in both chambers must be within $\Delta\phi = \pm 0.35$ rad of the midpoint of the hit in the eSC segment. Tracks are formed for every combination of eSC hit cluster and combined ePC1/ePC2 cluster, so clusters can be used in multiple tracks. The track objects are flagged in this situation, and the general case is to accept only the best-matching (smallest $\Delta\phi$) for tracks that share a hit cluster in a single eSC segment. However, tracks with separate eSC clusters but the same ePC clusters are accepted.

The 4-tube eSC clusters and 6-plane ePC clusters allow for many combinations, so several standard sets of detector coincidences and cuts are defined. The first distinction is whether the eSC hits are recorded with the WFDs or the TDC channels. Each eSC hit cluster may be a 3-fold or 4-fold coincidence of PMT signals. Additionally, inclusion of each of the planes of the ePC1 and ePC2 is indicated by “AnodeOnly” for no required cathode signal, “CathodeOR” for requiring at least one cathode plane in each detector, and “CathodeAND” for a full triple-plane coincidence in each ePC. The default standard electron definition is “CathodeOR WFD 4-fold”. In addition to these, the “eSC Only” definition makes no requirement for the ePCs and only looks at hits in the scintillator panels.

5.5 *Muon Decay Events*

After forming objects representing muon entrances, TPC tracks, and decay electron tracks, they are combined into decay events. The decay time is calculated from the electron and muon entrance objects, $\Delta t = t_e - t_\mu$, for each electron track in the ± 35 μs window surrounding each muon entrance object. Histograms of this quantity are produced for the various algorithms, analysis cuts, and data quality cuts identified in previous steps (Table 5.1). These histograms are then fit for the muon lifetime using an exponential function.

An additional cut is the impact parameter of the decay, b , calculated from the distance of closest approach of the electron track to the muon stop point. This cut suppresses electron detector backgrounds from sources other than muon decay, such as cosmic-ray muons or

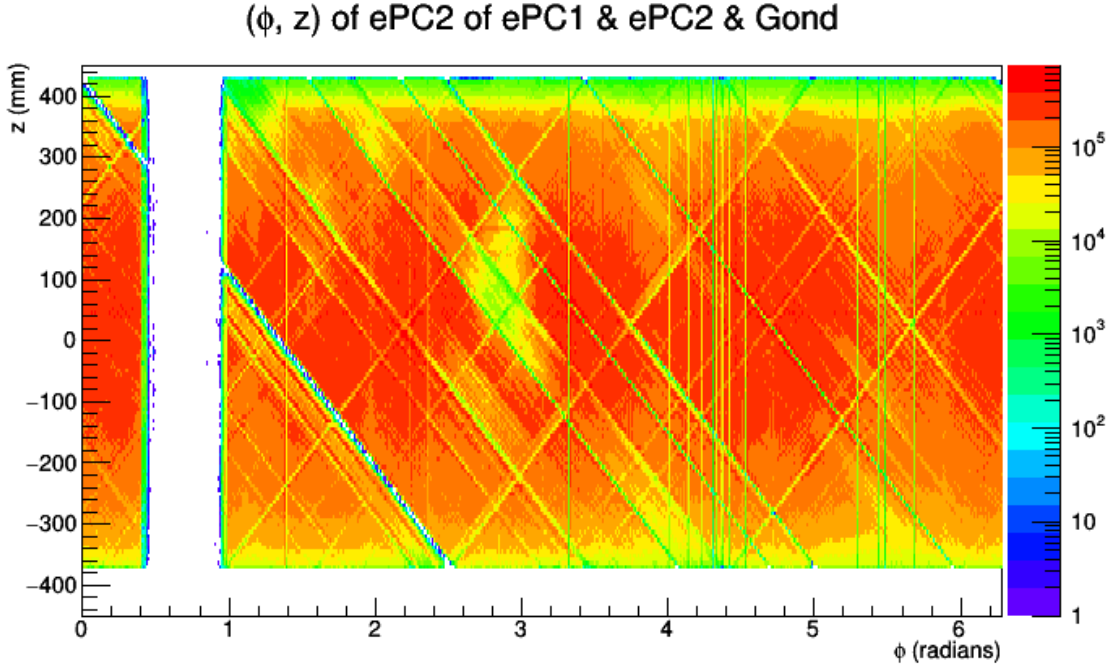


Figure 5.19: ϕ - z plot of electron tracks at ePC2, requiring a triple coincidence of ePC1/ePC2/eSC. The missing band shows the inactive region of ePC2 described in Sect. 4.4.1.

residual electron contamination in the muon beam scattering from downstream flanges. The impact parameter cut used in this analysis is

$$b \leq 120 \text{ mm.} \quad (5.8)$$

5.6 Lifetime Fitting

The analysis culminates in the generation of histograms of the decay time from accepted events (See Fig. 5.20). The muon lifetime is extracted by fitting a function derived from the full muon kinetics of Chapter 3, but for simplicity, a 3-parameter exponential fit is described here. For the small differences between the purely exponential fit and the full kinetics, see Sect. 6.4. The exponential fit function is

$$N_e(t) = N(t_0)\lambda w e^{-\lambda t} + B, \quad (5.9)$$




Figure 5.20: Fitted decay time distribution for R2011 data set. No systematic corrections have been applied, and the rate is subject to a blinding of the clock frequency. The fit residuals are displayed as a histogram and over the time range of the fit from 0.160 μs to 24.0 μs . The small oscillation in the residuals at late times is related to a background of beam electrons that can be subtracted.

where $N_e(t)$ is the number of observed decay electrons in a histogram bin of width w centered at time t , λ is the (blinded) muon disappearance rate, and B is a constant background term. $N(t_0)$ is the number of decays observed at time $t = 0$ ns. The bin width is a fixed parameter at 40 ns, selected to integrate over one period of the CAEN TDC interpolator, which is known to be non-linear.

Goodness of fit is assessed with the χ^2 statistic,

$$\chi^2 = \sum_i \frac{(H(t_i) - N_e(t_i))^2}{\sigma_i^2} \quad (5.10)$$

where $H(t_i)$ is the number of counts in the histogram bin at time t_i , $N_e(t_i)$ is the number

of decay counts predicted by the model in Eq. 5.9, and σ_i is the error in bin i , given by the $\sqrt{H(t_i)}$ because all bins are sufficiently populated that the errors are nearly Gaussian. The distribution of sums-of-squares of residuals for a set of k independent, normally-distributed variables is the Chi-squared distribution with degree k . Asymptotically for large k , the Chi-squared distribution is Gaussian with variance $2k$ and mean k . It is often easier to consider the quantity χ^2/k , called the reduced χ^2 where k is the number of degrees of freedom in the fit. The distribution of the reduced χ^2 for large k is normal with standard deviation $\sqrt{2/k}$ and mean 1. Thus the goodness of fit is assessed by comparing the reduced χ^2 of the fit to this distribution. A reduced χ^2 much larger than 1 indicates a poor fit, with additional behavior not included in the model, and $\chi^2/k < 1$ may indicate that the model includes too many parameters.

5.7 Blinding in MuSun

MuSun is a blinded experiment, meaning that the results obtained through data analysis cannot be directly converted into a measurement of the muon lifetime without additional information. Specifically, the master clock that provides the frequency standard for all detector systems was slightly detuned from the nominal value by PSI staff member Malte Hildebrandt, who is not in the MuSun collaboration. Once the data analysis is complete, the detuning offset will be revealed, and the blinding lifted. Blinding is especially valuable in experiments studying a value that has been measured or calculated before and experiments with many systematic error sources. There is always the possibility of bias towards agreement that can manifest in an unwillingness to continue searching for new sources of systematic error. Blinding the analysis until the final result is settled removes the ability to know whether it is in agreement with previous experiments. There are two blindings implemented for the R2011 data, a *hardware blinding*, where the master clock is detuned, and *software blinding*, where additional offsets are applied to data during the analysis.

The hardware blinding is a detuning of the master clock, which nominally runs at 500 MHz, to a lower frequency,

$$498.5 \text{ MHz} \leq f_{\text{R2011}} \leq 499.5 \text{ MHz} \quad (5.11)$$

The range of frequencies is smaller than 500 MHz to avoid running digitization electronics above their maximum clock frequency. Some electronics require clock frequencies of 100 MHz, provided by another clock slaved to the master clock. This slave clock was proportionally detuned to implement the blinding.

Software blinding includes an offset between μ^+ and μ^- data and offsets for each university or institution performing analysis. These offsets allow the separate unblinding of μ^+ data, which is only used for systematic checks and can be compared to the much more precise world average for the free muon lifetime. The software blinding is applied at the last stage of analysis, where lifetime fits are performed using the LtFit library. The library imposes the blinding as a part of the fit function.

The analysis presented in this document is performed with all of these blinding offsets applied. Removing the software blinding is as simple as modifying the fitting library to no longer impose the offsets. The hardware blinding is removed by scaling the measured rate by the ratio of the blind and assumed frequencies:

$$\Lambda_{real} = (f_{R2011}/500 \text{ MHz}) \times \Lambda_{blind} \quad (5.12)$$

Chapter 6

SYSTEMATIC ERRORS

Measurement of the muon lifetime to 10 ppm requires 10^{10} muon decays to achieve the required statistical precision and also careful analysis of potential systematic measurement errors. Correcting for the fusion interference systematic error is the primary analysis result of this thesis, discussed in Chapter 7. However, there are several other known sources of systematic error that are briefly outlined here.

6.1 Chemical Gas Impurities

In Sect. 3.3.6, the sensitivity to O_2 and N_2 gas impurities was quantified as $\Delta\lambda_{\mu^-} = 1.2 \text{ s}^{-1} \times \left(\frac{c_O}{\text{ppb}}\right)$ and $2.0 \text{ s}^{-1} \times \left(\frac{c_N}{\text{ppb}}\right)$, respectively, where c_N and c_O are the nitrogen and oxygen impurity concentrations. Periodic gas chromatography measurements performed during R2011 consistently measured oxygen and nitrogen impurities smaller than 1 ppb. This corresponds to a disappearance rate shift of less than 2 s^{-1} , though the consistency of the measurements was not at a level that reliably constrained the measurement uncertainty. Additionally, F. Wauters performed an in-situ measurement of nuclear recoils from muon capture on nitrogen. The capture process



results in a recoiling nucleus with atomic number $(Z - 1)$ with energy of a few hundred keV. The yield of recoils from this process is directly proportional to the concentration of impurities in the gas. The analysis was limited by the poor resolution of the TPC, causing the nitrogen capture spectrum to be largely obscured by noise and background from n- ^3He fusion. Nevertheless, a comparison with subsequent runs with upgraded TPC resolution gives a correction of $0 - 11 \text{ s}^{-1}$ for the R2011 data [59]. Fig. 6.1 illustrates the comparison of delayed TPC signal energies with and without a subsequent decay electron that determines the yield of capture recoils. Nuclear capture does not produce a decay electron, so requiring




Figure 6.1: Spectrum of delayed signals following a muon stop in the TPC in R2011 (top) corresponding to an impurity level of <4 ppb and the first run with upgraded TPC resolution in R2013 (bottom). A delayed electron requirement (black) suppresses recoils from nuclear capture relative to signals from the ${}^3\text{He}$ and $\mu^3\text{He}$ fusion products. The difference between the two spectra in the region $140 - 200$ keV is attributed to capture on gas impurities. The R2011 impurity level was substantially smaller than that of R2013, where several impurity dopings were performed to calibrate this effect. Figures and analysis by F. Wauters.

an observed electron following the TPC signal suppresses capture relative to the large MCF background.

6.2 Stops in Materials Other than Deuterium

Muons stopping in materials other than the target gas also form muonic atoms, but the capture rate on heavier nuclei is much larger, scaling as $\Lambda_Z \propto Z^4$. The MuSun TPC is an

active target designed to reject these events by tracking the 3-D muon stopping position, but nonetheless, rare events can be accepted by the tracking algorithm even though the muon stops in a material other than deuterium gas. The non-deuterium materials used to construct the TPC, pressure vessels, beam windows, and other mechanical structures are referred to collectively as “wall materials.”

A population of events that stop in a uniform material contributes an additional exponential distribution of decay electron times

$$e_Z(t) = \alpha N_\mu \lambda_0 e^{-(\lambda_0 + \Lambda_Z)t} \quad (6.2)$$

where $\alpha \equiv N_Z/N_\mu$ is the fraction of muons stopping in the material. Table 6.1 lists the fraction of accepted muon stops in ppm, $\alpha(2 \text{ s}^{-1})$, that would result in a 2 s^{-1} shift in the muon disappearance rate for carbon, silicon, and iron, using the standard fit start time of 160 ns. These elements make up the majority of the components of the TPC accessible by muons, with the notable omission of aluminum, which is similar to silicon because they differ by only one in atomic number.

Nucleus	Capture rate [s^{-1}]	$\alpha(2 \text{ s}^{-1})$ (ppm)
C	36×10^3	59
Si	850×10^3	20
Fe	4.4×10^6	48

Table 6.1: Capture rates of various materials used in the MuSun TPC and the fraction of stops necessary to cause a 2 s^{-1} shift in the measured lifetime with a fit start time of 160 ns [60].

In heavy nuclei such as Ag, W, and Au, the disappearance rate is dominated by the capture channel, and the lifetime of muons is 70 – 90 ns. Muons stopping in these materials disappear before the 160 ns start time of the lifetime fit, so they do not distort the observed decay time distribution or shift the measured disappearance rate. Some components of the R2011 TPC were constructed from steel and aluminum, such as the Frisch grid frame and

cathode, but these were replaced with tungsten and silver for the R2014 and R2015 runs.

The fraction of stops in wall materials can be studied with neutrons detected during normal operation. There are three primary sources of neutrons in MuSun: μd capture (1.2), muon-catalyzed fusion (3.5b), and capture on materials other than deuterium. Prompt neutrons will be dominated by muons stopping in wall materials, so the yield of neutrons can be plotted as a function of muon stop position to determine if there is a statistical excess near the edges of the TPC.

An analysis of neutrons in R2014 data by E. Muldoon uses the spectrum and time distribution of neutrons to determine the relative contributions of μd capture and high-Z captures as a function of the stop position within the TPC. The analysis is ongoing, but preliminary results suggest an upper limit on the fraction of muons reconstructed in the fiducial volume that actually stop in wall materials of $\alpha \lesssim 50$ ppm. Assuming the same limit in the R2011 dataset, this corresponds to a disappearance rate shift of 2 s^{-1} for stops in Fe (e.g. the Frisch grid wires and frame) and 4 s^{-1} for stops in Si (e.g. the MACOR posts and pad plane substrate). This shift is smaller than the uncertainty in the correction due to gas impurities.

6.3 *Electron Interference*

The decay electron rarely deposits enough energy in one anode pad to leave a signal above threshold in the TPC, but occasionally a small pulse is recorded. The small amount of energy can be important, however, if there is an energy threshold that determines whether or not a stop is accepted. In this case, electrons from early decays are more likely to overlap with the muon ionization cloud and add enough energy to cross the threshold. This effect has not been studied in great detail in MuSun, but some indicators have been examined. Energy thresholds that might be of importance include the pulse-finding threshold, the Road tracker threshold, and the minimum S-Energy threshold used by the PDir tracker. The pulse-finding threshold can affect the length of a track by adding small additional pulses at the end, potentially coupling to the fiducial volume cut. The overlap between electron and muon tracks depends strongly on the electron direction, so any lifetime difference for upward-going decays vs downward or horizontal decays may be evidence for this systematic. There has

been no observation of lifetime dependence on the direction of the decay electron in R2011 μ^+ or μ^- data.

6.4 *Deviation From Exponential Decay-Time Distribution*

The experimental decay-time distribution is nearly exponential, but there are small deviations. The hyperfine populations change over time and the capture rate is substantially different in the two states. Muons are usually recycled by the muon-catalyzed fusion reactions, but the muon sticks to the ^3He nucleus in the branch (3.5c) and the capture rate on ^3He is a factor of 5 larger than the doublet capture rate.

Table 6.2 summarizes the magnitude of these effects in a kinetic model of the decay time distribution [61]. The model defines a “true” capture rate, then simulates 10^{14} muon decays using the full kinetics model and the parameters listed in Table 3.2. The resulting decay time distributions are fitted with a fit function derived from the full kinetics described in Sect. 3.3.5, but with various processes modified or turned off. The fitted capture rate is then compared to the “true” rate defined in the model to determine the size of the distortion. The explored changes were: setting the (already small) quartet capture rate to zero ($\Lambda_q = 0$), setting the ^3He capture rate equal to the deuterium doublet capture rate, effectively quashing the distortion due to the sticking fusion branch (3.5c) ($\Lambda_{^3\text{He}} = \Lambda_d$), suppressing all fusions by setting the molecular formation rates to zero (“No Fusion”), and assuming a purely doublet population with no fusion, which is effectively a pure exponential with rate given by Eq. 3.1 (“Pure Exp.”).

