

Estimating the Aging Rate for the Cosmic Ray Veto Detector of the Mu2e Experiment at Fermilab

by

Mackenzie N. Devilbiss

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Physics)
in the University of Michigan
2023

Doctoral Committee:

Professor Myron Campbell, Chair
Professor Aaron Pierce
Associate Professor Joshua Spitz
Professor Sarah Veatch
Associate Professor Liuyan Zhao

Mackenzie N. Devilbiss

macndevv@umich.edu

ORCID iD: 0000-0002-3774-0470

© Mackenzie N. Devilbiss 2023

Acknowledgments

First and foremost, I would like to acknowledge my family for their support, guidance, and love during my time in graduate school and beyond. I would also like to thank my partner, JP, for his love, patience, and for sharing his experience to make my journey smoother. I would also like to acknowledge two groups of friends that I made over the past few years for encouraging me and continually reviving my spirits: the fellow students and staff at Fermilab that I connected with at the User's Center and the Great Lakes Flow community.

I also want to extend extreme gratitude to Dr. Myron Campbell, my advisor, for his steadfast encouragement and support since before I committed to graduate studies at University of Michigan. With Myron's guidance, I was able to begin my graduate journey with confidence and continue to grow continuously as a physicist and professional. I would like to also give special acknowledgement to Dr. Monica Tecchio, our research scientist, for being an invaluable resource for advice in navigating my place in a large collaboration.

I am very thankful for being a part of the Mu2e collaboration and for being able to work on-site at Fermilab for part of my graduate career. I want to personally thank Dr. Robert Bernstein for his countless hours of project planning, analysis assistance, interpreting results, and shaping my on-site work. I would like to show appreciation for Dr. Robert Kutschke, Dr. Ray Culbertson, and Dr. David Brown for their support in my code infrastructure and validation role. I have learned an immense amount from these few people. I also would like to extend appreciation to Dr. Alan Bross, Dr. Anna Pla, Brian Leung, and the rest of the Scintillation R & D group for their guidance and advice on my titanium dioxide cladding studies. I would also like to thank the Mu2e CRV team for their support with the Wideband test bench, namely Dr. Yongyi Wu, Dr. Yuri Oksuzian, and Dr. Ralf Ehrlich. It has been a pleasure to be a fellow collaborator with these great scientists.

Lastly, I would like to thank my committee members for their time and consideration as I conclude this journey. Thank you.

Table of Contents

Acknowledgments	ii
List of Figures	vi
List of Tables	x
List of Appendices	xi
List of Abbreviations	xii
Abstract	xiv
Chapter 1 Charged Lepton Flavor Violation	1
1.1 The Standard Model	2
1.2 Neutrino Oscillations within the Standard Model	5
1.3 Lepton Flavor Violation Beyond the Standard Model	7
1.3.1 Model Independent Analyses	8
1.3.2 Specific CLFV Models	9
1.4 CLFV Experimental Review	12
1.4.1 $\mu \rightarrow e\gamma$ Experiments	13
1.4.2 $\mu \rightarrow eee$ Experiments	16
1.4.3 $\mu N \rightarrow eN$ Experiments	17
Chapter 2 The Mu2e Experiment	22
2.1 Conceptual Design	22
2.2 Experimental Backgrounds	24
2.2.1 Intrinsic Backgrounds	25
2.2.2 Prompt Beam Processes	25
2.2.3 Delayed Beam Processes	26
2.2.4 Cosmic Ray Induced Backgrounds	27
2.2.5 Reconstruction Errors	28
2.3 Background Evaluation/Sensitivity Estimate	28
2.3.1 Run I Sensitivity Estimate (SU2020 paper)	29
2.3.2 Full Sensitivity Estimate	31
2.4 Outlook Beyond Mu2e	32
Chapter 3 Mu2e Detector Overview	34

3.1	Mu2e Beamline at Fermilab	35
3.2	Mu2e Solenoids	37
3.2.1	Production Solenoid	38
3.2.2	Transport Solenoid	40
3.2.3	Detector Solenoid	41
3.3	Straw Tube Tracker	44
3.4	Crystal Calorimeter	46
Chapter 4	Cosmic Ray Veto	51
4.1	CRV Counters and Dicounters	52
4.1.1	Scintillator Bars	52
4.1.2	Titanium Dioxide Cladding	54
4.1.3	Wavelength Shifting Fibers	54
4.1.4	Dicounters	55
4.2	CRV Modules	57
4.2.1	Sectors	59
4.3	Readout Electronics and Data Acquisition	61
4.3.1	Counter Motherboards	63
4.3.2	Front End Boards	65
4.3.3	Readout Controllers	68
4.3.4	Data Acquisition	70
Chapter 5	Mu2e Software Framework	72
5.1	Mu2e Framework Architecture	73
5.1.1	Event Building	73
5.1.2	Key Base Software Packages	74
5.2	Mu2e Online Software	76
5.2.1	Pattern Recognition	76
5.2.2	Track Reconstruction	79
5.3	Mu2e Trigger	79
5.4	Mu2e Offline Software	81
5.5	Studies Using Offline	82
5.5.1	Estimating the Antiproton Background	82
5.5.2	Improving Cosmic Trigger Efficiency	98
5.5.3	Optimizing the Cosmic Ray Veto (CRV) Dead-Time Window	108
Chapter 6	Estimating the CRV Aging Rate	118
6.1	Scintillator Aging	118
6.2	Wideband Test Bench Setup	119

6.2.1	Collecting Data with the Test Bench	123
6.2.2	Wideband Data Processing Code	124
6.3	Wideband Aging Measurement Strategy	125
6.3.1	Temperature and Bias Voltage Stability	129
6.4	Results	131
6.4.1	Light Yield Distributions	131
6.4.2	Investigating Data Quality	136
6.4.3	Final Aging Analysis	140
6.5	Conclusions	142
Chapter 7	Measuring Aging in Titanium Dioxide Cladding Material	144
7.1	Lab 6 Aging Measurement Strategy	145
7.1.1	Analysis Codes	146
7.1.2	Calibration and Standardization of the Spectrophotometer	146
7.1.3	Stability and Aging of NOvA and GoreDRP Standards	149
7.2	Description of Samples	154
7.2.1	Manufacturer-Produced Titanium Dioxide Cladding Coupons	155
7.2.2	Extruded Titanium Dioxide Cladding Surface Samples	159
7.3	Results	162
7.3.1	Aging of Titanium Dioxide Cladding Coupons	162
7.3.2	Aging of Titanium Dioxide Cladding Extrusions	165
7.3.3	Studies on Other Counter Components	167
7.4	Comparing with Simulations	168
7.5	Conclusions	174
Chapter 8	Mu2e Code Management	179
8.1	Code Maintenance and File Storage Systems	179
8.2	UPS Products	180
8.3	Internal Products	180
8.4	Integrating Updated Products	181
Chapter 9	Conclusion	184
Appendix	186

List of Figures

Figure 1.1	Elementary particles of the Standard Model	3
Figure 1.2	Feynman diagrams involving loops for CLFV processes in the minimally extended SM	7
Figure 1.3	Model independent BSM sensitivity reach of CLFV experiments	9
Figure 1.4	$\mu^- N \rightarrow e^- N$ BSM SUSY Feynman diagram	10
Figure 1.5	$\mu^- N \rightarrow e^- N$ BSM Feynman diagram with leptoquarks	10
Figure 1.6	$\mu^- N \rightarrow e^- N$ BSM Feynman diagrams with flavor violating Higgs interactions	11
Figure 1.7	$\mu^- \rightarrow e^- \gamma$ BSM Feynman diagram with heavy neutrinos	12
Figure 1.8	Summary of experimentally determined CLFV muon BR limits	14
Figure 1.9	Diagram of the MEG experiment detectors	15
Figure 1.10	$\mu^+ \rightarrow e^+ e^- e^+$ SM Feynman diagram	16
Figure 1.11	Diagram of the Mu3e experiment detectors	17
Figure 2.1	Dependence of $R_{\mu e}$ on Z	24
Figure 3.1	Mu2e detector overview	34
Figure 3.2	Fermilab accelerator overview	35
Figure 3.3	Proton pulse timing overview	36
Figure 3.4	Live window and background timing overview	37
Figure 3.5	Superconducting solenoid cable cross-section	38
Figure 3.6	Production target simulation	39
Figure 3.7	Transport Solenoid diagram	40
Figure 3.8	Muon stopping target	42
Figure 3.9	Straw tube details	45
Figure 3.10	Photograph of a Mu2e tracker panel	45
Figure 3.11	Straw tube tracker overview	46
Figure 3.12	CsI crystal calorimeter overview	47
Figure 3.13	Calorimeter crystal details	48
Figure 4.1	Photograph of a CRV counter cross-section	52
Figure 4.2	Photograph of a CRV scintillator extrusion	53
Figure 4.3	Emission and absorption spectra for Kuraray Y11 wavelength shifting fibers . .	55

Figure 4.4	Light yield vs fiber diameter	56
Figure 4.5	Fiber guide bar and counter motherboard details	57
Figure 4.6	CRV dicounter stacking diagram	58
Figure 4.7	CRV module schematic	58
Figure 4.8	CRV sector diagram	60
Figure 4.9	Block diagram of CRV electronics chain	62
Figure 4.10	Photograph of a Hamamatsu S13360-2050VE SiPM	64
Figure 4.11	Photograph of a prototype FEB	66
Figure 4.12	Block diagram of FEB data flow path	67
Figure 4.13	Photograph of a CRV ROC	68
Figure 4.14	Block diagram of ROC data flow path	69
Figure 5.1	Helix panel hits projected into X-Y and ϕ -Z planes	77
Figure 5.2	CPR selection of hits in X-Y plane	79
Figure 5.3	Trigger efficiency for CE events	81
Figure 5.4	Randomly stopped final state momentum distribution	86
Figure 5.5	Final state momentum distributions in different foils	88
Figure 5.6	Final state momentum distributions at different radii	88
Figure 5.7	ShieldingM stopped antiproton position and final state momentum distribution	91
Figure 5.8	Final state electrons using different physics lists	92
Figure 5.9	Roy thesis extracted annihilations	94
Figure 5.10	Geant4 pion spectrum	95
Figure 5.11	Roy vs Bernstein pion spectra	96
Figure 5.12	Roy vs Bernstein fit comparison	96
Figure 5.13	Initial reconstructed helix momentum distributions comparing CPR and TPR	99
Figure 5.14	Initial momentum resolution distributions comparing CPR and TPR	100
Figure 5.15	Final state reconstructed helix momentum distributions comparing CPR and TPR	103
Figure 5.16	Momentum resolution distributions comparing CPR and TPR	103
Figure 5.17	Reconstructed helix impact parameter distributions comparing CPR and TPR	104
Figure 5.18	λ distributions comparing CPR and TPR	105
Figure 5.19	λ distribution for TPR after ϕ Z modifications	106
Figure 5.20	Reconstructed helix momentum distribution for TPR after ϕ Z modifications	107
Figure 5.21	Momentum resolution distributions comparing TPR before and after changes	108
Figure 5.22	Time difference between track and CRV stub versus Z	111
Figure 5.23	Results of effective time clustering on cry31s91b0	112
Figure 5.24	dT vs. Z(CRV) for events with different numbers of effective time clusters	113
Figure 5.25	Propagation corrected time difference between track and CRV stub versus Z	114

Figure 5.26	Corrected time difference between track and CRV stub versus $Z(\text{CRV})$	115
Figure 6.1	Photographs of the Wideband building at Fermilab	120
Figure 6.2	Wideband test bench electronics diagram	121
Figure 6.3	Photograph of the Wideband test bench trigger paddles	122
Figure 6.4	Raw waveform of a cosmic event in the CRV	123
Figure 6.5	Test beam data for PE versus temperature SiPM corrections	126
Figure 6.6	General characteristics of Hamamatsu S13360-2050VE SiPM	128
Figure 6.7	Effect of temperature correction on a PE distribution	128
Figure 6.8	Example CMB and FEB temperature plots	129
Figure 6.9	Example FEB bias voltage plots	130
Figure 6.10	PE distribution example with fit	132
Figure 6.11	Low occupancy PE distribution example	132
Figure 6.12	Chi squared distributions of fits of PE plots	133
Figure 6.13	Example Wideband module 127 aging plot	134
Figure 6.14	Aging rates for module 127	135
Figure 6.15	Aging rates for FEB0 and FEB1 of module 127	136
Figure 6.16	Example module 127 aging plot	137
Figure 6.17	Average temperatures for an FEB and a CMB of module 127	138
Figure 6.18	Period 2 aging rates by channel for module 127	139
Figure 6.19	Period 3 aging rates by channel for module 127	140
Figure 6.20	Period 1 aging rates by channel for module 127	141
Figure 6.21	Period 1 aging rates for module 127	141
Figure 7.1	Photographs of the NOvA N-27-09-NC standard	148
Figure 7.2	Example NOvA N-27-09-NC reflectance spectrum	148
Figure 7.3	Photographs of the GoreDRP standard	149
Figure 7.4	Example GoreDRP reflectance spectrum	150
Figure 7.5	Standard deviation of repeated standard measurements over time at 390 nm	151
Figure 7.6	NOvA standard reflectance over time at 390 nm	152
Figure 7.7	NOvA standard aging rate summary plot	153
Figure 7.8	GoreDRP standard reflectance over time at 390 nm	154
Figure 7.9	GoreDRP standard aging rate summary plot	155
Figure 7.10	Photographs of an example TiO_2 coupon	156
Figure 7.11	Example TiO_2 KE12291 reflectance spectrum	157
Figure 7.12	Difference in reflectance spectra over time for two samples	158
Figure 7.13	Example spectra from both sides of a USPN115101 coupon sample	160
Figure 7.14	Photographs of an example extrusion sample	161

Figure 7.15	Example spectra from both sides of an extrusion sample	163
Figure 7.16	USPN115101 reflectance over time at 390 nm	164
Figure 7.17	USPN115101 coupon aging rate summary plot	165
Figure 7.18	Extrusion sample reflectance over time at 390 nm	166
Figure 7.19	Extrusion sample aging rate summary plot	167
Figure 7.20	CRV counter diagram with MC position labels	170
Figure 7.21	Simulations of light yield within CRV dicounters using a bismuth source	177
Figure 7.22	Simulations of light yield within CRV dicounters using a proton source	178
Figure B.1	Efficiency vs operating voltage for trigger paddles	190

List of Tables

Table 1.1	Experimentally determined limits on branching ratios for a selection of CLFV processes	13
Table 2.1	Run I background summary	30
Table 2.2	Total estimated background summary	31
Table 4.1	SiPM device specifications	65
Table 5.1	Estimated stopped antiprotons for different foil positions	89
Table 5.2	Estimated stopped antiprotons for different physics lists	93
Table 5.3	TPR improvements from CircleFit changes	102
Table 5.4	TPR improvements from FZFit changes	106
Table 5.5	Overall TPR algorithm improvements in CEMix	107
Table 5.6	Overall TPR algorithm improvements in CEPlusMix	107
Table 7.1	Light yield simulation results from bismuth source	171
Table 7.2	Light yield simulation results from proton source with 0.7 mm fiber holes	171
Table 7.3	Light yield simulation results from proton source with 0.9 mm fiber holes	172
Table 7.4	Light yield simulation results from proton source with 1.2 mm fiber holes	172
Table 7.5	Average values from bismuth source light yield simulations	173
Table 7.6	Average values from proton source light yield simulations	173
Table B.1	Optimal operating voltages for the PMTs used on the scintillating trigger paddles in Wideband.	189