Including the non-zero quartet capture rate, Λ_q in the fit function makes a negligible difference in the fitted rate. The build-up of a population of $\mu^3\text{He}$ atoms causes a 9 s^{-1} shift, so this must not be ignored for a 6 s^{-1} measurement. The “Pure Exp.” and “No Fusion” conditions will implicitly disable the ^3He capture channel, in addition to the contribution of a transient population in the quartet state. This means the disappearance rate shift due to the fraction of time the muon spends in the quartet state is approximately $9.4 - 4.5 = 4.9\text{ s}^{-1}$.

For each of these conditions, the same fit functions were applied to additional model data generated assuming distortions from fusion interference. Fusion interference is a systematic error in the lifetime measurement and will be described in detail in Chapter 7, but the

Fit model	Data w/o Fusion Int.		Data w/ Fusion Int.		Diff
	$\Delta\Lambda_d(\text{s}^{-1})$	Error	$\Delta\Lambda_d(\text{s}^{-1})$	Error	
Full kinetics	0.04	0.05	11.77	0.05	11.7
$\Lambda_q = 0$	0.16	0.05	12.03	0.05	11.9
$\Lambda_{3\text{He}} = \Lambda_d$	9.32	0.06	21.06	0.06	11.7
No Fusion	9.41	0.06	21.20	0.06	11.8
Pure Exp.	4.48	0.41	16.34	0.41	11.9

Table 6.2: Shift in fitted capture rate from various fit functions applied to model data generated using the full muon kinetic model. The fit start time used here is 160 ns. This analysis was performed by P. Kammel [61].

results are included here for reference. The model data uses a parameterized $\eta(t_f)$ function (Eq. 7.11) that approximates the results of simulation with an amplitude of 0.006. The distortion using any of the described fit functions is identical to within $< 0.5 \text{ s}^{-1}$, so it is not expected that the completeness of the kinetic model in the fit function has any bearing on the study of the fusion interference systematic error.

Chapter 7

FUSION INTERFERENCE

Fusion interference is the name given to a category of systematic error in the muon lifetime measurement arising from a difference in acceptance for *fusion events*, where a muon-catalyzed fusion reaction occurs following a stop in the TPC, and *non-fusion events*. I focus on fusion interference instead of other systematic issues in the R2011 analysis because the error could not be mitigated by hardware upgrades for the R2014 and R2015 production runs, and this analysis will remain relevant for those datasets¹. In this chapter, I will describe the mechanism of fusion interference, several important properties of events with fusion, a technique to correct the lifetime measurement, and suggestions for analysis strategies to improve this correction. This analysis is quite technical, combining particle physics with detector response details and the inevitable imperfection of analysis algorithms, so I will first give an overview of each section and the logical flow through the chapter.

Section 7.1: Fusion reactions convert two deuterons into charged particles that produce TPC signals that interfere with the tracking algorithm determining the muon stop position. A difference in acceptance between fusion and non-fusion events leads to a distortion of the observed decay time distribution. I introduce the differential fusion acceptance $\eta(t_f)$ to parameterize this discrepancy as a function of fusion time with respect to the muon entrance.

Section 7.2: The largest source of differential fusion acceptance is the fiducial volume cut vetoing events that would have been accepted had the fusion not caused a tracker error. The primary source of tracker errors comes from long-range protons from the p+t fusion branch in reaction 3.5, and these are dominated by errors in the z-coordinate of the stop position in the TPC. I establish several properties of tracker errors and the $\eta(t_f)$ function using MC simulation of the TPC response that justify the approach taken in the rest of the chapter.

¹For the program of upgrades inspired by R2011 analysis, see Sect. 8.1.

Section 7.4: A correction strategy is outlined based on the properties of $\eta(t_f)$. The lifetime shift is a combination of the fraction of extra events accepted and a proportionality factor that must be determined through the lifetime shifts in each of the z-slices.

Section 7.3: I develop a formalism to describe the true and observed count of muon stops in an arbitrary volume, N and \tilde{N} , and the same for fusion events, F and \tilde{F} . The migration, M , is the difference between observed and true distributions, and the fusion fraction, ϵ_f , is the ratio of F to N .

Section 7.5: The sum of the proton and triton energies (4.03 MeV) is sufficiently larger than other TPC signals to permit a low-background, high-efficiency tag identifying these events. The migration probability per fusion, $P_{mig}(z)$, is combined with the fusion tag and an interpolation of the muon stop distribution $n(z)$ to determine the fusion fraction and the number of migrating events in the z-slices and the whole TPC fiducial volume.

Section 7.6: The fractional migration and fitted lifetime of events in each of the z-slices establish the relationship between $\Delta\lambda$ and M . The final lifetime correction is calculated along with estimates of the uncertainty in the procedure. I discuss the limitations of this method and suggest improvements for the ongoing analysis of the R2014 and R2015 datasets.

7.1 *Distortion of the Observed Decay Time Distribution*

Fusion events have an average decay rate that is smaller than non-fusion events because the muon decay reaction must necessarily happen later than the time of the fusion reaction. There are several potential sources for a differential acceptance of fusion and non-fusion events, including electronic digitization thresholds and analysis cuts used in event selection. Here I introduce a mathematical formalism to describe this acceptance difference and the time distributions of relevant quantities.

For the analysis in this chapter, I consider a simplified kinetic model that ignores the difference in capture rate between the hyperfine states and ignores the distortion from the population of $\mu^3\text{He}$. This is a good approximation to the full kinetics of Chapter 3. The differential fusion acceptance $\eta(t_f)$ introduced in this chapter can be included in the kinetic model to study small perturbations [62].

Define the per-muon probability of survival to time t as $n_\mu(t)$, and its probability to

decay to an electron between t and $t + dt$ as $e(t)dt$. We have

$$e(t) = \lambda_0 n_\mu(t) = \lambda_0 e^{-\lambda_- t} \quad (7.1)$$

where λ_0 is the free muon decay rate and $\lambda_- = \lambda_0 + \Lambda_d$.

The per-muon probability of a fusion reaction is $\tilde{f}(t_f)$ and I define the conditional probability of a fusion given that the muon has survived until time t_f as $f(t_f)$:

$$\begin{aligned} \tilde{f}(t_f) &= \Pr(\text{Fusion at time } t_f) \\ f(t_f) &= \Pr(\text{Fusion at time } t_f \mid \text{Muon survives until } t_f). \end{aligned}$$

This relation can be written

$$\tilde{f}(t_f) \equiv f(t_f) n_\mu(t_f). \quad (7.2)$$

The probability of electron decay in Eq. 7.1 is the sum of the decay probabilities from events with and without a prior fusion, denoted by the superscripts f and nf , respectively,

$$e(t) = e^f(t) + e^{nf}(t). \quad (7.3)$$

To derive these probabilities, introduce as the conditional probability of an electron decay at time t given a prior fusion at time t_f :

$$\begin{aligned} e'(t, t_f) &= \Pr(\text{Decay at time } t \mid \text{Fusion at time } t_f) \\ &= e(t - t_f) \tilde{f}(t_f) \\ &= \lambda_0 f(t_f) n_\mu(t - t_f) n_\mu(t_f) \\ &\approx \lambda_0 f(t_f) n_\mu(t), \end{aligned} \quad (7.4)$$

where the last line is a consequence of the assumption that the decay time distribution is exponential. The probability of muon decay with a prior fusion is written

$$\begin{aligned} e^f(t) &= \int_0^t e'(t, t_f) dt_f \\ &= \lambda_0 n_\mu(t) \int_0^t f(t_f) dt_f \end{aligned} \quad (7.5)$$




Figure 7.1: (Left) Decay time distribution of decays from events with (blue) and without (green) p+t fusion. (Right) Decay time of fusion events on linear scale. The fusion-skimmed and fusion-vetoed decay time distributions are called $e^f(t)$ and $e^{nf}(t)$ in the text, respectively.

and the probability of decay without a prior fusion is the complement of this

$$e^{nf}(t) = e(t) - \int_0^t e'(t, t_f) dt_f \quad (7.6)$$

$$= \lambda_0 n_\mu(t) \left(1 - \int_0^t f(t_f) dt_f \right). \quad (7.7)$$

Figure 7.1 shows these two time distributions.

Accurately describing the observed fusion and electron time distributions requires that we introduce an event acceptance α , which includes the detection efficiency of the TPC and entrance detectors as well as the electron detectors. While the electron detector acceptance is unaffected by fusions, the TPC acceptance for fusion events can generally depend on the time of the fusion following the muon stop, t_f . The non-fusion acceptance, α , and overall time-dependent acceptance of fusion events, $\alpha^f(t_f)$, modifies Eqs. 7.5 and 7.6 to give the observed decay time distributions

$$e^{f,obs}(t) = \lambda_0 n_\mu(t) \int_0^t f(t_f) \alpha^f(t_f) dt_f \quad (7.8)$$

$$e^{nf,obs}(t) = \lambda_0 \alpha n_\mu(t) \left(1 - \int_0^t f(t_f) dt_f \right). \quad (7.9)$$

The observed distribution of decay electrons is the sum of these two terms,

$$\begin{aligned} e^{obs}(t) &= e^{f,obs}(t) + e^{nf,obs}(t) \\ &= \lambda_0 \alpha n_\mu(t) \left(1 + \int_0^t f(t_f) \eta(t_f) dt_f \right). \end{aligned} \quad (7.10)$$

Here I have introduced the differential fusion acceptance, $\eta(t_f)$,

$$\eta(t_f) \equiv \frac{\alpha^f(t_f) - \alpha}{\alpha}. \quad (7.11)$$

The range of the differential fusion acceptance is $-1 \leq \eta(t_f) < \infty$. The low extreme $\eta = -1$ corresponds to fusions causing all events to be rejected and the upper extreme $\eta \rightarrow +\infty$ means all non-fusion events would be rejected but fusion acceptance is non-zero. $\eta(t_f)$ is positive if fusions are more likely to be accepted, negative if less likely, and zero when there is no difference. Though any positive value is possible, in practice the acceptance of μ SC-identified muons by TPC tracking is $\sim 55\%$ and most of the losses are due to muons never entering the TPC active volume. In the study of fusion migration later in this chapter, special event selections and poor tracking algorithms lead to $|\eta(t_f)| \lesssim 0.1$ in the most extreme cases and much smaller for the final analysis cuts.

7.2 Fusion Migration and Tracker Errors

The largest contribution to fusion interference in MuSun is *tracker errors*, where fusion products influence the reconstructed stop position, causing a disproportionate number of fusion events to be vetoed by the fiducial volume cut. A tracker error occurs when the extra ionization from a fusion reaction causes a significantly different position reconstruction by the tracking algorithm. Note that this definition is referring to a discrepancy between the fusion event and a hypothetical event with exactly the same muon track, but no fusion reaction, even if the fusion event determines the stop position more accurately. An example of an event with a tracker error is shown in Fig. 7.2, where the fusion causes the Road tracking algorithm to determine that the stop z-coordinate is one pad upstream of the true stop position. Without the proton and triton ionization, the reconstructed stop pad would have been the actual muon stop pad.

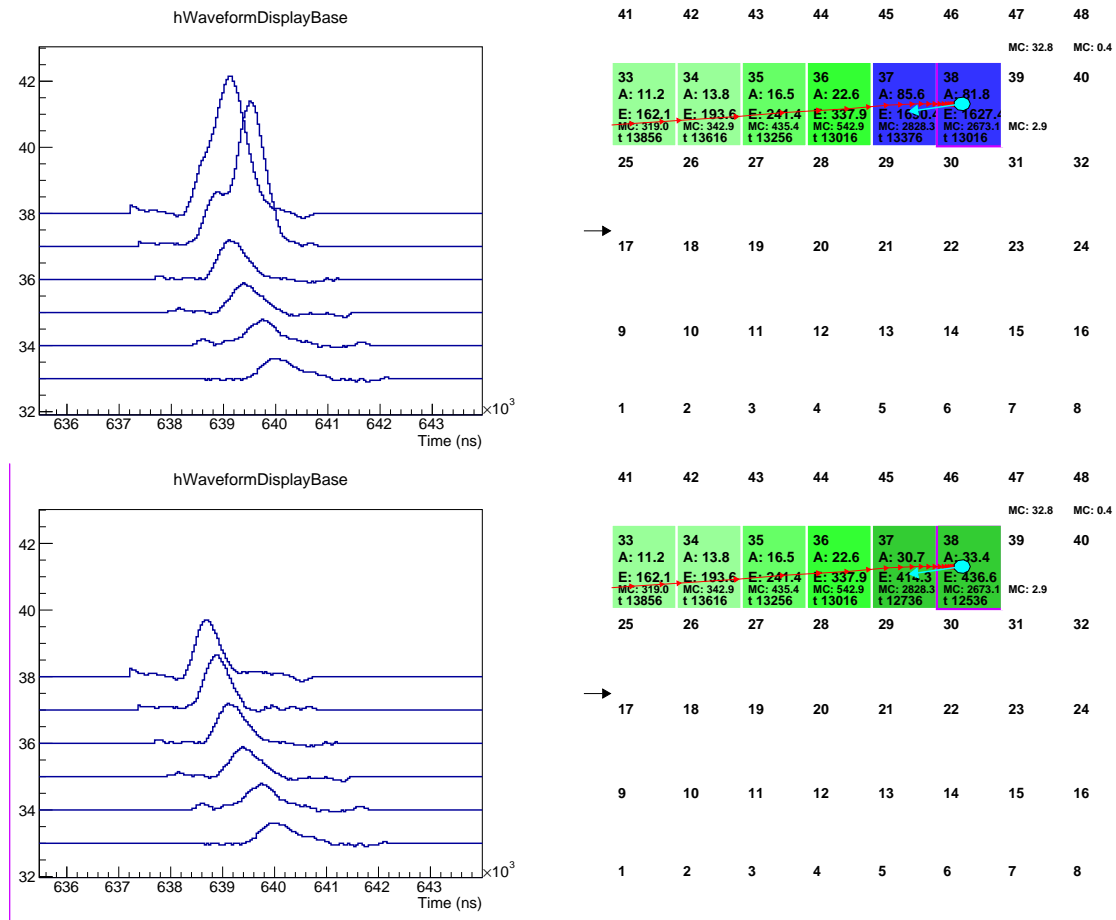


Figure 7.2: Event display of the same event in the fusion-skimmed dataset (top) and the fusion-suppressed dataset (bottom). On the pad plane display, the reconstructed energy in each pad is shown by darker green/blue pads. The muon path is marked by small red arrows and the p+t fusion path is marked by cyan circle (triton) and arrow (proton). The Road tracker determines that pad 38 is the stop pad without the fusion energy but the p+t signals put pad 37 above threshold and cause a tracker error.

I define *fusion migration* as the net effect of tracker errors on the population of stopped muon events. The magnitude of the lifetime shift from fusion interference depends on the net fusion migration into the fiducial volume, causing a non-exponential mixture of the fusion and non-fusion decay time distributions in Fig. 7.1 and thus a measured lifetime shift.

Simulation results are presented in the remainder of this section that show that the primary source of fusion migration is tracker errors in the z-coordinate caused by long-range protons. I define the function $\eta(t_f)$ for these tracker errors and explore some of its properties that will be used in the disappearance rate correction. In these studies and much of this chapter, we consider volumes of the TPC called *z-slices*, which are segments of the TPC active volume that all share a z-coordinate that were introduced in Sect. 5.3.4.

7.2.1 Monte Carlo Simulation Isolating Fusion Migration

The quantity $\eta(t_f)$ is very difficult to measure with purely experimental data, but can be extracted from a detailed Monte Carlo simulation. The MuSun simulation begins with physics events generated with GEANT4 that are passed through a stage `mcreponse` to simulate the detector, electronics, and digitization response. The resulting MIDAS files are analyzed with the standard toolchain. This analysis is concerned with fusion events, so these are skimmed before the `mcreponse` stage.

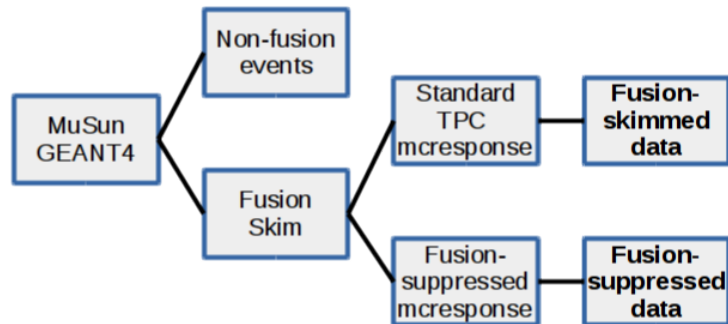


Figure 7.3: Data pipeline for fusion skim and TPC response suppression.