List of Appendices

Appendix A	Propagation Time within CRV Bars	186
Appendix B	Determining Optimal PMT Operating Voltage	187

List of Abbreviations

ADC analog-to-digital converter	IPA Inner Proton Absorber
BR branching ratio	LFV Lepton Flavor Violation
BSM Beyond the Standard Model	LHT Littlest Higgs Model with T-Parity
CE conversion electron	LQ Leptoquarks
CKM Cabibbo-Kobayashi-Maskawa	MBS Muon Beam Stop
CsI cesium iodide	MEG Mu to E Gamma Experiment
CLFV Charged Lepton Flavor Violation	MPPC multi-pixel photon counter
CMB Counter Motherboard	MSSM Minimal Supersymmetric Standard Model
CMS Compact Muon Solenoid	MVA multi-variate analysis
COMET Coherent Muon to Electron Transition Experiment	NICADD Northern Illinois University's Northern Illinois Center for Accelerator and Detector Development
CPR Calorimeter-based pattern recognition	PDE photon detection efficiency
CRV Cosmic Ray Veto	PE photoelectron
CVMFS CERN Virtual Machine File System	PMNS Pontecorvo-Maki-Nakagawa-Sakata
DAC digital-to-analog converter	PMT photomultiplier tube
DAQ Data Acquisition	PS Production Solenoid
DIF decay-in-flight	PSI Paul Scherrer Institute
DIO decay-in-orbit	OPA Outer Proton Absorber
DS Detector Solenoid	RC run control
DTC Data Transfer Controller	RMC radiative muon capture
FCC Feynman Computing Center	RMD radiative muon decay
FEB Front End Board	ROC Readout Controller
FGB fiber guide bar	RPC radiative pion capture
FNAL Fermi National Accelerator Laboratory	SES single event sensitivity
FPGA field programmable gate array	SiPM silicom photomultiplier
FTS file transfer service	SM Standard Model
GUI graphical user interface	SPE single photoelectron
HPGE high purity Germanium	
HV high voltage	

STM Stopping Target Monitor
TPR Tracker-based pattern recognition
TS Transport Solenoid

TSV through-silicon via
UPD Unix Product Distribution
UPS Unix Product Support
2HDM Two Higgs Doublet Model

Abstract

The Mu2e experiment at Fermilab will search for the neutrinoless conversion of a muon to an electron in the presence of an atomic nucleus. This process is an example of Charged Lepton Flavor Violation (CLFV), which is heavily suppressed in the Standard Model (SM). Due to this suppression, any observation of a CLFV process, like Mu2e conversion, is a clear sign of new physics. In order to achieve the expected single event sensitivity of 3×10^{-17} , the Mu2e detector apparatus has been designed to carefully mitigate backgrounds while selecting for the conversion signature of a monoenergetic 105 MeV/c electron emerging from the muon stopping target. Even in the case of no observation, the Mu2e experiment will improve the current upper limit for the muon conversion-to-capture ratio $R_{\mu e}$ by four orders of magnitude.

Cosmic ray induced backgrounds present the largest background mode contribution for Mu2e. The Cosmic Ray Veto (CRV) serves as an active shield against cosmic rays by surrounding the bulk of the Mu2e detector apparatus with four layers of scintillator, arranged in long bars. Cosmic ray particles generate light as they pass through the scintillator. When a signal is detected in three of the four layers, localized in space and time, it signifies that a cosmic ray has passed through the CRV and a veto is applied to the data in offline analysis. This assures that no cosmic ray or cosmic ray-induced events generate a false conversion electron signal. An overall efficiency of 99.99% is required for the CRV to reduce the total number of cosmic ray-induced background events to below one event for the duration of the experiment. If left unmitigated, cosmic ray particles would cause one conversion-like event per day.

The long-term response of the CRV is being studied at Fermilab as the Mu2e experiment prepares for operation. CRV modules are being stored in the Wideband building at Fermilab where a cosmic ray test bench has been collecting data using single modules since 2021. This data is the first cosmic ray data collected using an entire module and front end CRV electronics. Beyond using the Wideband test bench to check CRV modules for functionality, data can also be collected long-term to understand how the detector response changes over time. Plastic-based scintillating detectors can degrade over time and, as a result, generally produce less light in response to the same amount of energy deposition when cosmic particles pass through it. By measuring the light yield over time, an aging rate may be estimated for the CRV and the impacts to detector response as a result of aging can be anticipated and mitigated if necessary. Cosmic ray data collected with the

Wideband test bench was used to estimate an aging rate for the Mu2e CRV, in the first study of its kind, and the estimated aging rate was determined to be $8.06\% \pm 0.33\%$ per year in units of photoelectron yield in the channel lost annually.

The aging rate estimated from the study above prompted an independent investigation into the base components used to produce CRV counters to determine if one component could be responsible for most of the overall aging seen in the Wideband study. The counters have three main components: the plastic-based scintillator bulk, inner wavelength-shifting fibers, and outer reflective cladding. By monitoring characteristic properties of each of the three components individually, we can determine if any one component is changing over time in a way that would affect the CRV response overall. For this dissertation, titanium dioxide samples which are used in the reflective outer cladding for the CRV counters were monitored over time for decrease in surface reflectivity. It was observed that titanium dioxide lost surface reflectivity at an average rate on the order of 0.5% per year. Using this average rate in conjunction with simulations of the Mu2e CRV counters predicts roughly a 4% drop in light yield in a given CRV channel.

In addition to work on estimating the aging of the Mu2e CRV, several smaller studies and roles within the collaboration are discussed. Small software and simulation investigations are discussed that optimize the dead time of the CRV in response to a cosmic ray passing through the detector, improve the efficiency of detecting cosmic rays by using the timing information from both the CRV and the tracker, and estimating the background rate from antiproton annihilation. Maintenance and management of the Mu2e codebase is also discussed in the context of holding a service role for the Mu2e collaboration.

Chapter 1 Charged Lepton Flavor Violation

High energy physics is a vast field that explores the smallest fragments of the universe: elementary particles. For particle physicists to understand the mechanisms behind particle interactions, we must exploit our knowledge about the fundamental symmetries of the universe and question the rules that arise from our current best theory: the Standard Model (SM). The symmetries of the SM predict many exact conservation laws for particles, but are all of these laws reflected and obeyed by nature? What laws draw suspicion and how can we test them? The topic of this thesis is a test of lepton number conservation in the SM, but before we can understand the importance of lepton number conservation and how it can be used to search for new physics, we must first understand the basics of the SM.

The SM is the most successful theory to date in describing the classification of particles and the forces that govern their fundamental interactions. Leptons and quarks, together known as fermions, are the particles that give rise to matter of the universe. The SM organizes them into three generations of fermions, and each with a corresponding antiparticle. The three fundamental forces by which particles interact are the strong, weak, and electromagnetic interactions. Each force is mediated by a boson: gluons (g) mediate strong interactions, W^\pm and Z bosons mediate weak interactions, and photons (γ) mediate electromagnetic interactions. The Higgs boson (H) is a scalar boson that provides mass to all particles via the Higgs mechanism. At the core, an $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge theory defines the symmetries of the SM, where the subscript C denotes color, L describes the chirality of left-handed fermions, and Y denotes hypercharge.

By following its predictions, the SM has guided particle physicists to discovery for decades, which is what makes it such a compelling theory. The existence of the Higgs boson was theorized far before its discovery in 2012 [1]. Similarly, the W^\pm and Z bosons, the gluon, and the charm, top, and bottom quarks were all predicted by the SM before they were discovered. In addition, the measured value of the Z boson mass [2] aligns almost perfectly with the theoretical SM mass. Perhaps the strongest argument in favor of the SM is the measurement of the electron magnetic dipole moment, which agrees with the SM prediction to within 1 part in 100 billion [3]. This leads to our most precise determination of the fine structure constant, α .

Though the SM seems to provide a robust construct for particle interactions on Earth, there are also glaring shortcomings that do not have a place within this model. Gravity is an indisputable

fundamental force of nature, and it is currently irreconcilable with the SM. In similar opposition, neutrino oscillations have been experimentally observed, proving that neutrinos have mass [4–6]. Many other experimental searches also heighten tension with the SM. For example, recent results included the LHCb experiment at CERN reporting evidence for violation of lepton universality in B-meson decays [7] with a 3σ deviation from the SM and g-2 experiment at Fermi National Accelerator Laboratory (FNAL), also known as Fermilab, reporting a measured 4.2σ deviation [8] from the SM on the value of a_μ , the anomalous muon magnetic dipole moment. Until 2022, the measured mass of the W boson agreed very well with the SM, however, new results from the CDF-II experiment at Fermilab indicate that the W boson mass is much higher than the SM prediction, by 7σ [9]. A combination of decades of evidence against the SM leads us to believe that a better model must exist to completely describe nature as we know it.

One channel with strong support for investigating Beyond the Standard Model (BSM) physics lies in the conservation of lepton flavor number, or lack thereof. The conservation of lepton number naturally arises from the gauge invariance and renormalizability of the SM Lagrangian, and as a consequence processes that violate lepton number conservation are forbidden in the SM [10]. The non-conservation of lepton family number is called Lepton Flavor Violation (LFV). When particles with nonzero charge violate lepton family number conservation, the process is specifically called CLFV. In the SM, CLFV processes have an extremely small branching ratio (BR), deeming them theoretically undetectable by modern experimental technologies. However, this makes CLFV processes unique experimental targets; the smallness of the SM contributions to these processes means that any observation is a clear sign of new BSM physics. This chapter will contextualize LFV within the SM and motivate why specific CLFV processes, like $\mu \rightarrow e$ conversion, may be attractive channels in the search for BSM physics. The chapter will conclude with a brief review of experimental searches for muon CLFV decay modes.

1.1 The Standard Model

Within the SM, particles are organized into two main classes: fermions and bosons. Fermionic particles have spin-1/2, while bosonic particles have spin-0 or spin-1. In Figure 1.1 below, the spin, mass, and charge of each particle species is given in the upper left-hand corner of each box.

Fermions are organized into three generations, arranged according to their mass. The first generation is the lightest, and includes the up and down quarks (u , d), electrons (e), and electron neutrinos (ν_e). The second generation is heavier than the first, and includes the charm and strange quarks (c , s), muons (μ), and muon neutrinos (ν_μ). The third generation is the heaviest, and includes the top and bottom quarks (t , b), taus (τ), and tau neutrinos (ν_τ). Each fermion also has a

Standard Model of Elementary Particles




Figure 1.1: Elementary particles of the Standard Model. Fermions are shown on the left and bosons are shown on the right. The colors of the boxes indicate classifications of particles: fermions split into quarks and leptons and bosons split into scalar and vector types.

corresponding antiparticle, which are not shown in the figure above. Together, fermions comprise all known matter in the universe. The first generation of fermions are stable whereas the latter two generations are unstable, and particles comprised of these fermions spontaneously decay to more stable particles.

Quarks have color charge, fractional electric charge, and interact via the strong force through the exchange of gluons as force carriers. Quarks also interact with the electroweak force, and thus interact with all SM bosons. Composite particles that are comprised of multiple quarks are always color-neutral and free quarks are never observed in nature. Leptons have either $-e$ or neutral electric charge and no color charge, so they interact via the weak force through exchange of W and Z bosons. Neutrinos, all three flavors, have neutral charge. Electrons, muons, and taus have $-e$ charge, and may thus also interact via the electromagnetic force through photon exchange. All leptons may interact via the electroweak force. The Higgs boson is the only scalar boson in the SM. Instead of being a force carrier, the Higgs is responsible for giving mass to massive fermions and bosons via the Higgs mechanism as these particles interact with the scalar Higgs field.

In the SM, with the help of quantum field theory, particles are represented by elementary fields: fermionic fields, gauge fields, and the Higgs field. Fermions are represented by $SU(2)_L$ doublets if they are left-handed, and singlets if they are right-handed. The left-handed quark and

lepton fields may be written as

$$q_i = \begin{pmatrix} u_{Li} \\ d_{Li} \end{pmatrix}, l_i = \begin{pmatrix} \nu_{Li} \\ e_{Li} \end{pmatrix} \quad (1.1)$$

where the index i runs through the quark generation numbers. The right-handed quarks and leptons of the j^{th} generation are written as u_{jR} , d_{jR} and e_{jR} . There are no right-handed neutrinos because they do not exist in the SM. There is a single Higgs field represented by a complex $SU(2)_L$ doublet, Φ , where the conjugate is given by $\tilde{\Phi} = i\tau_2\Phi^*$ and τ_2 is the second Pauli matrix.

The full SM Lagrangian for particle interactions can be written in three pieces:

$$\mathcal{L}_{SM} = \mathcal{L}_{gauge} + \mathcal{L}_{Higgs} + \mathcal{L}_{Yukawa} \quad (1.2)$$

The gauge term describes interactions with the gauge bosons through the fields G_μ , A_μ , and B_μ . The fermion flavor mixing sector of the SM is contained within the Yukawa terms of the Lagrangian above, which can be written as [11]:

$$-\mathcal{L}_{Yukawa} = (Y_u)_{ij}\bar{q}_{Li}u_{Rj}\tilde{\Phi} + (Y_d)_{ij}\bar{q}_{Li}d_{Rj}\Phi + (Y_e)_{ij}\bar{l}_{Li}e_{Rj}\Phi + h.c. \quad (1.3)$$

The Higgs potential term above describes how fermions interact with the Higgs field to obtain mass via electroweak $SU(2)_L \times U(1)_Y$ symmetry breaking where the Higgs develops a nonzero vacuum expectation value, where

$$\langle\Phi\rangle^T = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix} \quad (1.4)$$

which numerically evaluates to $v \simeq 246$ GeV. Then, the fermion mass terms may take the form

$$(m_f)_{ij} = \frac{v}{\sqrt{2}}(Y_f)_{ij}, f = u, d, e \quad (1.5)$$

where v is the Higgs field vacuum expectation value and Y_f are complex 3×3 Yukawa matrices that contain coupling constants in flavor space for $f = u, d, e$. Any complex matrix may be diagonalized using a series of unitary rotations, so diagonalized Yukawa matrices may take the form

$$Y_f = V_f\hat{Y}_fW_f^\dagger, f = u, d, e \quad (1.6)$$

where the unitary rotations by V_f and W_f have no impact on the kinetic terms, neutral current interactions, or the coupling with the physical Higgs field h . The fermion couplings to the physical Higgs h are proportional to the mass matrix, so no flavor violation is induced in the Lagrangian

terms involving the Higgs [11]:

$$-\mathcal{L}_{h\bar{f}f} = \frac{m_f}{v} \bar{f}_L f_R h + h.c. \quad (1.7)$$

where m_f are the now diagonalized mass matrices. These mass matrices have no off-diagonal elements, and thus none of these terms allow for flavor violation. Similarly, the neutral current interactions that are mediated by photons or Z bosons are formulated in the same manner, and also result in only flavor conserving Lagrangian terms.