From the fusion-skimmed events, two datasets are produced: the *fusion-skimmed* dataset with normal `mcreponse` and the *fusion-suppressed* dataset that is treated with a modified `mcreponse` in which energy deposition from fusion products p , t , n , ${}^3\text{He}$ is ignored (see Fig. 7.3). These two datasets have identical events and kinetic time distributions, but the fusion-suppressed eliminates fusion migration. The two datasets are analyzed and compared

Dataset	Fusion-only	Fusion response suppressed	Statistics
Full MC	No	No	10^9
Fusion-skimmed	Yes	No	10^7
Fusion-suppressed	Yes	Yes	10^7
R2011 data	n/a	n/a	5×10^9

Table 7.1: Summary of simulated data sets used in this analysis. Full MC is a full-statistics set that closely mimics the R2011 data set. Fusion-skimmed and fusion-suppressed are selections of only events with fusion from the brute-force MC set. The fusion-suppressed data uses a modified TPC response that ignores the energy deposition from fusion products.

to determine changes in event categorization purely due to fusion products. See Table 7.1 for a summary of these datasets and Fig. 7.2 for an identical event in the fusion-skimmed and fusion-suppressed datasets.

7.2.2 Properties of Tracker Errors

Tracker errors can occur for all of the three tracking algorithms described in Chapter 5 (Basic, Road, and PDir), but the mechanisms and magnitude differ. Errors are characterized by the distribution of the difference in stop coordinates for events in the fusion-skimmed and fusion-suppressed datasets. For instance, the difference $z - z_{MC}$ between the true stop z -coordinate according to the Monte Carlo simulation and the reconstructed stop z is plotted in Fig. 7.4. The black distribution in the figure is from the fusion-suppressed dataset, so any discrepancy is due to non-fusion effects. The trackers only reconstruct the stop pad, not the position within the pad, so the entire region between -8 mm and $+8$ mm is populated by events where the true stop location is somewhere within the reconstructed stop pad. Contrast with the same distribution in the fusion-skimmed dataset (red), where a significant number of events have $z_{MC} > z$. Fig. 7.5 has similar plots for all three coordinate directions and all three tracking algorithms that will be referred to in the following discussion.

It is crucial to note here that the definitions of tracker errors and fusion migration do

not require that the tracking algorithm determine the correct stop position. There is no fusion interference lifetime shift from a tracker that determines the stop location incorrectly, as long as it is consistently incorrect for fusion and non-fusion events. In Figs. 7.4 and 7.5, the difference between the red and black distributions is the measure of fusion migration. For an ideal algorithm with no tracker errors, these distributions would overlap, meaning the fusion products have no net influence on the reconstructed stop position.

The Basic tracker experiences large tracker errors in all three coordinate directions. In x , the primary source of error is protons traveling to the left or right of the muon stop pad, leaving a larger pulse in a neighboring pad that is considered the muon stop. The y -coordinate is determined by the largest pulse on the stop z -slice. The proton and triton are much more energetic than the muon, so the corresponding pulses determine y . The TPC conflates distance in y with delays in time, so delayed fusions or protons traveling along the drift direction cause a wide range of y -coordinate discrepancies. Errors in z are due to the fusion proton traveling downstream of the muon stop and activating the next pad. The large magnitude of errors in all coordinates make this tracker unsuitable for the lifetime analysis, so it will not be discussed further.

The Road tracker was designed to minimize tracker errors, especially in the x - and y -coordinates, and this is reflected in the simulation data. As described in Sect. 5.3.4, nearly all of the tracker errors experienced by the Road tracker are due to upstream-going protons triggering the Road threshold and causing a discrepancy in the z -coordinate. There is almost no discrepancy between the fusion-skimmed and fusion-suppressed distributions for the x - and y -coordinates, because the pads affected by fusion are not used to determine these coordinates. The very slight difference between the y -distributions is due to muons that undergo a high-angle scatter in the last z -slice in their track, stopping several mm above or below the scattering site. Accompanying this is a upstream-going proton that triggers the Road threshold in the pad before the muon stop, so the pulses from the scatter pad are not used. Note that this implies that there is also a tracker error in the z -coordinate for these rare events.

The PDir tracker uses the same projection technique as the Road tracker, and the x - and y -coordinates have very similar levels and distributions of tracker errors. The discrep-




Figure 7.4: Discrepancy between MC truth and Road tracker reconstructed stop z-coordinate for the fusion-skimmed (red) and fusion-suppressed (black) datasets. The black, fusion-suppressed distribution extends nearly-uniformly from -8 mm to +8 mm, reflecting the single-pad resolution of the Road tracking algorithm. Since this data does not include the TPC response to fusions, the tracker gets the correct stop pad nearly 100 % of the time. The red, fusion-skimmed data does include the TPC response to fusions, so tracker errors are more common. The tail to the left side between -18 mm and -8 mm corresponds to protons traveling in the -z direction, causing the reconstructed stop position to be one pad upstream of the truth. This is less likely for muons that stop deep into the pad and more likely if the muon stops near the upstream edge of a pad. The extreme tails for $Z - Z_{MC} < -18$ mm and $Z - Z_{MC} > 15$ mm are due to fusion neutrons from the $n+{}^3\text{He}$ fusion branch scattering off of a deuteron, triggering the Road threshold. On the left side, the scattering occurs upstream from the muon stop. On the right side, the scattering occurs downstream from the muon stop, and has the gap-peak shape because this tracker error can only occur if the muon itself doesn't trigger the road threshold, so it must have stopped a short distance into a pad.




Figure 7.5: Muon stop position reconstruction errors for each stop coordinate and tracking algorithm for the fusion-skipped (red) and fusion-suppressed (black) datasets.

ancy in the z-coordinate is reduced compared to the Road tracker because the energies E_0 and E_1 are used to discriminate downstream- and upstream-going protons. This procedure is not perfect, but does reduced the overall magnitude of z-coordinate error by $2/3$, at a cost of distributing these errors in both the upstream and downstream directions. The additional complexity of bi-directional tracker errors and somewhat complex energy cuts make modeling the migration with the tracking algorithm too difficult, so the Road tracker is used for the rest of the discussion in this chapter.

7.2.3 Fusion Time Dependence of $\eta(t_f)$

The Monte Carlo simulation is also used to determine the time dependence of the differential fusion acceptance $\eta(t_f)$ (Eq. 7.11). Fusion migration is the net effect of all tracker errors leading to events being counted as part of the fiducial volume (FV) (defined in Table 5.3). This can be broken down into uni-directional migration across each of the individual faces of the rectangular prism defining the FV. Events are categorized by their reconstructed stop position and the simulation truth position, grouped into several volumes as seen in the examples in Fig. 7.6. The format for the categorization lists the coordinate direction, then “High” or “Low” for the upper or lower boundary along that coordinate axis, and then “IN” or “OUT” indicating the direction from the MC stop to the tracker reconstructed stop point. For example, **X Low-IN** means the true stop was outside the FV on the -x side, but the reconstruction was inside the FV. There are two additional categories, **BothFV** and **NeitherFV** for the cases where both the simulation and reconstruction agree.

This categorization is performed separately for the fusion-skimmed and the fusion-suppressed datasets as a function of the fusion time t_f to yield the time distributions of events in each category, for both datasets: $N_{Z\text{-High-IN}}^{\text{skimmed}}(t_f)$, $N_{Y\text{-Low-IN}}^{\text{suppressed}}(t_f)$, etc. These time distributions are all likely to be non-zero, since the trackers do not perfectly reconstruct the stop position, but not all of these mistakes are due to fusion. The time distributions of each of the categories are subtracted and normalized by the time distribution of all fusions to get the normalized per-fusion probability of tracker errors for each category,

$$n_{Z\text{-High-IN}} = \frac{N_{Z\text{-High-IN}}^{\text{skimmed}}(t_f) - N_{Z\text{-High-IN}}^{\text{suppressed}}(t_f)}{\tilde{f}(t)}. \quad (7.12)$$

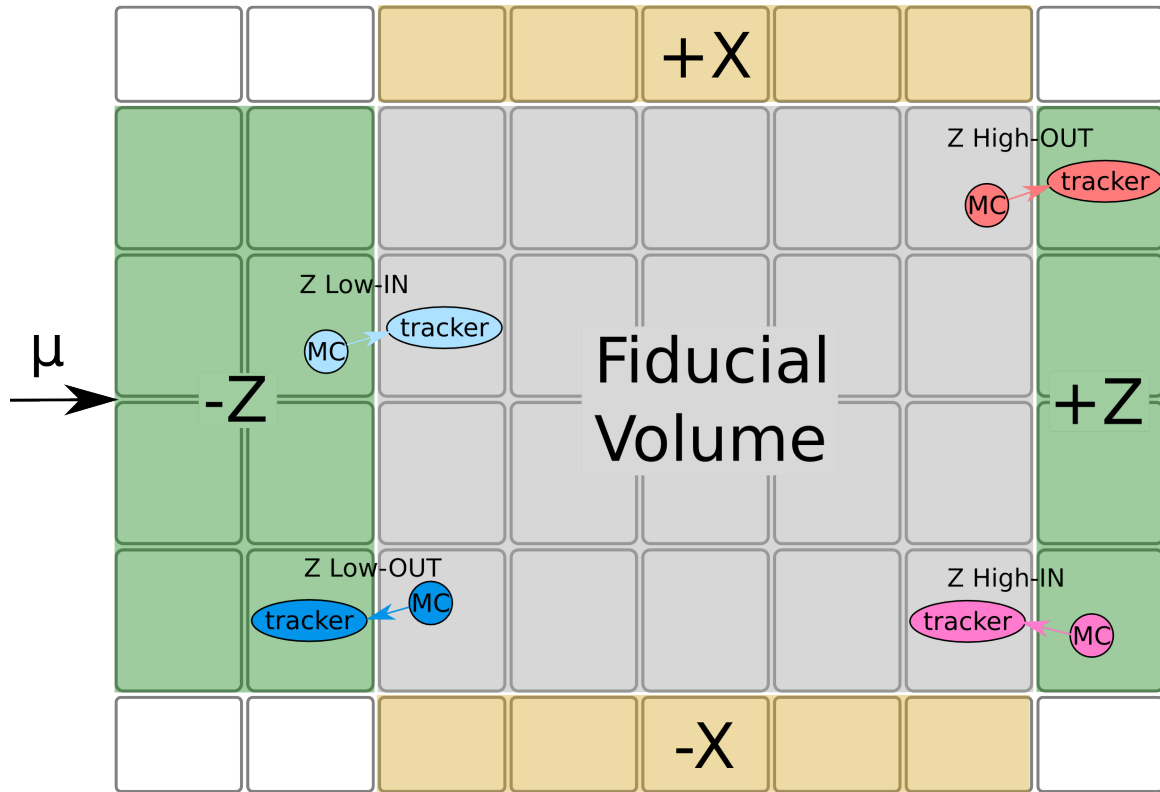


Figure 7.6: Diagram of tracker error categorization with the muon beam entering from the left. A slice of the TPC in the x-z plane is shown with the fiducial volume labeled as well as the 4 regions just outside the FV ($-Z$, $+Z$, $-X$, $+X$). Tracker errors are categorized by the location of the true stop location according to MC simulation and the reconstructed tracker stop location, which are colored ellipses in the figure. There are four categories of error shown. For example, **Z High-IN** means the true stop is outside the FV but the reconstructed stop is inside the FV, and the migration is inward across the $+Z$ boundary (“Z High”).

Distributions of tracker errors for each of the coordinates and fiducial boundaries and directions of migration are shown in Fig. 7.7 for events with p+t fusion. Although not broken down by coordinate, the distributions for ^3He fusion are plotted in Fig. 7.8. Nearly all of this migration is in the z-direction.

The category **StopFV** is the net change in the number of events reconstructed inside

Category	Description
Z Low-IN	MC stop too far upstream (-z), reconstruction in FV
Z Low-OUT	MC stop in FV, reconstruction too far upstream (-z)
Z High-IN	MC stop too far downstream (+z), reconstruction in FV
Z High-OUT	MC stop in FV, reconstruction too far downstream (+z) ... same for x, y
Z IN-OUT	(Z Low-IN) + (Z High-IN) - (Z Low-OUT) - (Z High-OUT) ... same for x, y
BothFV	MC stop in FV, reconstruction in FV.
NeitherFV	MC stop outside FV, reconstruction outside FV.
Net IN	MC stop outside FV, reconstruction in FV.
Net OUT	MC stop in FV, reconstruction outside FV.
StopFV	reconstruction inside FV (MC stop ignored)

Table 7.2: Categorization of tracker errors in the GEANT MC simulation, comparing tracker reconstruction and the known simulated stop position. “BothFV” and “NeitherFV” indicate agreement between simulation and reconstruction, and the remaining categories are based on the nature of the discrepancy with respect to the fiducial volume boundary.

the FV. This distribution is identified with the differential fusion acceptance, $\eta(t_f)$, due to the fiducial volume cut,

$$\begin{aligned}
 \eta(t_f) &= n_{\text{StopFV}}(t_f) \\
 &= n_{\text{Net IN}}(t_f) + n_{\text{BothFV}}(t_f).
 \end{aligned}
 \tag{7.13}$$

Note that all events must be placed one of these categories for both the fusion-skimmed and fusion-suppressed datasets. The categories for the two datasets are subtracted, so the sum of all categories must be zero,

$$\sum_i n_i = 0.$$

This means not all the categories in Table 7.2 are independent. For instance, the **StopFV**

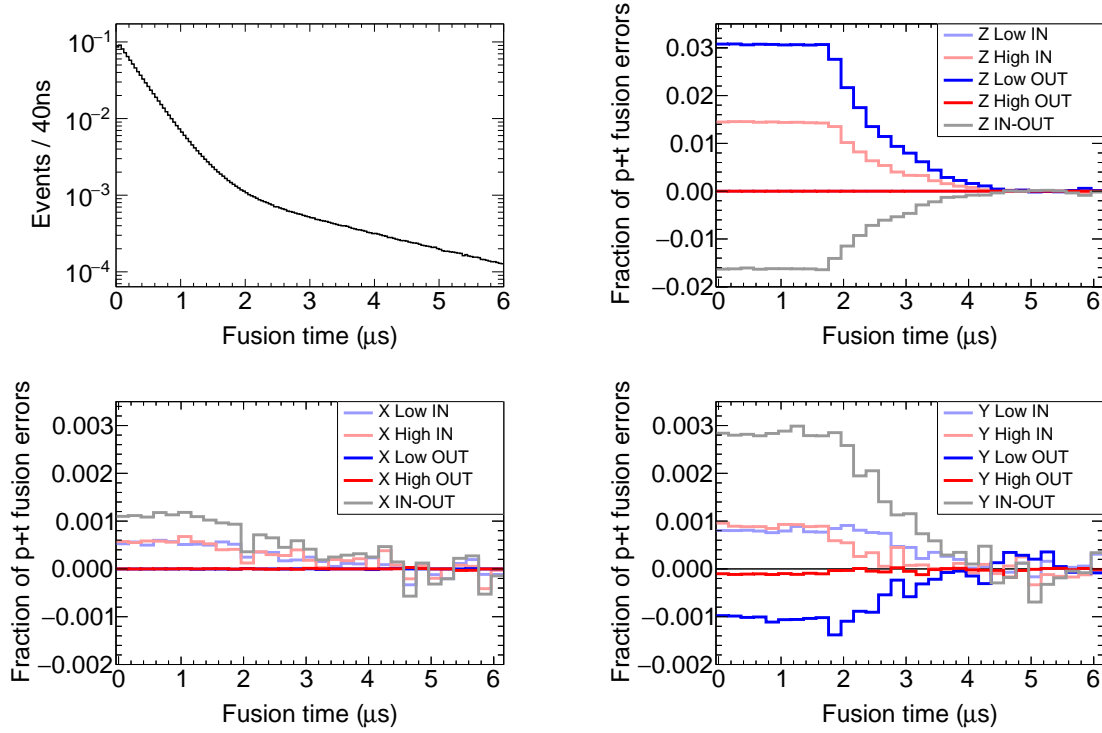


Figure 7.7: (Upper left) Time distribution of p+t fusions. Most fusions occur in the first 2 μs . The remaining panels show the fraction of fusions that induce tracker errors in and out of each of the 6 faces of the rectangular prism that defines the FV by time of the fusion reaction. The dominant contribution is across the z-boundaries (note the factor of 10 difference in the vertical scale for the x and y errors). Each bin at a given fusion time indicates the fraction of fusions occurring at that time that cause reconstruction errors of that category. The categories of migration are described in Table. 7.2 and the functional form is explained in Fig. 7.9. Negative values indicate that fusion induces fewer of that type of reconstruction error.

category is the difference of net inward and net outward migration:

$$n_{\text{StopFV}} = n_{\text{Net IN}}(t_f) - n_{\text{Net OUT}}(t_f).$$

The proton fusion migration probabilities in Figs. 7.7 and 7.8 show a similar functional form for all boundaries of the fiducial volume in each coordinate direction, though the mag-

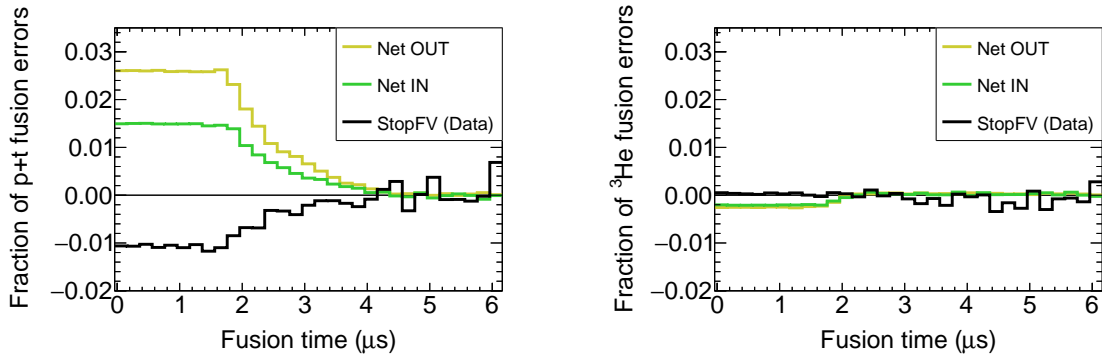


Figure 7.8: (See Fig. 7.7 caption.) Net in-migration and out-migration due to p+t (left) and ${}^3\text{He}$ (right) fusion-induced tracker errors. The category labeled **StopFV** is the difference of **IN** and **OUT** and is the differential fusion acceptance $\eta(t_f)$ (see text). Helium fusion induced errors events have time dependence that is similar to p+t but without the tail after 2.0 μs , and a much smaller overall magnitude. Nearly all of this migration is in the z-direction.

nitude of migration varies substantially. The shape is understood as the effect of signal clustering in the tracking algorithms. The probability of overlap between the muon and fusion signals decreases with increasingly delayed fusions. See the illustration and explanation in Fig. 7.9.

The magnitude of the fusion migration varies substantially, with net migration across the z boundary dominating the x or y fiducial boundaries. The specific numerical values are not directly applicable to experimental data, since the total migration depends on the detailed shape of the stop distribution and the selection of a fiducial volume. However, the relative magnitude of the migration across individual boundaries is indicative of which are likely to be the largest problems.

Since the net migrations can be small due to cancellation of large migrations across individual boundaries, the magnitude of uni-directional migrations should be compared. These indicate that z migration occurs 5-10 times more often than y migration and 15-30 times more often than x migration.

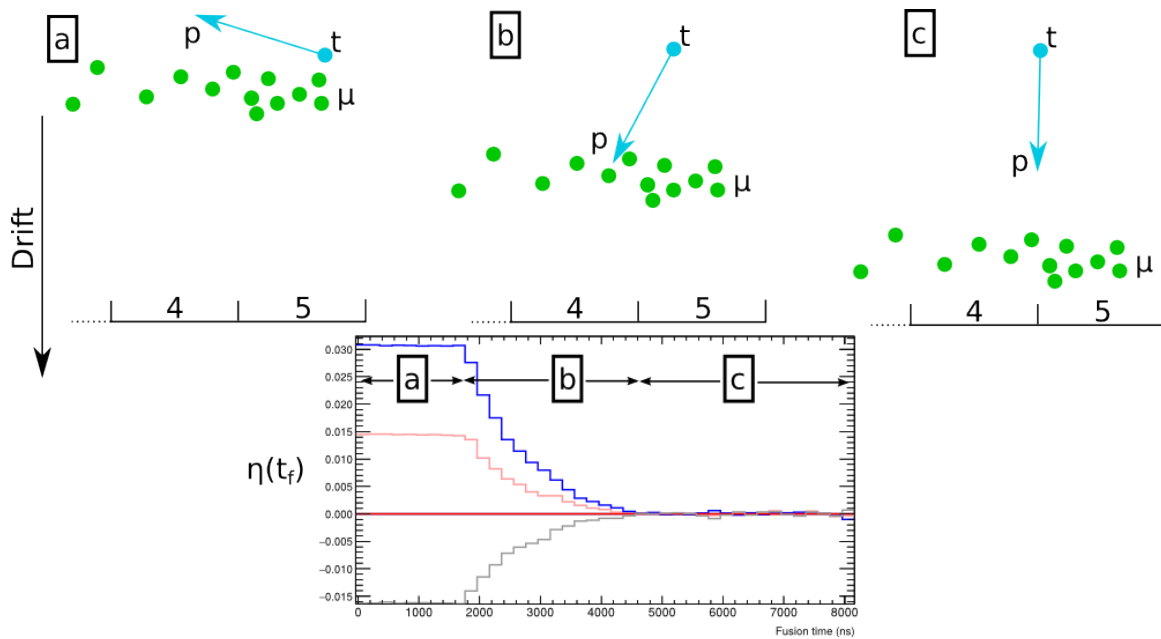


Figure 7.9: Event topology diagram leading to the shape of $\eta(t_f)$ for p+t fusion events. Each of the upper diagrams contains the drifting ionization cloud from a muon that stopped in pad 5. From left to right, the time of the p+t fusion is increasingly delayed relative to the muon stop. Each diagram corresponds to a distinct region in the lower plot of $\eta(t_f)$. (a) The triton overlaps with the muon ionization, and the proton energy deposition in pad 4 is enough to trigger the Road threshold, so there is a tracker error with the reconstructed stop in pad 4. (b) A somewhat delayed p+t-fusion will only overlap for tracks where the proton goes downward to catch up with the drifting ionization cloud. Downward-pointing protons have less solid angle to cross over the boundary into pad 4. (c) The fusion is delayed enough that even a downward-going proton cannot catch up to the muon ionization. The tracker finds separate clusters for these signals, and there is no fusion interference ($\eta = 0$).

Migration in the z-direction is only upstream (categories **Z High-IN** and **Z Low-OUT**). This confirms the results of Fig. 7.5 and the nature of the Road tracking algorithm that p+t fusion induced migration is upstream.

The time dependence of the helium fusion migration probability in Fig. 7.8 is also due to pulse clustering. Since the ${}^3\text{He}$ is short-ranged, clustering does not occur beyond 1.7 μs .

The magnitude is 10-15 times smaller than z migration from protons.

The remainder of this chapter discusses fusion migration in the z-direction specifically from p+t fusions. The other coordinates and ^3He migration are not completely negligible, and should be the subject of further study.

7.2.4 Separation of $\eta(t_f)$ Shape and Magnitude

The normalized time dependence of the fusion migration is the same for all fiducial volume boundaries in Figs. 7.7 and 7.8, despite considerable variation in the migration magnitude. This motivates factoring the time-dependent migration into a time-independent magnitude and a function defining the time dependence. For a volume V_j such as the FV or a single z-slice,

$$\eta_j(t_f) = \eta_j \tilde{\eta}(t_f). \quad (7.14)$$

The factor η_j is proportional to the per-muon probability of migrating into the volume V_j , and the function $\tilde{\eta}(t_f)$ is identified with the probability of clustering with the muon track for protons that cause tracker errors. At early times, all tracks are clustered ($\tilde{\eta}(0) = 1$) and long-delayed fusions are never clustered ($\tilde{\eta}(t_f \rightarrow \infty) = 0$). The intermediate behavior is explained in Fig. 7.9.

The hypothesis that $\eta(t_f)$ has a fixed shape for net proton migration is tested with the distribution of net migration for sub-volumes of the TPC FV, such as z-slices. Net migration into the FV (**StopFV**) for individual z-slices 2 through 5 is shown in Fig. 7.10. The right panel contains the same distributions, but unit-normalized to emphasize that the functional form does not depend on the z-slice, and the parameterization in Eq. 7.14 is justified for individual z-slices.

7.2.5 Lifetime Shift from Net Migration

The separation of the time dependence and migration magnitude suggests that the probability of fusion migration can be determined largely independent of the time of the p+t fusion, as long as we consider only fusions that are clustered with the muon track. Substitute the factored Eq. 7.14 into the observed decay time distribution, Eq. 7.10, for decays from a




Figure 7.10: $\eta(t_f)$ function for each of the z-slices in the fiducial volume (left) and normalized (right). Positive η indicates net fusions gained and negative values indicate net losses. Figures by D. Salvat.

volume V_j ,

$$e_j^{obs}(t) = \lambda_0 \alpha n_\mu(t) \left(1 + \eta_j \int_0^t f(t_f) \tilde{\eta}(t_f) dt_f \right). \quad (7.15)$$

The total probability of a p+t fusion reaction occurring in an event is $\sim 3\%$, and the amplitude η_j is on the order of 0.1 for situations considered in this chapter. Thus, the second term in parenthesis is a small perturbation on the undisturbed decay time distribution and the resulting lifetime shift is linear in η_j to good approximation:

$$\Delta\lambda_- \propto \eta_j. \quad (7.16)$$

This is verified empirically by vetoing increasing fractions of fusion events and fitting for the lifetime, or with a simple Monte Carlo or analytical simulation of the muon kinetics. To set the scale for the size of a fusion interference effect, these studies indicate that a constant $\eta(t_f) = 0.005$ leads to a disappearance rate shift of $\sim 16 \text{ s}^{-1}$, and a more realistically-shaped $\eta(t_f)$ function leads to $\sim 10 \text{ s}^{-1}$. A fit start time of 160 ns is used to be consistent with the data analysis, but a later start time will suppress the error and correction.

The proportionality constant in Eq. 7.16 generally depends on the $\tilde{\eta}$ function and the fit

start time,

$$\Delta\lambda = \eta_j \times k[\tilde{\eta}, t_{\text{start}}]. \quad (7.17)$$

7.3 Migration Formalism

The time dependence of the $\eta(t_f)$ function has been shown to be independent of the choice of particular sub-volume in the TPC. The magnitude η_j in Eq. 7.14 is proportional to the per-muon probability of an event migrating into volume V_j , which could be a z-slice or the fiducial volume, for instance. We will define some theoretical distributions that distinguish between events as they are reconstructed and events as they *would have been reconstructed* in the absence of fusion-induced tracker errors. The distribution of muons reconstructed in V_j is \tilde{N}_j . This includes events where a fusion-induced tracker error causes an erroneous reconstruction in V_j . The distribution N_j is the number of stops that would be reconstructed in V_j in the absence of fusion-induced tracker errors. Define the analogous distributions for events in which a fusion reaction occurs as \tilde{F}_j and F_j . These and other quantities that will be used in this chapter are summarized in Table. 7.3.

Throughout this section, refer to Fig. 7.11 for illustrations of these distributions and the relationship between them. The number of events migrating due to fusion, M_j , is the difference in distribution that accounts for fusion migration and the distribution that ignores the effect of fusion-induced tracker errors,

$$M_j \equiv \tilde{N}_j - N_j = \tilde{F}_j - F_j. \quad (7.18)$$

This migration is proportional to the disappearance rate shift for events in z-slice j ,

$$\Delta\lambda_j \propto \frac{M_j}{N_j}. \quad (7.19)$$

This ratio is also proportional to the quantity η_j in Eq. 7.14. Fusion migration due to the Road tracker is nearly all upstream (see Sect. 7.2.2), so the *net* migration can also be described as the difference between uni-directional migration into and out of the z-slice,

$$M_j = U_{j+1} - U_j. \quad (7.20)$$

Quantity	Definition
V_j	Either the FV or a single z-slice of the TPC
N_j	muon stops in the volume V_j , absent fusion migration
\tilde{N}_j	muon stops in V_j , including fusion migration
F_j	muon stops w/ clustered p+t fusions in V_j , absent fusion migration
\tilde{F}_j	muon stops w/ clustered p+t fusions in V_j , including fusion migration
M_j	Net change in muon stops in V_j caused by fusion migration
ϵ_f	fraction of events with a clustered p+t fusion
U_j	Uni-directional upstream fusion migration from z-slice j to $j - 1$
$n_j(z)$	muon stop density within z-slice j
$f_j(z)$	analogue of $n_j(z)$ counting only events w/ clustered fusion
$P_{mig}(z)$	prob. of upstream migration by muon penetration depth
CluE9	sum of pulse energies near the muon stop used to tag p+t fusions

Table 7.3: Definitions for quantities used in this chapter. The distributions with a tilde are distinguished from the non-tilde distributions by hypothetically turning the fusion-induced tracker errors on or off. Non-tilde distributions are the “true” number of events, as if there were no fusion-induced tracker errors. The tilde distributions include the effect of these tracker errors. This means \tilde{N}_j and \tilde{F}_j are directly observable by counting reconstructed stops and stops with a CluE9 tag, respectively, but all other quantities must be modeled or derived from these.

Here, the quantity U_j is the number of fusion events migrating upstream from the j th z-slice. For the Road tracker, the U_j are all positive, but the net migrations M_j may be negative.

All of these quantities with index j describe individual z-slices, but analogous quantities are used for the whole fiducial volume: N_{FV} , F_{FV} , etc. The number of muons and fusion events in the FV is the sum of the z-slices, $N_{FV} = \sum_j N_j$ (and likewise for \tilde{N} , F , and \tilde{F}). In the same vein, the net migration is simply the sum of z-slices, but this can also be written in terms of the upstream migration into the last z-slice and the migration out of the first




Figure 7.11: Diagram of event distributions relevant for fusion migration. (Left) \tilde{N}_i is the observed number of stops in a given bin i . N_i is count of stops in that bin as if there were no fusion-induced tracking errors. (Right) F_i and \tilde{F}_i are defined analogously for fusion events. The difference between the observed stops and stops ignoring fusion migration in a bin is the net fusion migration into the bin: $M_i \equiv \tilde{F}_i - F_i = \tilde{N}_i - N_i$. The small fraction of fusion events makes N and \tilde{N} indistinguishable and F difficult to see on the left plot. The fusion fraction, ϵ_f , is the ratio between the number of fusion stops and muon stops in a bin, $F_i = \epsilon_f N_i$.

z-slice. The FV comprises the z-slices j , for $j \in [2, 6]$, counting the first z-slice as zero, so the migration into the FV is from z-slice 7 and the migration out is from z-slice 2,

$$M_{FV} = \sum_j M_j = U_7 - U_2. \quad (7.21)$$

The fusion fraction, ϵ_f , is the ratio between the number of fusions and muon stops and does not depend on the choice of volume,

$$\epsilon_f \equiv \frac{F_j}{N_j}. \quad (7.22)$$

It can be written in terms of the observed muon stops and fusions as

$$\epsilon_f = \frac{\tilde{F}_j - M_j}{\tilde{N}_j - M_j}. \quad (7.23)$$

The fusion fraction is an important quantity because if it is known, the fractional migration can be determined from just the observed muon and fusion distributions

$$\begin{aligned}\frac{M_j}{N_j} &= \frac{\tilde{F}_j - F_j}{N_j} \\ &\approx \frac{\tilde{F}_j}{\tilde{N}_j} - \epsilon_f.\end{aligned}\tag{7.24}$$

Here we use the approximation that $\tilde{N}_j \approx N_j$, since fusions are a small fraction of events, and migrations are only a fraction of fusion events.

7.4 Correction Strategy

I have characterized the fusion-induced tracker errors of the Road tracking algorithm, concluding that the migration is dominantly in the z-direction and that the resulting lifetime shift is proportional to the ratio of migrating fusion events to total events in the TPC fiducial volume. In the remainder of the chapter, the disappearance rate correction for the R2011 dataset will be determined by the following procedure:

1. Identify p+t fusion events \tilde{F}_i in the data with a tag based on the large total energy of the proton and triton (4.03 MeV).
2. Calculate the probability of migration as a function of the penetration depth of muons within a z-slice, $P_{mig}(z)$.
3. Approximate the continuous muon stop density with a gaussian function $n(z)$ derived from the stops in each z-slice, N_i . The difference between the N_i and the integral of $n(z)$ is ± 1 %.
4. Combine the migration probability and the stop density to determine the upstream migrations from each z-slice, $U_i = \int \epsilon_f n(z) P_{mig}(z) dz$ and the net migration $M_i = U_{i+1} - U_i$.
5. Correct the observed counts in each z-slice (\tilde{N}_i and \tilde{F}_i) by subtracting the net fusion migration, M_i . Since N_i and $n(z)$ are unobservable, this correction is performed iteratively, where N_i and F_i are first approximated by \tilde{N}_i and \tilde{F}_i to derive M_i .
6. Determine the constant k in Eq. 7.17 from the lifetime of events in individual z-slices.

Each z-slice has a different fraction of migrating events and therefore a proportionally different lifetime shift.

7. The disappearance rate correction in the FV is given by $\Delta\lambda = kM_{FV}/N_{FV}$.

The remaining sections describe each of these steps in detail, including an estimation of the uncertainty in the correction. There is an additional complication that the CluE9 tag is unavailable for the full R2011 dataset, and the fusion fraction must be determined with a partial dataset. This is discussed in Sect. 7.5.4.

7.5 Modeling Fusion Migration

In this section, a tag for fusion events is developed based on total pulse energy and the distributions listed in Sect. 7.3 are determined with a model of tracker error probability and the muon stop distribution. These are combined into a procedure to extract the fusion fraction from a subset of the R2011 data and determine the net fusion migration into individual z-slices and the fiducial volume as a whole.