From the rotations in Equation 1.6, the only flavor-violating terms are the charged-current interaction terms with W bosons.

$$\mathcal{L}_{CC} = \frac{g}{\sqrt{2}} (\bar{u}_L \gamma^\mu (V_u^\dagger V_d) d_L + \bar{\nu}_L \gamma^\mu (V_\nu^\dagger V_e) e_L) W_\mu^+ + h.c. \quad (1.8)$$

From here, the first term illustrates that flavor violation in the quark sector arises because $V_u \neq V_d$ in order to make both Y_u and Y_d diagonal [12]. This misalignment gives rise to flavor-violating transitions parameterized by the Cabibbo-Kobayashi-Maskawa (CKM) matrix [13, 14], $V_{CKM} = V_u^\dagger V_d$. On the other hand, the second term describes flavor mixing in the leptonic sector, where one may choose that $V_\nu = V_e$ because no other term involves the lepton doublets. Since neutrinos in the SM are massless, lepton flavor is exactly conserved.

Given the SM condition of massless neutrinos, it is now clear how lepton flavor number is strictly conserved. The Lagrangian \mathcal{L}_{Yukawa} is invariant under three independent global $U(1)$ rotations that correspond with each lepton family, making the lepton family numbers L_e , L_μ , and L_τ independently conserved quantities. However, this review of lepton number conservation in the SM hints at conditions which may allow for LFV. Introducing additional terms into \mathcal{L}_{Yukawa} that involve the lepton fields in some non-trivial basis in flavor space can generate flavor violation in the lepton sector. Some example extensions may be the addition of a neutrino mass term or a second Higgs doublet [11].

1.2 Neutrino Oscillations within the Standard Model

Given that neutrino oscillations have now been experimentally observed, it is necessary to extend the SM to include these interactions that we know to exist. The SM consequence of vanishing neutrino mass was under question for decades before it was discovered that neutrinos can change flavor. The first hypothesis of neutrino oscillation came from Pontecorvo in 1957 [15]. Then, in the 1960's, the 'solar neutrino problem' attracted much attention as some experiments observed a deficit in the flux of solar electron neutrinos compared to what the SM predicted. Only electron neutrinos, no other neutrino flavors, are produced as a byproduct of solar nuclear fusion processes. The

Homestake Experiment observed in the late 1960's that the number of electron neutrinos incident on Earth's surface was much lower than it should be if all solar neutrinos were electron neutrinos. Homestake used a chlorine-based detector [16] which was only sensitive to electron neutrinos, so there was no direct evidence for neutrino oscillation at this time. In the following decades, similar results of an electron neutrino deficit were observed by radiochemical detectors [2, 5]. The next generation of experiments used water-based Cherenkov detectors which were sensitive to all neutrino flavors, gathering more evidence for neutrino oscillations [17, 18]. Finally, neutrino oscillation was confirmed in 1998 by the Super-Kamiokande (Super-K) Experiment [4]. Super-K precisely measured atmospheric neutrinos to determine the ν_μ/ν_e flux ratio over a wide energy range as well as the zenith angle distributions for 'e-like' and ' μ -like' events [4, 5]. The combination of these results, combined with similar findings from other experiments at the time, effectively resolved the atmospheric and solar neutrino problems and concluded that neutrino flavor oscillations were responsible for the observed discrepancies.

The discovery of neutrino oscillations ultimately meant that the SM must be extended to include neutrino mass terms. Two ways to minimally extend the SM and reconcile this are to introduce right-handed neutrino fields or non-renormalizable operators [11]. With the addition of a right-handed neutrino singlet ν_R , a new set of Yukawa Lagrangian terms arise that give rise to Dirac-type neutrino mass terms, much like other fermions:

$$\mathcal{L}_D = -(Y_\nu)_{ij} \bar{\nu}_{Ri} \tilde{\Phi}^\dagger L_{Lj} + h.c. \implies (m_\nu^D)_{ij} = \frac{\nu}{\sqrt{2}} (Y_\nu)_{ij} \quad (1.9)$$

where the Y_ν Yukawa couplings that dictate the magnitude of neutrino masses would have to be very small to explain experimentally measured masses. With the addition of non-renormalizable operators, Majorana mass terms arise for left-handed neutrinos:

$$\mathcal{L}_M = -\frac{1}{2} (m_\nu^M)_{ij} \bar{\nu}_{Li}^C \nu_{Lj} + h.c. \quad (1.10)$$

where $\bar{\nu}_L^C$ are charge-conjugated neutrino fields. Operators that are allowed by the SM to mediate these terms must be dimension $d = 5$ [19], and effective Lagrangian terms can be written as:

$$\mathcal{L}_{M,eff} = \frac{C_{ij}}{\Lambda} (\bar{L}_L^C \tau_2 \Phi) (\Phi^T \tau_2 L_{Lj}) + h.c. \implies (m_\nu^M)_{ij} = \frac{C_{ij} \nu^2}{\Lambda} \quad (1.11)$$

where Λ is the mass scale of the extra degrees of freedom and C_{ij} is a charge conjugation matrix. The resulting Majorana mass terms will thus be small if $\Lambda \gg \nu$.

Despite the exact mechanism that extends the SM to include neutrino masses, the final consequence is a diagonal unitary matrix that relates the weak neutrino eigenstates to the mass

eigenstates, analogous to the CKM matrix of the quark sector. The matrix that describes lepton sector mixing is called the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, $U_{PMNS} \equiv V_e^\dagger V_\nu$. The interactions that are allowed via U_{PMNS} can give rise to CLFV processes. However, these processes must occur through loop diagrams involving neutrinos and W bosons. Feynman diagrams for $\mu \rightarrow e\gamma$ and $\mu N \rightarrow eN$ are shown in Figure 1.2. As expected, the Feynman diagram for $\mu - e$ conversion in the presence of a nucleus looks almost identical to that of $\mu \rightarrow e\gamma$, with the additional photonic exchange with quarks in the nucleus.




Figure 1.2: Feynman diagrams involving loops for CLFV processes in the minimally extended SM. The process $\mu^- \rightarrow e^- \gamma$ is shown on the left and $\mu^- N \rightarrow e^- N$ is shown on the right.

After understanding neutrino oscillations and the loops of the Feynman diagrams in Figure 1.2, the suppression of CLFV processes in the SM becomes clear. For example, the BR for $\mu \rightarrow e\gamma$ processes may be calculated as [20–22]:

$$BR(\mu \rightarrow e\gamma) = \frac{3\alpha}{32\pi} \left| \sum_{i=2,3} U_{\mu i}^* U_{ei} \frac{\Delta m_{i1}^2}{M_W^2} \right|^2 \sim \mathcal{O}(10^{-54}) \quad (1.12)$$

where α is the fine structure constant, $U_{\mu i}$ and U_{ei} are elements in the PMNS matrix, Δm_{i1}^2 is the difference between the involved neutrino masses squared, and M_W is the W boson mass. Numerically, this BR is $\mathcal{O}(10^{-54})$, impossibly small. This stems from the factor $|\Delta m_{i1}^2/M_W^2|^2$, where the meager neutrino mass difference squared is divided by the large W boson mass squared, making this factor infinitesimal and thus forbidding CLFV in the SM. Though it is true that the rate of CLFV in the SM is not exactly zero, the rate is so small that it is *effectively* zero. Any BR rate this small is beyond the reach of conceivable experiments. Though the challenge of searching for such a rare process may seem insurmountable, CLFV modes can provide a unique way to search for BSM physics. Because the rates for these processes are predicted to be so small in the SM, any observation of CLFV is a clear sign of new physics.

1.3 Lepton Flavor Violation Beyond the Standard Model

Though the conservation of lepton number is a direct consequence of the SM symmetry breaking properties, theories that extend the minimal SM predict many possible sources of LFV.

This makes rare muon decay mode experiments attractive candidates for the potential discovery of BSM physics.

1.3.1 Model Independent Analyses

After electroweak symmetry breaking, new terms can be written to mediate CLFV processes. These Lagrangian terms include effective operators of dimension five or greater. The effective Lagrangian terms for $e\mu$ modes are [23]:

$$\mathcal{L}_{CLFV} = \frac{m_\mu}{(1 + \kappa)\Lambda^2} \bar{\mu}_R \sigma_{\mu\nu} e_L F^{\mu\nu} + \frac{\kappa}{(1 + \kappa)\Lambda^2} \bar{\mu}_L \gamma_\mu e_L \left(\sum_{q=u,d} \bar{q}_L \gamma^\mu q_L \right) + h.c. \quad (1.13)$$

where m_μ is the muon mass, $F^{\mu\nu}$ is the photon field strength, and the subscripts L and R indicate the chirality of the fermion fields. The coefficients on the two terms contain Λ , the effective mass scale of new physics and κ , a dimensionless parameter that determines the relative size of the two terms. The first term is a 'dipole' interaction term, a loop-type, magnetic moment operator which would directly mediate $\mu \rightarrow e\gamma$ processes and at order α mediate $\mu \rightarrow eee$ and $\mu N \rightarrow eN$ processes. The second term is a 'contact' interaction term, a four-fermion operator which would directly mediate $\mu N \rightarrow eN$ processes, and mediate $\mu \rightarrow e\gamma$ processes at one-loop level [23]. When κ is large, four-fermion operators are the dominant mode for CLFV, and when the converse is true, dipole operators dominate.

Using the Λ - κ parameter space defined by the Lagrangian above, we may visualize the sensitivity reach of different CLFV experiments [21]. The highest effective mass scales can be probed by $\mu N \rightarrow eN$ experiments, which can almost always probe higher effective mass scales than $\mu \rightarrow e\gamma$ experiments. For values of $\kappa \gg 1$, $\mu N \rightarrow eN$ experiments are most sensitive and can explore effective mass scales $\mathcal{O}(10^4 \text{ TeV})$. If $\kappa \ll 1$, $\mu \rightarrow e\gamma$ experiments instead become the more sensitive search, though they may only probe mass scales below $\mathcal{O}(10^4 \text{ TeV})$. The restricted discovery reach of $\mu \rightarrow e\gamma$ experiments compared to $\mu N \rightarrow eN$ experiments is attributed to the inability to suppress accidental $e\gamma$ coincidences, which makes reducing the sensitivity of these experiments below $\mathcal{O}(10^{-14})$ challenging [24]. One conclusion that may be drawn from Figure 1.3 is that an observation in a $\mu \rightarrow e\gamma$ experiment would be a harbinger of a $\mu \rightarrow e$ conversion observation. Another, non-observation from a $\mu \rightarrow e\gamma$ search does not preclude observation in $\mu N \rightarrow eN$ searches. Brief discussions of CLFV experiments like Mu to E Gamma Experiment (MEG) at PSI and other $\mu \rightarrow e$ conversion searches like the Coherent Muon to Electron Transition Experiment (COMET) at J-PARC will be discussed in the next chapter.

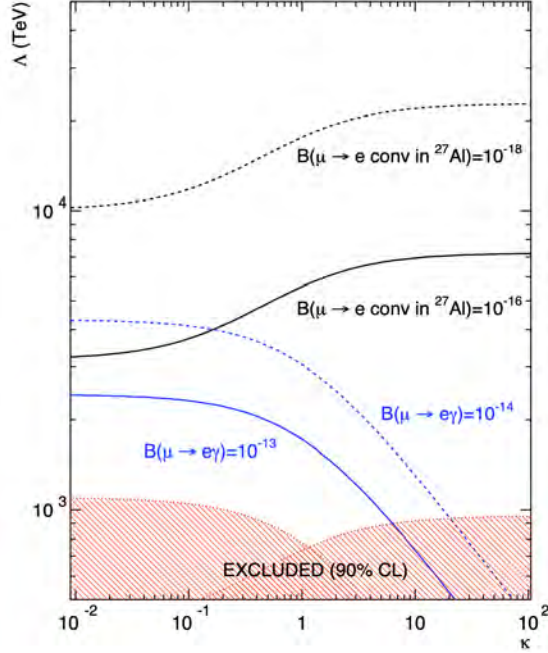


Figure 1.3: Model independent BSM sensitivity reach of CLFV experiments. The sensitivity of $\mu \rightarrow e$ conversion and μ to $e\gamma$ experiments to the mass scale of new physics, Λ , as a function of the parameter κ . The red shaded area of the parameter space has been excluded by past experiments and the solid and dashed lines indicate the reach of proposed experiments with different sensitivity projections[21].

1.3.2 Specific CLFV Models

In addition to building model independent Lagrangians involving BSM physics, a number of specific theories predict enhanced CLFV rates. All new physics models aim to address one of the shortcomings of the SM, and sometimes flavor violation arises as a consequence. A small selection of those theories will be discussed below, with further reviews in [10, 11].

1.3.2.1 $SO(10)$ SUSY Grand Unified Model

Supersymmetric (SUSY) models are popular because they can address a number of SM complications, such as the gauge hierarchy problem. These models may also predict enhanced CLFV rates as a result of the new interactions allowed by SUSY particles. In SUSY theories, each SM particle has a corresponding SUSY partner. The fermionic SUSY particles are called sleptons (\tilde{l}) and squarks (\tilde{q}), while the bosonic SUSY particles are called gluinos (\tilde{g}), winos (\tilde{W}), binos (\tilde{B}) and the Higgsino (\tilde{H}). Charginos ($\tilde{\chi}_i^\pm$) and neutralinos ($\tilde{\chi}_i^0$) arise as mass eigenstates from the mixing of bosonic superpartners. Many different SUSY theories exist depending on the specific mechanism behind SUSY breaking. A SUSY Feynman diagram for $\mu \rightarrow e$ conversion can be seen in Figure 1.4.




Figure 1.4: $\mu^- N \rightarrow e^- N$ BSM SUSY Feynman diagram. This BSM Feynman diagram depicts $\mu \rightarrow e$ conversion mediated by SUSY particles sleptons (\tilde{l}_j) and neutralinos ($\tilde{\chi}_i^0$).

One version of SUSY models that lead to enhanced CLFV rates are SO(10) models, where different Yukawa couplings are considered for neutrinos that can be 'PMNS-like' or 'CKM-like' [25, 26]. These theories involve very massive right-handed neutrinos from the new Yukawa couplings that are able to mediate CLFV interactions, while also accounting for the light Higgs mass, LHC search limits for SUSY particles, and the observed values for neutrino mixing angles. Calculations for BRs of $\mu \rightarrow e\gamma$ processes from these models are anywhere from $\mathcal{O}(10^{-8})$ to $\mathcal{O}(10^{-18})$ depending on the specific model type and value of $\tan \beta$, the ratio of the vacuum expectation value of the two Higgs doublets [26]. Corresponding $\mu N \rightarrow e N$ rates are within reach of Mu2e observation.

1.3.2.2 Scalar Leptoquark Models

Leptoquarks (LQ) are theoretical particles that mediate lepton-quark interactions, hence the name 'lepto-quark'. These particles would be color-triplet bosons that carry both lepton and baryon numbers and could have different interaction strengths with different flavors of leptons. Since LQs directly couple quarks and leptons, CLFV processes can be mediated at tree-level via interactions shown in Figure 1.5. In a particular LQ model which enhances the rate for $\mu \rightarrow e$




Figure 1.5: $\mu^- N \rightarrow e^- N$ BSM Feynman diagram with leptoquarks. This tree-level BSM Feynman diagram depicts $\mu \rightarrow e$ conversion mediated by leptoquarks (LQ).

conversion, a scalar LQ at the TeV scale couples to the top quark. This enhances the top quark mass

while conforming to known constraints from collider and quark-flavor physics experiments [27]. Depending on the specific coupling strength and mass of the TeV-range LQ, experiments like MEG and Mu2e can constrain much of the parameter space for LQ theories.