7.5.1 Fusion Event Tagging

A method for tagging events that contain a p+t fusion is a prerequisite to modeling fusion migration in the R2011 data set. The energy released by muon-catalyzed fusion is substantial: 3 MeV for the proton and 1 MeV for the triton. The largest energy deposition from a muon stopping in an anode pad is just over 1 MeV, so the total energy deposition is a reasonable method of tagging p+t fusions. To this end, define the quantity *CluE9* as the sum of all pulse energies on the stop pad and the 8 surrounding pads, considering only pulses in the muon track cluster. Protons have a range that is smaller than the length or width of a pad, so their energy deposition is necessarily contained in this selection of pads. The distribution of E9 in the simulation data set is shown in Fig. 7.12. The quantity E9 is similar to CluE9, except all pulses in the event are included, not just clustered pulses. There is a relatively clean separation between the events with no fusion or only ^3He fusion and those with a p+t fusion (red). Based on this separation, fusion events are defined as those having $\text{CluE9} > 3.0 \text{ MeV}$.




Figure 7.12: (Left) Distribution of E9 energy sum in GEANT simulation and (Right) CluE9 distribution in a subset of the R2011 dataset. The orange line indicates the CluE9 > 3.0 MeV tag that identifies p+t fusions. The black distribution includes all others and additionally events with multiple fusions. Differences between E9 and CluE9 come from fusions delayed beyond the time window in which TPC pulses are clustered, which are a small fraction of all fusion events. The GEANT simulation over-estimates ion-electron recombination leading to a left-side tail for p+t fusions, meaning the simulation-based efficiency of the CluE9 cut is underestimated. Left figure created by D. Salvat.

Efficient identification of fusion events in the fiducial volume is critical for determining the migration correctly and it is important that the CluE9 cut efficiency not be time-dependent. Using simulation data, the energy-based tag is compared to the true number of events with p+t fusion and only 0.2 % of true fusion events in the fiducial volume are missed by the CluE9 tag with no strong dependence on the time of the fusion (see Fig. 7.13). The CluE9 cut can falsely categorize events as p+t fusions when no p+t fusion occurred, most commonly because the 2.45 MeV neutron from the ^3He -n fusion branch scatters on a deuteron. These background events amount to 1 % of events tagged with CluE9.

The effect of fusion migration in the z-direction is apparent when the E9 and CluE9 tags are used select events. Figure 7.14 shows the ratio of fusion-tagged events to total muon stops in each of the fiducial volume z-slices. The upstream slices have a large net excess of events and the downstream slices have fewer, consistent with substantial upstream migration. The effect of this migration is to substantially shift the lifetime measured by




Figure 7.13: (Left) Fraction of ^3He events tagged by the E9 fusion in the fiducial volume. (Right) The inefficiency for tagging p+t fusion events in the fiducial volume has no strong time dependence, remaining $\lesssim 0.2\%$. There is little data after 1.5 μs because most fusions occur before this time. Figures created by D. Salvat.

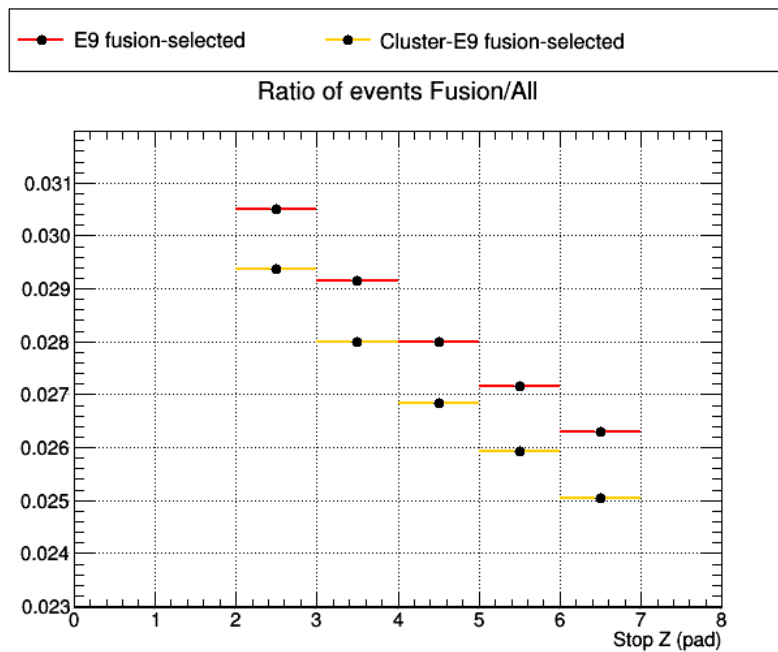


Figure 7.14: Observed fraction of events in each z-slice in the fiducial volume.




Figure 7.15: (Left) Scan of fitted decay rate versus z-slice in the TPC using simulated data. (Right) Fit residuals for the second z-slice. The early-time deviation is due to fusion interference.

events in individual z-slices. The simulated decay time distributions are fit for each z-slice and the disappearance rate fit parameter is shown in Fig. 7.15. Excess fusion events cause the average disappearance rate to become smaller for upstream z and missing fusion events cause a higher rate for downstream z-slices.

7.5.2 Migration Probability Based on Muon Penetration Depth

Fusion migration for the Road tracker is primarily due to long range protons traveling upstream and pushing the next-to-last z-slice in the muon track above the threshold energy. A secondary source of migration is muons that barely penetrate into a z-slice so that the energy deposition is below the pulse-finding threshold of ~ 150 keV. A ^3He fusion particle will deposit enough energy to bring this pulse above threshold and change the reconstructed stop. In this section, I present an analytical model for the probability for fusion migration as a function of muon stop depth within a z-slice, $P_{mig}(z)$. The model is compared with the GEANT4-based MC simulation.

Analytical Model

The mechanism of tracker errors in the z -coordinate due to protons can be explained with a simple geometric argument plus additional consideration of energy thresholds. The simplest model of proton migration assumes migration for any event in which the proton stops in a pad one upstream from the muon stop, regardless of energy deposition and thresholds. Protons are emitted isotropically from the $d\mu d$ molecule and the range straggling in deuterium is small, 3-4 %. The endpoints of proton tracks uniformly populate a spherical shell with radius equal to the proton range in deuterium, $R_p = 12.9$ mm. For a muon stopping at a depth z into the pad, the fraction of protons crossing the boundary is the ratio of the surface area of a spherical cap ($A^{\text{cap}} = 2\pi Rh$) of height $(R_p - z)$ to the surface area of the spherical shell,

$$P_{\text{mig}}(z) = \frac{A^{\text{cap}}}{4\pi R_p^2} = \frac{1}{2} \left(1 - \frac{z}{R_p} \right). \quad (7.25)$$

This predicts that the number of errors for stops near the boundary will approach 50 % of clustered p+t fusion events and linearly decrease with distance from the pad boundary, which can be seen in Fig. 7.4, albeit on a log scale.

In the Road tracker, upstream-going protons must combine with the muon on the next-to-last pad to deposit enough energy to exceed the Road threshold, $E_{\text{Road}} = 1100$ keV. This can be translated into an effective proton range, modified by the energy required to meet the threshold and the energy-range relation for protons and muons, $E_p(r_p)$ and $E_\mu(r_\mu)$. Define the muon and proton energies on the last pad as E_0^p and E_0^μ and second-to-last pad as E_1^p and E_1^μ . The purely geometric model uses the simple condition that $E_1^p > 0$, but incorporating the Road threshold modifies this to

$$E_1^p + E_1^\mu > E_{\text{Road}}. \quad (7.26)$$

For a muon traveling parallel to the z -axis and stopping at a penetration depth z into the pad,

$$E_1^\mu = E_\mu(L + z) - E_\mu(z), \quad (7.27)$$

where $L = 16$ mm is the length of an anode pad. Thus, upstream migration occurs when $E_1^p(z) > E_{\text{Road}} - E_1^\mu(z)$. Inverting the energy-range relation for protons allows this energy

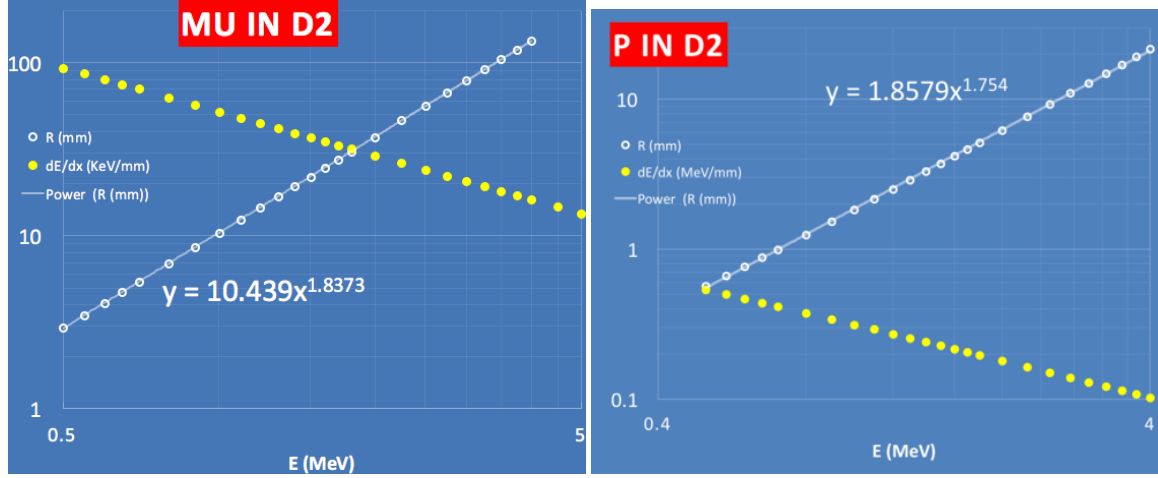


Figure 7.16: SRIM-based simulation of muons (left) and protons (right) stopping in deuterium gas. The energy-range relations (open white circles) are approximated with a power law. Protons from muon-catalyzed fusion have an initial energy of 3.02 MeV.

to be transformed into a required proton penetration into the upstream pad,

$$R_p^{\text{eff}}(z) = R_p - r_p(E_{th} - E_1^\mu(z)). \quad (7.28)$$

The muon and proton energy-range relations are determined via simulation with the SRIM package² and a power-law approximation (Fig. 7.16),

$$E_\mu(r_\mu) = (0.279 \pm 0.006 \text{ MeV}) \left(\frac{r_\mu}{\text{mm}} \right)^{0.54} \quad (7.29)$$

$$E_p(r_p) = (0.70 \pm 0.02 \text{ MeV}) \left(\frac{r_p}{\text{mm}} \right)^{0.57} \quad (7.30)$$

$$r_\mu(E_\mu) = (10.4 \pm 0.5 \text{ mm}) \left(\frac{E_\mu}{\text{MeV}} \right)^{1.84} \quad (7.31)$$

$$r_p(E_p) = (1.86 \pm 0.08 \text{ mm}) \left(\frac{E_p}{\text{MeV}} \right)^{1.75}. \quad (7.32)$$

The uncertainty in the particle ranges comes from longitudinal range straggling, which is consistently within 3.1 – 4.3 % for fusion protons and 3.8 – 4.6 % for stopping muons over their whole range in deuterium. The uncertainty in energy $\delta E/E = (1/p)\delta r/r$, is simply

²www.srim.org

derived from the approximate power-law relationship $r = AE^p$. Substitute these relations and the proton range $R_p = 12.9$ mm into Eq. 7.28 to get an effective proton range, which replaces R_p in Eq. 7.25 to get the probability of upstream migration as a function of z ,

$$P_{mig}(z) = \frac{1}{2} \left(1 - \frac{z}{R_p^{\text{eff}}} \right). \quad (7.33)$$

Consideration of the Road threshold including the energy-range relations of the proton and muon leads to only a 3 % difference in the integral of the migration probability function in the region away from the pad boundary, where threshold effects are more important. (See the difference between the orange line and the shaded region in Fig. 7.18.) The uncertainties assigned to the energy-range relation parameters are therefore negligible compared to the final uncertainty assigned to $P_{mig}(z)$ of ~ 4.5 % in the next section and would in fact be a small effect even with nearly 100 % uncertainty.

Monte Carlo Estimate

The analytical estimation shows that the main features of the p+t fusion migration are captured by simple geometric arguments. The full GEANT4-based MuSun simulation includes additional effects like beam divergence and the pulse-finding threshold. Fig. 7.17 shows the probability of reconstructing the stop position one pad upstream for a muon stopping at a given depth z into the pad. The pulse-finding threshold is evident in the first plot, where muons deposit insufficient energy in the pad to exceed the pulse-finding threshold. Events with ^3He are selected in the middle plot, where this threshold effect is diminished, effectively amounting to a small *downstream* fusion migration. Events with p+t fusion are selected for the third plot and it is obvious that the protons dominate the overall tracker errors due to their long range. Tracker errors in the first 0.5-1.0 mm are not substantially different from the non-fusion case because the proton and triton are emitted back-to-back, so there is no configuration with the muon near the boundary that does not have a fusion particle traveling upstream to exceed the road threshold in that pad.

The fusion migration from each channel is the difference of the mis-reconstructions for fusion events and non-fusion events, shown in Fig. 7.18. This plot includes visualizations of the naive geometric model and the extension based on energy-range relations for muons and




Figure 7.17: GEANT simulation of the probability of reconstructed stop position being one pad upstream from the actual muon stop, as a function of depth of muon stop along z . (Left) Non-fusion muons with a small enough energy deposition do not exceed the pulse-finding threshold. (Middle) ^3He from fusion adds energy, reducing this mis-reconstruction. However, some particles travel upstream, exceeding the Road threshold on the previous pad. (Right) Protons have much longer range, reaching the upstream pad from as far away as 13 mm. Near the boundary, one of the back-to-back triton and proton is very likely to reach the upstream pad, moving the reconstructed stop location upstream. Pad-to-pad variation is likely due to z -dependent variation in the average angle between muon tracks and the z -axis.

protons.

The uncertainty in the migration probability $P_{mig}(z)$ is determined by difference between the MC probability and the model calculation, using the simulation z -slice with the worst agreement. The total probability of migration given a uniform stop distribution across the pad is 0.174 in the simulation and 0.182 in the analytical model, representing a 4.5 % difference,

$$\frac{\delta M_j}{M_j} = 4.5 \%, \quad (7.34)$$

where M_j is the number of events migrating into sub-volume j in the TPC.




Figure 7.18: GEANT simulation of the probability of upstream fusion migration as a function of penetration depth into the muon stop pad (data points), generated from differences of the p+t and ^3He migration probability in Fig. 7.17 with the migration probability for muons. (Left) ^3He fusion-induced migration. Negative numbers indicate net downstream fusion migration. (Right) Migration caused by p+t fusions. The filled overlay is the analytical model with thresholds (Eq. 7.33) and the orange line is the purely geometric migration probability (Eq. 7.25).

7.5.3 Interpolated Muon Stop Density

In the previous section, I derived analytical and MC simulation estimates of the probability of fusion migration as a function of muon penetration depth within a pad, $P_{mig}(z)$. Define the density of muon stops, $n(z)$ and the density of stops for fusion events alone, $f(z) \equiv \epsilon_f n(z)$. The uni-directional upstream migration from the j th z -slice is derived from these quantities,

$$U_j = \int_0^L f(z) P_{mig}(z) dz. \quad (7.35)$$

The tracker does not determine the stop z -coordinate with better resolution than the anode pad, however, so the distribution $f(z)$ (or $n(z)$) must be inferred from an interpolation of the distribution of stops in each z -slice.

Assume that the muon stop density distribution is well-approximated by a gaussian. Minimize the least-squares difference between the integral of a gaussian distribution and the




Figure 7.19: (Left) Interpolation using a gaussian fit to the stop distribution, binned by anode pad. The minimum track length is 2 pads, so the first populated bin is 2 (starting at zero) The final pad 7 is contaminated with tracks that stop past the end of the TPC so it is not included in the fit. (Right) Residuals from gaussian fit, indicating that this interpolation over- or under-estimates total populations by 0.6 %.

counts at each z-slice using fixed weights. The relative residual difference between the z-slice counts and the gaussian interpolation is $\lesssim 1\%$ (see Fig. 7.19).

Uncertainty in U_j arises from error in this interpolation and the uncertainty in the probability of migration (Eq. 7.34). Deriving the interpolation error is difficult, because we have no other measure of the muon stop density within the pad, so I estimate the error as the maximum of the z-slice residuals. The residual counts, δN_j , are given by the difference between observed stops and the interpolated distribution,

$$\delta N_j = \tilde{N}_j - \int_j n(z) dz. \quad (7.36)$$

I take the relative size of the largest of these as an uncertainty in U_j ,

$$\frac{\delta U_i}{U_i} = \max_j \frac{\delta N_j}{N_j}. \quad (7.37)$$

Examining the residuals, we get $\max_j \delta N_j / N_j = 0.6\%$, but to be conservative, this error is inflated to 1.0 %. The errors in individual z-slices are very likely not independent of each other, so the error in the net migration $M_j \equiv U_{j+1} - U_j$ is taken as the linear sum,

$$\delta M_j = (1.0\%) (U_{j+1} + U_j). \quad (7.38)$$

Table 7.4 summarizes this and the uncertainty from the migration probability.