1.3.2.3 Higgs Induced Flavor Violation

Many BSM theories predict extensions to the SM Higgs sector of that allow for flavor violation. Such theories are attractive because the single SM Higgs boson is only weakly coupled to leptons and the lighter quarks by small Yukawa couplings, so new physics contributions can easily dominate SM predictions [28]. Some of these theories involve more than one Higgs field, the simplest of which is the Two Higgs Doublet Model (2HDM), which includes two Higgs fields. In 2HDMs, both Higgses obtain a nonzero vacuum expectation value and couple to fermions, generating new Lagrangian Yukawa mass terms from the new vacuum expectation values, shown below.

$$\mathcal{L}_Y = -m_i \bar{f}_{Li} f_{Ri} - (Y^a)_{ij} \bar{f}_{Li} f_{Rj} h^a + h.c. + \dots \quad (1.14)$$

where a runs over all scalars. There are also new Yukawa coupling terms where the single Higgs results are now modified by generally replacing the single SM Higgs field with the sum over several Higgses [28]. The nonzero off-diagonal terms in $(Y^a)_{ij}$ give rise to flavor violation in the Yukawa couplings, meaning that CLFV processes like $\mu \rightarrow e$ conversion and $\mu \rightarrow eee$ may proceed at tree-level. Some tree-level and one-loop Feynman diagrams for flavor violating Higgs interactions are shown in Figure 1.6. $\mu \rightarrow e$ conversion experiments are the most sensitive probe into determining




Figure 1.6: $\mu^- N \rightarrow e^- N$ BSM Feynman diagrams with flavor violating Higgs interactions. Several BSM Feynman diagrams depict $\mu^- \rightarrow e^-$ conversion contributions involving flavor violating Higgs interactions [28]. The flavor violating Higgs Yukawa couplings are denoted by $Y_{e\mu}$ and $Y_{\mu e}$.

the coupling $|Y_{e\mu}|$, and Mu2e will further constrain the bounds on this coupling constant.

1.3.2.4 Littlest Higgs Model with T-Parity

In the Littlest Higgs Model with T-Parity (LHT), the Higgs boson emerges as a pseudo-Nambu-Goldstone boson under several symmetries [29]. If all symmetries are broken, referred to here as collective symmetry breaking, then the Higgs can gain an additional mass contribution. T-parity is the symmetry introduced to accomplish this, and functions as a discrete symmetry analogous to R-parity in the Minimal Supersymmetric Standard Model (MSSM). As a consequence of T-parity, these models predict the existence of new particles on the scale of 1 TeV that is consistent with present precision electroweak results [30]. Parameters in these models predict measurable rates for both $\mu \rightarrow e\gamma$ and $\mu \rightarrow e$ conversion such that a Mu2e non-observation would fully exclude the validity of these models.

1.3.2.5 Heavy Neutrinos

Models containing new neutrino mass states are potential candidates for enhanced CLFV rates. If new heavy neutrino mass states arise from Lagrangian terms, processes such as $\mu \rightarrow e\gamma$ may be mediated as shown in the figure below. Some of these models predict BR rates of $\mu \rightarrow e\gamma$ as




Figure 1.7: $\mu^- \rightarrow e^- \gamma$ BSM Feynman diagram with heavy neutrinos. This BSM Feynman diagram depicts $\mu \rightarrow e\gamma$ mediated by heavy neutrinos ν_{iL} [11].

large as $\mathcal{O}(10^{-13})$, which implies that the BR for $\mu \rightarrow e$ conversion would be around $\mathcal{O}(10^{-14})$ [31]. In this case, the Mu2e experiment should observe many conversion electron (CE) events.

1.4 CLFV Experimental Review

Many different experiments have been designed and executed to search for evidence of CLFV since these processes produce clear signatures for new physics. So far, no experiment has been successful in observing CLFV, but many constraints have been put on the BRs of these processes. A selection of these experimentally determined upper limits are presented in Table 1.1. Upon inspection of Table 1.1, rare muon decay experiments provide the most sensitive searches for CLFV, and result in the lowest BR upper limits. Figure 1.8 summarizes the evolution of BR

constraints on rare muon processes over time, and includes an outlook on the reach of the planned next generation of muon CLFV experiments.

The experimental sensitivity for rare muon decay experiments has improved greatly over time, benefiting from more intensely produced muon beams and advanced detector technologies. However, these experiments are subject to a wide spectrum of experimental backgrounds that must be heavily mitigated to have any chance at observing CLFV. A detailed discussion of the different background modes that contribute to $\mu \rightarrow e$ conversion experiments will follow in Chapter 2. For this section, only a brief review of the backgrounds to each type of experiment will be mentioned.

For the rest of this chapter, we will review past experiments probing the muon sector of CLFV, namely experiments searching for $\mu \rightarrow e\gamma$, $\mu \rightarrow eee$, and $\mu N \rightarrow eN$, and also present a brief discussion on some next-generation CLFV experiments for these channels. These three rare muon channels are complementary CLFV searches and may be used in combination to measure and further constrain the model-independent BSM parameters Λ and κ .

1.4.1 $\mu \rightarrow e\gamma$ Experiments

The state-of-the-art $\mu \rightarrow e\gamma$ experiment is the MEG Experiment at Paul Scherrer Institute (PSI) in Zurich, Switzerland. MEG was proposed in 1999, began taking physics data in 2008, and reported the full analysis of its data in 2016, where the limit on the $\mu^+ \rightarrow e^+\gamma$ BR was set at 4.2×10^{-13} [33]. The data analyzed for the final analysis contained 7.5×10^{14} stopped muons on target.

All modern rare muon experiments generally start with a muon beam directed onto a stopping

Table 1.1: Experimentally determined limits on branching ratios for a selection of CLFV processes [11]. All listed results below have a confidence level of 90%.

Process	BR Upper Limit	Experiment
$\mu^- \text{Au} \rightarrow e^- \text{Au}$	$< 7.0 \times 10^{-13}$	SINDRUM II
$\mu^- \text{Ti} \rightarrow e^- \text{Ti}$	$< 6.1 \times 10^{-13}$	SINDRUM II
$\mu^+ \rightarrow e^+ e^- e^+$	$< 1.0 \times 10^{-12}$	SINDRUM
$\mu^+ e^- \rightarrow \mu^- e^+$	$< 8.3 \times 10^{-11}$	SINDRUM
$\mu^+ \rightarrow e^+ \gamma$	$< 4.2 \times 10^{-13}$	MEG
$\tau \rightarrow e\gamma$	$< 3.3 \times 10^{-8}$	BaBar
$\tau \rightarrow \mu\gamma$	$< 4.4 \times 10^{-8}$	BaBar
$\tau \rightarrow eee$	$< 2.7 \times 10^{-8}$	Belle
$\tau \rightarrow \mu\mu\mu$	$< 2.1 \times 10^{-8}$	Belle
$\pi^0 \rightarrow \mu e$	$< 3.6 \times 10^{-10}$	KTeV
$K_L^0 \rightarrow \mu e$	$< 4.7 \times 10^{-12}$	BNL E871
$B^0 \rightarrow \mu e$	$< 2.8 \times 10^{-9}$	LHCb