The uncertainties are categorized by the dependence on the upstream migrations into three groups. Statistical error is negligible in all circumstances discussed here. The uncertainties from the migration probability and the fusion-tagging are correlated when calculating different U_j , so they partially cancel in the net migration M_j . Conversely, the $n(z)$ interpolation uncertainty depends on the total number of events migrating because we assume the worst-case correlation between errors in the z-slice, so there is no cancelation when the net migration is small.

Uncertainty description	δU_i	δM_i
$n(z)$ interpolation	$1.0 \% \times U_i$	$\delta U_i + \delta U_{i+1} = 1.0 \% \times (U_{i+1} + U_i)$
Migration probability $P(z)$	$4.5 \% \times U_i$	$ \delta U_i - \delta U_{i+1} = 4.5 \% \times U_{i+1} - U_i $
Statistical	$\sqrt{U_i}$	$\sqrt{U_i + U_{i+1}}$
CluE9 cut background (1%)	$1.0 \% \times U_i$	$ \delta U_i - \delta U_{i+1} = 1.0 \% \times U_{i+1} - U_i $
CluE9 inefficiency (0.1%)	$0.1 \% \times U_i$	$ \delta U_i - \delta U_{i+1} = 0.1 \% \times U_{i+1} - U_i $
Total	$\delta M_i = (5.6 \% \times (U_{i+1} + U_i)) + (1.0 \% \times U_{i+1} - U_i)$	

Table 7.4: Uncertainty contributions to net fusion migration M_i in volume V_i based on the uncertainty in uni-directional upstream migrations U_i . The interpolation error is assumed to be highly-correlated between z-slices, so the sum of errors is taken for the net migration M_i . The statistical error is ignored in the total because it is negligible in all practical cases.

7.5.4 Determining the Fusion Fraction

We now have the necessary tools to determine the fusion fraction and the net fusion migration in each of the z-slices and the fiducial volume. First, though, a technical detail about the R2011 data analysis must be addressed. The full R2011 dataset is divided into many smaller datasets to ease analysis of well-defined subsets of data. The latest full analysis pass of the full dataset on the XSEDE computing cluster did not include the tools necessary to tag p+t fusions with the CluE9 cut described in Sect. 7.5.1. Two of the smaller datasets, called

prod_rn4_mu_g0 and **prod_rn4_mu_h0**, or g0 and h0 for short, were re-processed to include the tools for the CluE9 tag. The practical consequence is that the fusion fraction ϵ_f can only be determined using the g0 and h0 datasets, but this value can then be applied to the full R2011 dataset. In this section, I will use the smaller datasets, but in Sect. 7.6, the full R2011 dataset is used.

The migration probability $P_{mig}(z)$ and muon stop density $n_j(z)$ in z-slice j derived in Sections 7.5.1 and 7.5.2 are combined with the fusion fraction ϵ_f to obtain the uni-directional upstream migrations

$$U_j = \int_j \epsilon_f n_j(z) P_{mig}(z) dz. \quad (7.39)$$

For the purposes of this section, define this method as the function W that takes the input of the number of fusions in all z-slices, and produces the upstream migration from each z-slice,

$$\vec{U} \equiv W(\vec{F}) = W(\epsilon_f \vec{N}) = \epsilon_f W(\vec{N}). \quad (7.40)$$

Here I use the vector notation to indicate that the model requires the input of all F_i and produces a vector with migration values U_i as its components. The modeling is not necessarily linear because of the muon stop density interpolation procedure, but it does permit scalar multiplication, allowing $\epsilon_f \vec{N}$ to substitute for \vec{F} .

The observable quantities are not N and F , but instead \tilde{N} and \tilde{F} , which are polluted by fusion migration. We will write down first approximations to each of these quantities, derive the migrations from Eq. 7.40, and use these to improve the approximation iteratively. Start with an approximation for the fusion fraction, called ϵ_{FV} to distinguish it from the true fusion fraction,

$$\begin{aligned} \epsilon_{FV}^{(0)} &= \frac{\tilde{F}_{FV}}{\tilde{N}_{FV}} \\ &= \frac{F_{FV} + U_7 - U_2}{N_{FV} + U_7 - U_2}. \end{aligned} \quad (7.41)$$

This approximation using the FV is better than the fusion fraction in any particular z-slice \tilde{F}_j/\tilde{N}_j , since the U_j are smaller near the edges and because the denominator N_{FV} is much bigger than any of the N_j . The first approximation to the number of stops is simply

$$N_j^{(0)} = \tilde{N}_j. \quad (7.42)$$

A first approximation to the F_j is made using the fusion fraction

$$F_j^0 = \epsilon_{FV}^{(0)} \tilde{N}_j. \quad (7.43)$$

Write the iteration from step $k - 1$ to step k as

$$\vec{U}^{(k)} = W(\vec{F}^{(k-1)}) = \epsilon_{FV}^{(k-1)} W(\vec{N}^{(k-1)}) \quad (7.44)$$

$$F_i^{(k)} = \tilde{F}_i - M_i^{(k)} = \tilde{F}_i + U_i^{(k)} - U_{i+1}^{(k)} \quad (7.45)$$

$$N_i^{(k)} = \tilde{N}_i + U_i^{(k)} - U_{i+1}^{(k)} \quad (7.46)$$

$$\epsilon_{FV}^{(k)} = \frac{F_{FV}^{(k)}}{N_{FV}^{(k)}} = \frac{\tilde{F}_{FV} + U_2^{(k)} - U_7^{(k)}}{\tilde{N}_{FV} + U_2^{(k)} - U_7^{(k)}}. \quad (7.47)$$

The iteration converges because the corrections for each step are proportional to the fusion fraction, ϵ_f , which is small.

The results of applying this iterative procedure to the datasets g0 and h0 is summarized in Table 7.5. The uncertainty in ϵ_{FV} is the combination of errors in Table 7.4,

$$\begin{aligned} \delta(\epsilon_{FV}) &= \frac{\delta M_{FV}}{N_{FV}} \\ &= \frac{1}{N_{FV}} [1.0 \% \times (U_7 + U_2) \oplus 5.6 \% \times |U_7 - U_2|]. \end{aligned} \quad (7.48)$$

The errors are summed linearly.

The fusion fraction is depends only on muon atomic and molecular kinetics and the pulse clustering algorithm, so it is calculated for each z-slice independently with the expectation of consistency. The results for data set **prod_rn4_mu_h0** are shown in Fig. 7.20. Within this uncertainty, the fusion fraction in individual z-slices, ϵ_i , agrees with that of the fiducial volume. Additionally, consistent fusion fractions are determined using each of the datasets g0 and h0.

$$\epsilon_{FV} = (2.6909 \pm 0.0013)\% \quad (\text{dataset h0}) \quad (7.49)$$

$$\epsilon_{FV} = (2.6891 \pm 0.0018)\% \quad (\text{dataset g0}) \quad (7.50)$$

7.5.5 Migration in R2011 Data Set

The fusion fraction is fixed by the analysis of datasets g0 and h0 in Section 7.5.4 to the average $\epsilon_{FV} = 2.690 \%$ and the iterative migration model (Eq. 7.44) is applied to the full

Dataset g0	Fid. Vol.	$z = 2$	3	4	5	6	7
Muon stops, N_i	$8.23 \cdot 10^8$	0.13923	0.19970	0.23596	0.23529	0.18982	-
p+t fusions, $F_i \times 10^3$	26.89	3.76	5.37	6.34	6.34	5.08	-
Upstr. mig., $U_i \times 10^3$	n/a	0.628	0.943	1.171	1.201	1.019	0.714
Net mig., $M_i \times 10^3$	0.086	0.315	0.227	0.030	-0.183	-0.305	-
Total mig. $\delta M_i \times 10^6$	13.426	15.719	21.144	23.722	22.201	17.330	-
Net mig. $\delta M_i \times 10^6$	4.803	17.642	12.738	1.701	10.221	17.057	-
Fusion frac., F_i/N_i	2.6891%	2.704%	2.689%	2.686%	2.693%	2.677%	-
Mig. err., $\delta M_i/N_i$	0.0018%	0.024%	0.017%	0.011%	0.014%	0.018%	-

Dataset h0	Fid. Vol.	$z = 2$	3	4	5	6	7
Muon stops, N_i	$7.96 \cdot 10^8$	0.14163	0.20532	0.24071	0.23271	0.17963	-
p+t fusions, $F_i \times 10^3$	26.91	3.84	5.53	6.46	6.26	4.82	-
Upstr. mig., $U_i \times 10^3$	n/a	0.637	0.970	1.199	1.200	0.975	0.642
Net mig., $M_i \times 10^3$	0.005	0.333	0.228	0.002	-0.225	-0.333	-
Total mig. $\delta M_i \times 10^6$	12.794	16.076	21.690	23.990	21.755	16.173	-
Net mig. $\delta M_i \times 10^6$	0.280	18.660	12.775	0.107	12.623	18.639	-
Fusion frac., F_i/N_i	2.6909%	2.710%	2.692%	2.684%	2.689%	2.686%	-
Mig. err., $\delta M_i/N_i$	0.0013%	0.025%	0.017%	0.010%	0.015%	0.019%	-

Table 7.5: Migration model data for datasets **prod_rn4_muMinus_g0** and **prod_rn4_muMinus_h0**. All numbers except the total fiducial volume stops and fusion fraction and error are normalized to the total number of fiducial volume stops. Note the multiplicative factor of 10^3 for fusions and migrations and 10^6 for migration errors. There are error contributions from the total number of events migrating and from the net number of events migrating (see Table 7.4). The measured number of muon stops in the $z = 7$ slice is over-estimated due to punch-through muons, so these are an extrapolation of the gaussian stop density approximation.




Figure 7.20: Fraction of events with a clustered p+t fusion for the fiducial volume and each z-slice individually for dataset g0. The fusion fraction is determined with the migration model by iteratively correcting for the number of fusion migrations in each z-slice. It is independent of position in the TPC, so this is a check on the consistency of the model. The net migration (and its uncertainty) is roughly the same in the FV as an individual z-slice, but the total number of events is much larger, so the fusion fraction can be determined much more precisely.

R2011 dataset. The results of the migration model for the full R2011 data set are shown in Table 7.6.

Surprisingly, the net migration into the fiducial volume is much smaller than any of the migration between z-slices. This is not a property of the migration model, but rather a consequence of a fortuitous muon stop distribution that happened to nearly balance the in-migration with out-migration when using the Road tracking algorithm. The portion of the systematic error attributed to uncertainty in the migration probability function $P_{mig}(z)$ is nearly zero because it is proportional to the net migration. The error due to interpolation of the stop distribution $n(z)$ depends on the absolute number of events migrating rather than the net, so it dominates the total error (see Table 7.4),

$$\frac{M_{FV}^{R2011}}{N_{FV}^{R2011}} = (2.3 \pm 1.4) \times 10^{-5}. \quad (7.51)$$

R2011 Dataset	Fid. Vol.	$z = 2$	3	4	5	6	7
Muon stops, N_i	$1.00 \cdot 10^{10}$	0.14223	0.20404	0.23920	0.23272	0.18182	-
Upstr. mig., $U_i \times 10^3$	n/a	0.623	0.941	1.160	1.168	0.961	0.646
Net mig., $M_i \times 10^3$	0.023	0.318	0.219	0.008	-0.207	-0.315	-
Total mig. $\delta M_i \times 10^6$	12.693	15.640	21.009	23.286	21.299	16.075	-
Net mig. $\delta M_i \times 10^6$	1.280	17.783	12.281	0.471	11.598	17.657	-
Fusion frac. (fixed)	2.6900%						
Mig. err., $\delta M_i/N_i$	0.0014%	0.024%	0.016%	0.010%	0.014%	0.019%	-

Table 7.6: Migration model data for the full R2011 dataset. See Table 7.5 caption. The fusion fraction is fixed to the average of ϵ_{FV} from the fusion-tagged analysis using datasets g0 and h0.

7.6 Disappearance Rate Correction

The concepts and results of the previous sections in this chapter are combined to determine a disappearance rate correction for fusion migration in the z direction due to p+t fusions. The rate correction is proportional to the net migration into the fiducial volume with a proportionality constant that depends on the time dependence of the migration probability and the start time of the lifetime fit,

$$\Delta\lambda = -\frac{M_{FV}}{N_{FV}} \times k[\tilde{\eta}, t_{\text{start}}]. \quad (7.52)$$

The fiducial volume migration can be found in Table 7.6, but the proportionality constant must be determined by relating the net migration into each of the individual z -slices to the fitted rate for stops in that slice (see Fig. 7.21). A given z -slice will have inward and outward migrations that are larger and less balanced than the fiducial volume as a whole and additionally the smaller volume reduces the denominator N_i , so the ratio M_i/N_i can vary considerably among the z -slices, allowing precise determination of k .

With the net migrations M_i determined by the model described in Sect. 7.5.2, a fit to the decay time distribution is performed for each z -slice, and the parameter k extracted using




Figure 7.21: Disappearance rate fit for individual z-slices. The standard x- and y-coordinate fiducial volume cuts are applied. (Left) The shift due to fusion interference is much more pronounced in a given z-slice than in the fiducial volume as a whole because the ratio of migrating events to non-migrating events is larger. (Right) The residuals for the second z-slice, showing the effect of a net in-migration of fusion events.

relation 7.52. Though the disappearance rate λ_- is unknown, a shift in its value would only shift the intercept of the linear relation and the slope k would be unchanged. In Fig. 7.22, a fit of the linear relationship determines the slope as

$$k = (-85.8 \pm 7.7) \times 10^3 \text{ s}^{-1}. \quad (7.53)$$

The uncertainty in k is driven by both the statistical uncertainty in the fitted disappearance rates and the systematic uncertainty in the fractional migration from Table 7.4.

With the net migration from Eq. 7.51, and the proportionality constant k from Eq. 7.53, the central value for the correction is

$$\begin{aligned} \Delta\lambda &= -k \frac{M_{FV}}{N_{FV}} \\ &= (85.8 \times 10^3 \text{ s}^{-1}) \times (0.023 \times 10^{-3}) \\ &= 2.0 \text{ s}^{-1}. \end{aligned} \quad (7.54)$$

The relative error in the lifetime shift is determined by the uncertainties in the net migration




Figure 7.22: Linear fit to the disappearance rate vs. fraction of fusion events in the fiducial volume z-slices. The error in migration is given by Table 7.4, whereas the error in Λ_d is determined by the lifetime fit.

and k (Eq. 7.19),

$$\begin{aligned} \frac{\delta(\Delta\lambda)}{\Delta\lambda} &= \frac{\delta k}{k} \oplus \frac{\delta M_{FV}}{M_{FV}} \oplus \frac{\delta N_{FV}}{N_{FV}} \\ &\approx \frac{\delta k}{k} \oplus \frac{\delta M_{FV}}{M_{FV}}. \end{aligned} \quad (7.55)$$

Here \oplus indicates a quadratic sum of uncorrelated errors, or a conservative linear sum with systematic errors. I've neglected δN because the statistical error is negligible compared to the migration model errors in the other terms. Multiplying through by $\Delta\lambda$,

$$\delta(\Delta\lambda) = \delta k \frac{M_{FV}}{N_{FV}} \oplus k \frac{\delta M_{FV}}{N_{FV}}. \quad (7.56)$$

The uncertainty in the correction (7.54) is

$$\begin{aligned}
 \delta(\Delta\lambda) &= (7.7 \times 10^3 \text{ s}^{-1}) \times (2.3 \times 10^{-5}) \oplus (85.8 \times 10^3 \text{ s}^{-1}) \times (1.4 \times 10^{-5}) \\
 &= (0.18 \oplus 1.2) \text{ s}^{-1} \\
 &= 1.4 \text{ s}^{-1}.
 \end{aligned}
 \tag{7.57}$$

7.7 Results

I report the calculation of the fusion interference systematic error in the R2011 MuSun dataset. The dominant contribution is from long-range protons from the p+t branch of muon-catalyzed fusion causing the z-coordinate of the muon stop to be reconstructed one pad upstream of the true stop location. The correction to the measured disappearance rate is

$$\Delta\lambda = (2.0 \pm 1.4) \text{ s}^{-1}. \tag{7.58}$$

The unexpected small central value is a consequence of the fortuitous balancing of in-migration and out-migration of fusion events. A portion of the uncertainty is proportional to the net migration and will be larger for a less balanced distribution. This balancing depends in general on the muon stop distribution and the specifics of the tracking algorithm, so a small correction should not be expected for the R2014 or R2015 datasets. Based on typical population differences between z-slices, one could expect a correction that is 10 times as large and an uncertainty twice as large for an arbitrary realistic stop distribution.

Another consequence of the stop distribution balancing is that the absolute correction may be comparable for the $^3\text{He-n}$ fusion channel or p+t migration in the x- and y-directions, depending on whether there is a similar balancing effect. These additional sources of fusion migration were considered secondary to the z-direction fusion migration and require further analysis. However, the small correction also means that the fit for the lifetime of the full dataset should have only a very small perturbation due to fusion interference. In Fig. 7.23, the early-time residuals are compared for the first and last z-slices in the fiducial volume, which have a net gain and loss of fusion events, respectively, and for the fit to all events in the fiducial volume.




Figure 7.23: Residual comparison for fits to decays from the fiducial volume (top) and z-slices 2 (bottom left) and 6 (bottom right). Z-slice 2 has a net in-migration of fusion events, and z-slice 6 has a net out-migration. The fiducial volume has a very small net in-migration, so the fit residuals show no trend at early times.

The analysis in this chapter is dependent on MC simulation for the model of fusion migration and the efficiency of the CluE9 fusion tag. Additional model-dependent assumptions enter in the gaussian functional form of the the muon stop density distribution. Estimating the error for these models is difficult to do rigorously, and I have assumed conservatively

large uncertainties to reflect this.

7.8 Outlook

I present here an outlook for how this analysis could be improved and extended to all sources of fusion migration.

7.8.1 Direct Measurement of $\eta(t_f)$

The differential fusion acceptance $\eta(t_f)$ is not a directly-observable quantity in this analysis, but it is an important conceptual quantity for understanding the properties of fusion migration. The primary difficulty in measuring $\eta(t_f)$ is that the time of p+t fusion is very difficult to extract from the waveforms. The proton travels up to 13 mm in the D₂ gas, corresponding to a drift time delay of up to 2.6 μ s, which is very long compared to the pulse clustering times and the required resolution on $\eta(t_f)$. One could imagine using waveform templates to extract the fusion time from waveforms of vertical p+t tracks, but tracks parallel to the beam direction are likely hopeless (Fig. 5.7, left and center displays, respectively).

7.8.2 Fusion Tag

This analysis used the data subsets g0 and h0 to determine the fusion fraction because the CluE9 fusion tag was available. In a full analysis of R2014 or R2015 or re-analysis of R2011 data, this tag would be available for the full dataset. Large improvements in the tag efficiency are not likely, since the energies involved are so well separated from other TPC signals, but improved tracking may allow better separation of p+t waveforms from muon track signals, reducing the number of events capable of causing migration.

7.8.3 Fusion Fraction

With fusion tagging for the full dataset, a determination of the true fusion fraction allows direct extraction of the net migration through Eq. 7.24, though care must be taken to distinguish clustered fusions from all events with fusion. This analysis determined the fusion fraction by modeling and correcting for the migrations at the fiducial volume boundaries, but

another strategy is to select a subset of events with a reduced population of events near the boundaries. With ideal stop distribution shaping, the width of the muon stop distribution would be small compared to the fiducial volume dimensions, and no events would stop near enough to the boundaries to allow migration. In this ideal circumstance, the fusion fraction would simply be the fraction of fusion-tagged events with no correction for migration. Since the fusion fraction depends only on the thermodynamic gas properties, it is generalizable to the full dataset without stop distribution shaping.

An initial analysis by R. Ryan uses cuts on the TPC pulses in the very first TPC z-slice to shape the stop distribution. As long as the cuts are performed on z-slices that are 2 pad lengths upstream of selected muon stops, the cuts cannot be influenced by the fusion protons. Requiring that the initial pulses in the muon track be centered in the x- and y-coordinates strongly influences the extent of the stop distribution in those coordinates. Initial results show that the population near the FV boundaries, and therefore the migration near these boundaries, can be reduced by nearly an order of magnitude. The width of the stop distribution parallel to the beam axis can be shaped to a lesser extent by limiting range of accepted energies on the first z-slice. The population reduction is not as dramatic with this technique, but the correction for migrating events can be reduced.

7.8.4 Delayed Fusion Tag for Muon Stop

Muon-catalyzed fusion reactions do not consume the muon, so there is a subset of events with a muon stop and overlapping p+t fusion followed by a ${}^3\text{He-n}$ fusion with enough delay that the ${}^3\text{He}$ waveform is well-separated. This allows the muon stop location to be known independent of the tracking algorithm because the ${}^3\text{He}$ almost always leaves a signal on only the muon stop pad. This technique would be useful for determining the muon stop distribution irrespective of tracker errors and for distinguishing the event topologies for upstream- and downstream-going protons. Since the TPC conflates the y-coordinate with the arrival time of signals, there is an inherent ambiguity between delayed signals and errors in the y-coordinate that does not appear for the x- and z-coordinates. Detection of the neutron from the ${}^3\text{He-n}$ reaction can resolve this because the time of the fusion is

the difference between the neutron time and the muon entrance, $t_f = t_n - t_{\mu\text{SC}}$ and the y-coordinate can be reconstructed from the arrival time of the ${}^3\text{He}$ pulse,

$$t_{3\text{He}} = \frac{y_{\mu}}{v_{\text{drift}}} + t_f. \quad (7.59)$$

Delayed fusion reactions are substantially rarer than prompt fusions, so backgrounds from other rare events become more important.

A delayed tag can also be used to select events to determine the p+t fusion fraction via CluE9-tagged events selected by the delayed ${}^3\text{He}$ position instead of the muon tracking algorithm.

Chapter 8

CONCLUSION

In the final chapter of this thesis, I will describe experimental upgrades inspired in part by R2011 data analysis and summarize my contributions to the MuSun experiment.

8.1 Upgrade Program Informed by R2011 Analysis

Several changes to the MuSun detector systems were made between R2011 and the subsequent runs in 2014 and 2015. Among these were upgraded TPC materials and custom-built cryogenic pre-amplifiers installed inside the insulation vacuum. These upgrades were substantially informed by the analysis of R2011 data.

8.1.1 Impurity Monitor

The possibility of a large chemical impurity drove the effort to improve the TPC resolution to enable an in-situ measurement of the chemical impurity concentration to supplement the existing gas chromatography purity monitor. Delayed muon captures on nitrogen nuclei produce a TPC signal with spectrum peaked in energy between the ^3He fusion energy and low-energy noise and decay electron signals (Fig. 6.1). The R2011 pre-amplifiers mounted outside the vacuum and connected with a long cable run were replaced by custom-designed cryogenic pre-amplifiers mounted just outside the TPC cryo-vessel inside the insulation

Data set	μ^-	μ^+
R2011	5×10^9	0.5×10^9
R2014	6×10^9	1.0×10^9
R2015	7×10^9	1.0×10^9

Table 8.1: Muon decay events collected in each production dataset.

vacuum. The amplifiers were cooled with a LN circulation loop and the input capacitance was reduced by removing the long cable run. The baseline electronic resolution was improved from 30 keV to 10 keV, based on an injected pulser, and the background counts in the capture spectrum region of interest were significantly reduced. In addition, the operating temperature was reduced from 34 K to 31 K to take advantage of the strong temperature dependence of the vapor pressures of O₂ and N₂ (Fig. 3.7).

The primary objective of the resolution improvement was the impurity monitor, but several advantages are gained as a side-effect. The digitization and pulse-finding thresholds are significantly lower (from 150 keV to 80 keV) which reduces the threshold effect, where a particle penetrating barely into a pad but depositing so little energy that no pulse is detected. Fusion migration from ³He will be reduced since it is primarily caused by the threshold effect (see Sect. 7.5.2). Additionally, both the Road and PDir trackers use specific energy cuts or thresholds to categorize events. Improved resolution effectively makes these cuts sharper, allowing better distinction between, e.g. upstream-going and downstream-going protons in the PDir tracker. Some of the variance in measured particle energies is due to chamber effects and other physical processes, limiting this advantage.

A potential disadvantage of the reduced threshold is that the low-energy decay electrons are more likely to leave a signal in the TPC. This opens up the possibility of *electron interference* with the muon tracking, where the prompt decays are more likely to be accepted. This was a leading systematic error in the MuCap experiment [13], but there is no evidence for it in MuSun.

8.1.2 Wall Stop Monitor

The other major upgrade was to construct the TPC out of high-Z materials such as Ag, W, and Au, which all have extremely high rates of muon capture. The large capture rate means that muons stopping in these materials are captured with lifetime ~ 80 ns. The lifetime measurement using a fit starting at 160 ns is mostly undistorted by these events. This is not obviously the optimal strategy. For instance, in the MuCap experiment, μd atoms would form over time from trace deuterium impurities in the protium gas and rapidly

diffuse to the TPC walls, where the muon would capture. Minimizing this lifetime distortion entailed using low- Z materials to suppress the capture branch. In MuSun, however, the fusion systematic necessitates muon tracking algorithms with limited resolution, avoiding the information available on the muon stop pad in order to remain unaffected by subsequent fusion reactions. The degraded resolution increases the likelihood that a muon promptly stopping in the wall is mistakenly reconstructed as a good stop in deuterium.

For R2014 and R2015, several detector components that could potentially be reached by stopping muons were upgraded. The anode plane was constructed from an alumina substrate plated with a 30 μm layer of silver. The Frisch grid was strung with tungsten-rhenium wires on a composite tungsten and stainless steel frame. The cathode plate was covered by a 100 μm silver foil. The field-shaping wires were built from tungsten and sections of these that the muon beam passes through were cut and strung with thin silver wires. The major TPC components that could not be converted to high- Z materials are the MACOR support posts and the high voltage-dividing resistors. The resistors were relocated to be behind the MACOR post to minimize the solid angle of medium- Z material visible to the muon beam.

Additionally, large-volume neutron detectors with a much larger dynamic range were installed to improve the efficiency for detecting neutrons from muon capture on high- Z materials.

8.1.3 Stop Distribution Balancing

There were no feasible hardware improvements that would directly reduce the dominant fusion interference systematic error caused by migration in the z -direction. However, the balancing of migration discussed in Sect. 7.5.5 motivated adjusting the stop distribution (via the beam momentum) to achieve balanced migration on the upstream and downstream boundaries of the fiducial volume.

8.2 Conclusion

MuSun will be the first experiment to measure the doublet muon capture rate on the deuteron precisely enough to meaningfully constrain the weak axial low-energy constant in the two-nucleon system. The experiment requires both large statistics and exceptional attention to

systematic errors to measure the muon lifetime to 10 ppm. While the total systematic error is unlikely to be constrained at this level for the R2011 dataset, several upgrades to the experiment were informed by this analysis, and the results from datasets R2014 and R2015 should be capable of determining all systematic errors to the required precision.

In summer 2011, I participated in the beam run to collect the R2011 dataset. In addition to general operation of the detector system and $\pi E3$ beamline, I was responsible for the electron detector wire chambers (ePCs) and the maintenance and operation of the beam kicker system. I developed several software tools to aid flexible analysis strategies. These include the definition of generic pulse-finding and track-finding interfaces, which allowed the collaboration to iteratively develop algorithms that excel at various facets of the analysis. The Road tracker and threshold pulse-finder that I created to address fusion interference are specific examples, but at least five pulse finding algorithms and six tracking algorithms have been used in MuSun analyses. In addition to these tools, I defined interfaces to create histograms of muon track properties and decay time distributions with a standard set of event selection cuts that facilitate comparison of results using the various algorithms.

The focus of my work was constraining fusion interference, a leading systematic error caused by a discrepancy in muon tracking efficiency between fusion and non-fusion events. Fusion interference fundamentally arises from the difficulty of developing a muon tracking algorithm that is insensitive to the presence of fusion at the 0.1 % level. While most experiments tracking particles in a TPC aim for the highest-precision position determination, the challenge in MuSun is instead to maintain the same precision and bias for fusion and non-fusion events. This is an algorithmic challenge, and cannot be directly improved by experimental upgrades, so the fusion interference correction remains relevant for the final analysis of the R2014 and R2015 datasets.

The Road tracking algorithm is key to controlling this systematic, because it nearly eliminates the fusion/non-fusion discrepancy in the x- and y-coordinates, leaving only a correction in the z-coordinate. Additionally, the interference with the Road tracker generated by fusions does not depend strongly on the time of the fusion or the shape of the underlying stop distribution, which allows for modeling the probability of fusion migration. A formalism and model were used to reconstruct the stop distribution for non-fusion events, tag $p+t$

fusion events with high efficiency, and quantify the net fusion migration into each of the TPC sub-volumes. The proportionality between the disappearance rate variation and the net migration in these sub-volumes allowed me to determine the lifetime correction for the entire TPC fiducial volume. Finally, I reported this correction with uncertainty to be applied to the final lifetime measurement.

After significant experimental upgrades, motivated in part by my analysis of R2011, the MuSun collaboration has collected statistics of 1.3×10^{10} muon decays in two major production runs in 2014 and 2015. This data is currently being analyzed to determine the remaining systematic errors and produce a final lifetime measurement.

GLOSSARY

Bragg curve: distribution of energy loss for an ionizing particle as a function of distance through matter.

Bragg peak: peak at the end of the Bragg curve where the energy loss per unit distance of an ionizing particle is largest.

χ PT: Chiral perturbation theory. A low-energy effective field theory constructed with a Lagrangian that includes all terms consistent with the approximate chiral symmetry of QCD.

CluE9: sum of pulse energies on the muon stop pad and the 8 surrounding pads, counting only pulses in muon track cluster. Also, the tag used to identify p+t fusion events based on this energy.

clustered fusion: fusion reaction that occurs early enough after the muon stop that the resulting pulses are in the same cluster as the muon track. Contrast with **delayed fusion**.

delayed fusion: fusion reaction that occurs late enough in time after the muon stop that the resulting pulses are in a separate cluster from the muon track. Contrast with **clustered fusion**.

disappearance rate: rate of muons disappearing from the target. Most often, this is the sum of the decay rate and the nuclear capture rate. Inverse of the lifetime (e.g. $\lambda_{\mu^-} = 1/\tau_{\mu^-}$). Because of this simple relationship, the rate and lifetime are often interchanged in discussion.

downstream: parallel to the muon beam, in the direction the muons are traveling before they come to rest. This is the positive z -direction in experimental coordinates. Contrast with **upstream**.

E_n : the sum of energies of pulses in a muon track on a z -slice that is n z -slices upstream of the muon stop. E_0 is the sum of pulse energies on the stop z -slice.

E9: sum of pulse energies on the muon stop pad and the 8 surrounding pads, counting all pulses in the muon event, even if they are in different pulse clusters than the muon track.

$\eta(t_f)$: average difference in acceptance between events with fusion and events without fusion, divided by the non-fusion event acceptance, as a function of fusion time, t_f .

fiducial volume (FV): the central volume of the TPC used to select good muon stops. The FV boundary is one anode pad from the edge of the TPC active volume in x and z and 15 mm in y .

fusion: see **muon-catalyzed fusion**.

fusion event: event containing a muon stop in deuterium followed by at least one muon-catalyzed fusion reaction.

fusion interference: a shift in the measured disappearance rate due to the net effect of fusions on muon tracking, primarily through fusion migration.

fusion migration: net effect of tracker errors on a population of muon stops, causing a difference in the number of events in some category (such as fiducial volume stops).

fusion product: a non-muon particle produced in a muon-catalyzed fusion reaction, typically proton, triton, neutron or a ^3He nucleus.

fusion time: time difference between the muon stopping in the deuterium gas and the occurrence of a fusion reaction.

FV: see **fiducial volume**.

lifetime: symbol τ . Average time after which a fraction $1/e$ of the initial population of unstable particles remains. Inverse of the disappearance rate (e.g. $\tau_{\mu^-} = 1/\lambda_{\mu^-}$). Because of this simple relationship, the rate and lifetime are often interchanged in discussion.

lifetime shift: difference in the measured lifetime and true lifetime due to a systematic error in the measurement process.

muon-catalyzed fusion: nuclear fusion of two deuterons from a $d\mu d$ molecule via any of the reactions 3.5a, 3.5b, or 3.5c.

n- ^3He fusion: muon-catalyzed fusion through the reaction $d\mu d \rightarrow n + ^3\text{He} + \mu$ or $d\mu d \rightarrow n + \mu + ^3\text{He}$ (Eqs. 3.5b and 3.5c).

non-fusion event: event containing a muon stop in deuterium without a subsequent muon-catalyzed fusion reaction.

OPE: One-pion exchange. Peace on Earth. Purity of Essence. First-order approximation to the nucleon-nucleon force in which a single pion is exchanged.

p+t fusion: muon-catalyzed fusion through the reaction $d\mu d \rightarrow p + t + \mu$ (Eq. 3.5a).

prod_rn4_mu_g0 and prod_rn4_mu_h0: subsets of the R2011 dataset containing approximately 8 % of the total number of events each. These datasets were analyzed separately from R2011 to include the CluE9 fusion tag. Referred to as “g0” and “h0” for short.

R2011: dataset collected during the Summer 2011 beam run at PSI. This term is sometimes used to refer to this time period as well as the dataset itself.

R2014 and R2015: production datasets in which the total number of events necessary for the final precision were collected.

rate correction: calculated rate that is added or subtracted from the measured muon disappearance rate to correct for a systematic error that causes a lifetime shift.

recycle: following muon-catalyzed fusion, the muon forms another atom with a deuterium nucleus, restarting the kinetic evolution.

TDC: Time-to-digital converter. Digitization electronics that records the leading or trailing edge of a digital logic signal.

TPE: Two pion exchange. Class of graphs in χ PT in which two pions are exchanged between nucleons.

tracker: algorithm that determines the muon stop location based on signals in the TPC.

tracker error: change in muon stop position reconstruction caused by fusion products.

upstream: anti-parallel to the muon beam, opposite the direction the muons are traveling before they come to rest. This is the negative z-direction in experimental coordinates. Contrast with **downstream**.

WFD: Waveform digitizer. Data-acquisition electronics element that records the input voltage over time, typically in fixed-length traces called waveforms.

wall materials: shorthand for any materials in the beam windows, pressure vessels, or those used to construct the TPC detector such as the MACOR support posts, stainless steel Frisch grid, steel field cage wires, aluminum cathode plate, or gold-plated copper anode pad plane.

wall stops: events where the muon stops in any of the wall materials. Usually refers to the subset of these events that are erroneously accepted by the muon tracking algorithm, leading to a systematic error.

z-slice: the volume bounded by the x and y fiducial volume cuts, and a single pad in the z-direction. The fiducial volume is divided into 5 z-slices, for $z \in [2, 6]$ ($z = 0$ is the upstream most z-slice in the TPC).

BIBLIOGRAPHY

- [1] 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.

- [2] C. Patrignani et al. Review of Particle Physics. *Chin. Phys.*, C40(10):100001, 2016.

- [3] Bruno Rossi and David B. Hall. Variation of the Rate of Decay of Mesotrons with Momentum. *Phys. Rev.*, 59:223–228, Feb 1941.

- [4] G. W. Bennett, B. Bousquet, H. N. Brown, G. Bunce, R. M. Carey, P. Cushman, G. T. Danby, P. T. Debevec, M. Deile, H. Deng, S. K. Dhawan, V. P. Druzhinin, L. Duong, F. J. M. Farley, G. V. Fedotovitch, F. E. Gray, D. Grigoriev, M. Grosse-Perdekamp, A. Grossmann, M. F. Hare, D. W. Hertzog, X. Huang, V. W. Hughes, M. Iwasaki, K. Jungmann, D. Kallow, B. I. Khazin, F. Krienen, I. Kronkvist, A. Lam, R. Larsen, Y. Y. Lee, I. Logashenko, R. McNabb, W. Meng, J. P. Miller, W. M. Morse, D. Nikas, C. J. G. Onderwater, Y. Orlov, C. S. Özben, J. M. Paley, Q. Peng, C. C. Polly, J. Pretz, R. Prigl, G. zu Putlitz, T. Qian, S. I. Redin, O. Rind, B. L. Roberts, N. Ryskulov, Y. K. Semertzidis, P. Shagin, Yu. M. Shatunov, E. P. Sichtermann, E. Solodov, M. Sossong, L. R. Sulak, A. Trofimov, P. von Walter, and A. Yamamoto. Measurement of the Negative Muon Anomalous Magnetic Moment to 0.7 ppm. *Phys. Rev. Lett.*, 92:161802, Apr 2004.

- [5] A. M. Baldini, Y. Bao, E. Baracchini, C. Bemporad, F. Berg, M. Biasotti, G. Boca, M. Cascella, P. W. Cattaneo, G. Cavoto, F. Cei, C. Cerri, G. Chiarello, C. Chiri, A. Corvaglia, A. de Bari, M. De Gerone, T. Doke, A. D’Onofrio, S. Dussoni, J. Egger, Y. Fujii, L. Galli, F. Gatti, F. Grancagnolo, M. Grassi, A. Graziosi, D. N. Grigoriev, T. Haruyama, M. Hildebrandt, Z. Hodge, K. Ieki, F. Ignatov, T. Iwamoto, D. Kaneko, T. I. Kang, P.-R. Kettle, B. I. Khazin, N. Khomutov, A. Korenchenko, N. Kravchuk, G. M. A. Lim, A. Maki, S. Mihara, W. Molzon, Toshinori Mori, F. Morsani, A. Mtchedilishvili, D. Mzavia, S. Nakaura, R. Nardò, D. Nicolò, H. Nishiguchi, M. Nishimura, S. Ogawa, W. Ootani, S. Orito, M. Panareo, A. Papa, R. Pazzi, A. Pepino, G. Piredda, G. Pizzigoni, A. Popov, F. Raffaelli, F. Renga, E. Ripiccini, S. Ritt, M. Rossella, G. Rutar, R. Sawada, F. Sergiampietri, G. Signorelli, M. Simonetta, G. F. Tassielli, F. TENCHINI, Y. Uchiyama, M. Venturini, C. Voena, A. Yamamoto, K. Yoshida, Z. You, Yu. V.

- Yudin, and D. Zanello. Search for the lepton flavour violating decay ... with the full dataset of the MEG experiment. *The European Physical Journal C*, 76(8):434, 2016.
- [6] Wilhelm H. Bertl et al. A Search for muon to electron conversion in muonic gold. *Eur. Phys. J.*, C47:337–346, 2006.
- [7] Laura E Marcucci, A Kievsky, S Rosati, R Schiavilla, and M Viviani. Chiral Effective Field Theory Predictions for Muon Capture on Deuteron and He 3. *Physical review letters*, 108(5):052502, 2012.
- [8] David B Kaplan, Martin J Savage, and Mark B Wise. Two-nucleon systems from effective field theory. *Nuclear Physics B*, 534(1-2):329–355, 1998.
- [9] J Menéndez, D Gazit, and A Schwenk. Chiral two-body currents in nuclei: Gamow-Teller transitions and neutrinoless double-beta decay. *Physical review letters*, 107(6):062501, 2011.
- [10] J Engel, F Šimkovic, and P Vogel. Chiral two-body currents and neutrinoless double- β decay in the quasiparticle random-phase approximation. *Physical Review C*, 89(6):064308, 2014.
- [11] P Ackerbauer, D.V Balin, V.M Baturin, G.A Beer, W.H Breunlich, T Case, K Crowe, H Daniel, J Deutsch, J Govaerts, Yu.S Grigoriev, F.J Hartmann, P Kammel, R King, B Lauss, E.M Maev, V.E Markushin, J Marton, M Mühlbauer, C Petitjean, Th Petitjean, G.E Petrov, R Prieels, W Prymas, W Schott, G.G Semenchuk, Yu.V Smirenin, A.A Vorobyov, N.I. Voropaev, and P Wojciechowski. A precision measurement of nuclear muon capture on ^3He . *Physics Letters B*, 417(3-4):224 – 232, 1998.
- [12] DV Balin, VA Ganzha, SM Kozlov, EM Maev, GE Petrov, MA Soroka, GN Schapkin, GG Semenchuk, VA Trofimov, AA Vasiliev, et al. High precision study of muon catalyzed fusion in D2 and HD gas. *Physics of Particles and Nuclei*, 42(2):185–214, 2011.
- [13] 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 1% Precision and Determination of the Pseudoscalar Coupling g_P . *Phys. Rev. Lett.*, 110:012504, Jan 2013.
- [14] F. Halzen and A.D. Martin. *Quarks and leptons: an introductory course in modern particle physics*. Wiley, 1984.

- [15] Steven Weinberg. Charge Symmetry of Weak Interactions. *Phys. Rev.*, 112:1375–1379, Nov 1958.
- [16] E.D. Commins and P.H. Bucksbaum. *Weak Interactions of Leptons and Quarks*. Cambridge University Press, 1983.
- [17] A Mooser, S Ulmer, Klaus Blaum, Kurt Franke, H Kracke, C Leiteritz, W Quint, Cricia C Rodegheri, C Smorra, and J Walz. Direct high-precision measurement of the magnetic moment of the proton. *Nature*, 509(7502):596–599, 2014.
- [18] Peter J. Mohr, David B. Newell, and Barry N. Taylor. CODATA recommended values of the fundamental physical constants: 2014. *Rev. Mod. Phys.*, 88:035009, Sep 2016.
- [19] M. P. Mendenhall, R. W. Pattie, Y. Bagdasarova, D. B. Berguno, L. J. Broussard, R. Carr, S. Currie, X. Ding, B. W. Filippone, A. García, P. Geltenbort, K. P. Hicker-son, J. Hoagland, A. T. Holley, R. Hong, T. M. Ito, A. Knecht, C.-Y. Liu, J. L. Liu, M. Makela, R. R. Mammei, J. W. Martin, D. Melconian, S. D. Moore, C. L. Morris, A. Pérez Galván, R. Picker, M. L. Pitt, B. Plaster, J. C. Ramsey, R. Rios, A. Saunders, S. J. Seestrom, E. I. Sharapov, W. E. Sondheim, E. Tatar, R. B. Vogelaar, B. VornDick, C. Wrede, A. R. Young, and B. A. Zeck. Precision measurement of the neutron β -decay asymmetry. *Phys. Rev. C*, 87:032501, Mar 2013.
- [20] D. Mund, B. Märkisch, M. Deissenroth, J. Krempel, M. Schumann, H. Abele, A. Petoukhov, and T. Soldner. Determination of the Weak Axial Vector Coupling $\lambda=g_A/g_V$ from a Measurement of the β -Asymmetry Parameter A in Neutron Beta Decay. *Phys. Rev. Lett.*, 110:172502, Apr 2013.
- [21] G. Ecker. Effective field theories, 2005.
- [22] Howard Georgi. Effective field theory. *Annual review of nuclear and particle science*, 43(1):209–252, 1993.
- [23] Peter Kammel and Kuniharu Kubodera. Precision muon capture. *Annual Review of Nuclear and Particle Science*, 60:327–353, 2010.
- [24] Steven Weinberg. Effective chiral Lagrangians for nucleon-pion interactions and nuclear forces. *Nuclear Physics B*, 363(1):3–18, 1991.
- [25] Steven Weinberg. Nuclear forces from chiral lagrangians. *Physics Letters B*, 251(2):288 – 292, 1990.
- [26] T-S Park, LE Marcucci, R Schiavilla, M Viviani, A Kievsky, S Rosati, K Kubodera, D-P Min, and M Rho. Parameter-free effective field theory calculation for the solar proton-fusion and hep processes. *Physical Review C*, 67(5):055206, 2003.

- [27] Yu.A. Akulov and B.A. Mamyrin. Half-life and value for the bare triton. *Physics Letters B*, 610(1–2):45 – 49, 2005.
- [28] David B Kaplan, Martin J Savage, and Mark B Wise. A New expansion for nucleon-nucleon interactions. *Physics Letters B*, 424(3):390–396, 1998.
- [29] Jiunn-Wei Chen, Takashi Inoue, Xiangdong Ji, and Yingchuan Li. Fixing two-nucleon weak-axial coupling L_1, A from μ -d capture. *Physical Review C*, 72(6):061001, 2005.
- [30] Hilla De-Leon, Lucas Platter, and Doron Gazit. Tritium β -decay in pionless effective field theory. *arXiv preprint arXiv:1611.10004*, 2016.
- [31] Jiunn-Wei Chen, Karsten M Heeger, and RG Hamish Robertson. Constraining the leading weak axial two-body current by recent solar neutrino flux data. *Physical Review C*, 67(2):025801, 2003.
- [32] KIT Brown, MN Butler, and DB Guenther. Constraints on proton-proton fusion from helioseismology. *arXiv preprint nucl-th/0207008*, 2002.
- [33] 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. Axial-Current Matrix Elements in Light Nuclei from Lattice QCD. *arXiv preprint arXiv:1611.00344*, 2016.
- [34] R. Machleidt and D.R. Entem. Chiral effective field theory and nuclear forces. *Physics Reports*, 503(1):1 – 75, 2011.
- [35] Martin Savage. Lecture notes, PHYS 560, Oct 1999.
- [36] Paul Buettiker and Ulf-G Meißner. Pion–nucleon scattering inside the Mandelstam triangle. *Nuclear Physics A*, 668(1):97–112, 2000.
- [37] Nadia Fettes, Ulf-G. Meißner, and Sven Steininger. Pion-nucleon scattering in chiral perturbation theory (I): Isospin-symmetric case. *Nuclear Physics A*, 640(2):199 – 234, 1998.
- [38] Doron Gazit, Sofia Quaglioni, and Petr Navratil. Three-nucleon low-energy constants from the consistency of interactions and currents in chiral effective field theory. *Physical review letters*, 103(10):102502, 2009.
- [39] Jr Adam, M Tater, E Truhlik, E Epelbaum, R Machleidt, and P Ricci. Calculation of doublet capture rate for muon capture in deuterium within chiral effective field theory. *Physics Letters B*, 709(1):93–100, 2012.

- [40] S Ando, T-S Park, K Kubodera, and F Myhrer. The μ -d capture rate in effective field theory. *Physics Letters B*, 533(1):25–36, 2002.
- [41] P Ricci, E Truhlik, B Mosconi, and J Smejkal. Muon capture in deuterium. *Nuclear Physics A*, 837(1-2):110–144, 2010.
- [42] 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.
- [43] M. Cargnelli et al. Measurement of the muon capture rate in gaseous deuterium. In *Proceedings of the XXIII Yamada Conf. on Nuclear Weak Processes and Nuclear Structure, Osaka, Japan*, 1989.
- [44] G. Bardin et al. A Measurement of the Rate of Muon Capture in Liquid Deuterium by the Lifetime Technique. *Nucl. Phys.*, 1986.
- [45] 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. *Physical Review*, 105:1127–1128, February 1957.
- [46] V. A. Andreev, R. M. Carey, V. A. Ganzha, A. Gardestig, T. Gorringer, F. E. Gray, D. W. Hertzog, M. Hildebrandt, P. Kammel, B. Kiburg, S. Knaack, P. A. Kravtsov, A. G. Krivshich, K. Kubodera, B. Lauss, K. R. Lynch, E. M. Maev, O. E. Maev, F. Mulhauser, F. Myhrer, 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. Muon Capture on the Deuteron – The MuSun Experiment, 2010.
- [47] Daniel Salvat. Muon Kinetics in MuSun. internal report, Dec 2015.
- [48] Arnaud Dupays. Isotopic Effects in the Muon Transfer from p μ and d μ to Heavier Atoms. *Physical review letters*, 93(4):043401, 2004.
- [49] I-T. Wang, E. W. Anderson, E. J. Bleser, L. M. Lederman, S. L. Meyer, J. L. Rosen, and J. E. Rothberg. Muon Capture in $(p\mu d)^+$ Molecules. *Phys. Rev.*, 139:B1528–B1538, Sep 1965.
- [50] A. Bertin, A. Vitale, A. Placci, and E. Zavattini. Muon Capture in Gaseous Deuterium. *Phys. Rev. D*, 8:3774–3793, Dec 1973.
- [51] O Bunemann, TE Cranshaw, and JA Harvey. Design of grid ionization chambers. *Canadian journal of research*, 27(5):191–206, 1949.

- [52] 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.
- [53] Nandita Raha. Measurement of the $d\mu d$ Quartet-to-Doublet Molecular Formation Rate Ratio ($\lambda q: \lambda d$) and the μd Hyperfine Rate (λqd) Using the Fusion Neutrons from μ -Stops in D2 Gas. 2015.
- [54] I Alekseev, Ev Arkhipov, S Bondarenko, O Fedorchenko, V Ganzha, K Ivshin, P Kammel, P Kravtsov, C Petitjean, V Trofimov, et al. Cryogenic distillation facility for isotopic purification of protium and deuterium. *Review of Scientific Instruments*, 86(12):125102, 2015.
- [55] VA Ganzha, PA Kravtsov, OE Maev, GN Schapkin, GG Semenchuk, VA Trofimov, AA Vasilyev, ME Vznuzdaev, SM Clayton, P Kammel, et al. 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.
- [56] S Ritt, P Amaudruz, and K Olchanski. MIDAS (Maximum Integration Data Acquisition System). URL <https://midas.psi.ch>, 2001.
- [57] Ian Foster. Globus Online: Accelerating and democratizing science through cloud-based services. *IEEE Internet Computing*, 15(3):70–73, 2011.
- [58] John Towns, Timothy Cockerill, Maytal Dahan, Ian Foster, Kelly Gaither, Andrew Grimshaw, Victor Hazlewood, Scott Lathrop, Dave Lifka, Gregory D Peterson, et al. XSEDE: accelerating scientific discovery. *Computing in Science & Engineering*, 16(5):62–74, 2014.
- [59] Frederik Wauters. MuSun Purity Analysis Report. internal report, Aug 2015.
- [60] Peter Kammel. MuSun Systematics Run. internal report, Jan 2016.
- [61] Peter Kammel. Kinetics fit. internal report, Jul 2016.
- [62] Daniel Salvat. MuSun No-Fusion Analysis. internal report, May 2016.