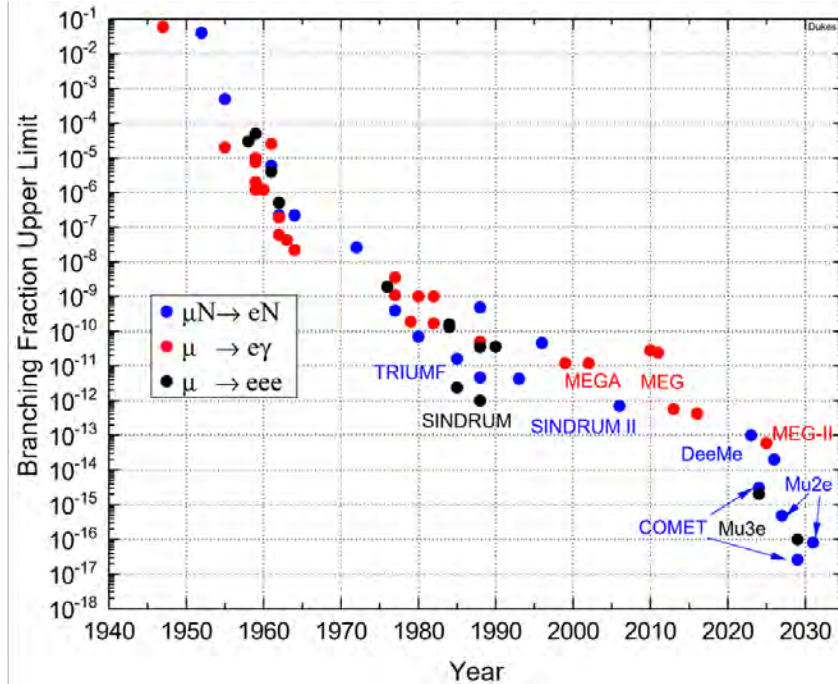


Figure 1.8: Summary of experimentally determined CLFV muon BR limits. Overview of experimentally determined branching ratio limits placed on CLFV processes since 1940 and the outlook of planned future experiments [32]. The only processes included in this plot are rare muon decay modes.

target, where muons are brought to rest and may be monitored by detectors. MEG is no different; the $\pi E5$ beamline at PSI served a continuous μ^+ beam to the experiment, with a muon stopping rate of $3 \times 10^7 \mu^+/s$ [33]. The muon stopping target is a thin, plastic ellipse, which is surrounded by a superconducting magnet. The superconducting COBRA magnet generates a graded magnetic field to guide the positrons coming from the stopping target towards the tracking detectors. To detect and track the positrons, there is a series of drift chambers and a timing counter mounted near the stopping target. The 16 gas-filled drift chambers that are used are aligned in a semicircle and use sensitive trip wires to detect passing charged particles. The timing counter is comprised of a series of scintillating bars read out by PMTs. To detect the photon from the stopping target, a liquid xenon scintillation detector also forms a semicircle around the stopping target. See Figure 1.9 below for a diagram of the MEG experiment detector apparatus.

Unlike $\mu \rightarrow e$ conversion experiments, $\mu \rightarrow e\gamma$ experiments are searching for a two-body decay. This feature of the decay can be exploited and used as criteria for event selection: the electron and photon from the stopped muon decay must be reconstructed back to the same point in space and time to be considered a signal event. A two-body decay from a muon at rest means that both the electron and the photon will carry off half of the muon rest energy, $E = 52.83 \text{ MeV}$. With the positron trajectory and momentum from the drift chambers, positron impact time and position from the timing counter, and the photon interaction vertex, interaction time, and energy from the liquid




Figure 1.9: Diagram of the MEG experiment detectors. A schematic view of the MEG detector apparatus with labels for the major components. A cross-sectional view is shown on the right. The incoming muon beam is shown in green coming in from the left and a simulated event is shown as the solid red and dashed blue lines [33].

xenon scintillation detector, MEG is able to efficiently reconstruct events from the stopping target and search for signal positrons.

The two main types of backgrounds for $\mu \rightarrow e\gamma$ experiments are radiative muon decay (RMD) and accidental backgrounds. RMD occurs when the muon decays with the added production of neutrinos, conserving lepton number, $\mu^+ \rightarrow e^+\gamma\nu_e\bar{\nu}_\mu$. When the neutrinos carry away only a small amount of energy, this process can appear to be a back-to-back positron and photon emission from the stopping target. RMD scales with the number of stopped muons and may be mitigated by improving the energy resolution of future detectors. Accidental backgrounds are formed by the coincidence of two independent processes that can form positrons and photons. For example, a positron from a Michel decay could be detected at the same time as a photon from an e^+e^- pair annihilation. This background also scales with the muon stopping rate and can be reduced with improved energy resolution [33].

The next generation $\mu \rightarrow e\gamma$ experiment will be MEG-II, an upgrade to MEG. The MEG-II experiment will include several upgrades to MEG: some of the PMTs in the MEG liquid xenon scintillation detector will be replaced with VUV MPPCs, the segmented MEG drift chamber system will be replaced with a large, modular, central drift chamber, the timing counter resolution will be upgraded by replacing scintillator bars with scintillator tiles, and a new radiative decay counter will be added to reduce accidental background [34]. With the aforementioned upgrades, MEG-II aims to enhance its sensitivity and reach a $\text{BR}(\mu^+ \rightarrow e^+\gamma) < 6 \times 10^{-14}$.

1.4.2 $\mu \rightarrow eee$ Experiments

The last experiment to set an upper limit for $\mu \rightarrow eee$ was the SINDRUM experiment in 1987, where it was reported that $\text{BR}(\mu \rightarrow eee) < 1.0 \times 10^{-12}$ [35]. This result was actually produced using data from the SINDRUM II apparatus, which included two drift chambers and a series of both scintillator and Cherenkov hodoscopes to track electrons and positrons ejected from the stopping target. Similar to our other rare muon experiments, the detectors are housed in a superconducting solenoid that generates an electric field to steer particles from the stopping target towards the detectors.

Much like $\mu \rightarrow e\gamma$ and $\mu \rightarrow e$ conversion experiments, $\mu \rightarrow eee$ experiments also benefit from stopping muons in a target to constrain the total energy and momentum of the final state particles. For Mu3e, this means that the sum of all decay particle momentum should vanish and, in this case, the sum of the decay particle energies should be equal to the muon mass. In a $\mu \rightarrow eee$




Figure 1.10: $\mu^+ \rightarrow e^+e^-e^+$ SM Feynman diagram. This tree-level SM Feynman diagram depicts a $\mu \rightarrow e^+e^-e^+$ decay [36].

decay, the maximum amount of energy that an electron or positron can have is half of the muon rest energy, so detectors for this search should be tuned for low momentum electrons and positrons. For Mu3e, an antimuon beam is used, so the signal event will be two positrons and one electron from a common vertex.

The backgrounds for $\mu \rightarrow eee$ experiments are similar to those of $\mu \rightarrow e\gamma$ experiments. There is a lepton flavor conserving muon decay with the added production of neutrinos that can mimic a signal event or a combination of coincident accidental processes can combine to look like signal event decay products. The process $\mu^+ \rightarrow e^+e^-e^+\bar{\nu}_\mu\nu_e$ has a BR of 3.4×10^{-5} [37]. This background can look like a signal event if the neutrinos carry off very little energy. Accidental backgrounds may be formed by the overlap of positrons from two Michel decays and an electron from Compton scattering, for example.

The next generation experiment for $\mu \rightarrow eee$ searches is the Mu3e experiment at PSI. Mu3e was proposed in 2012 and will utilize the same $\pi E5$ beamline of the MEG and MEG-II experiments. There is a planned upgrade to the beamline which will increase the intensity of stopped muons to

$2 \times 10^9 \text{ Hz}$ [37]. Mu3e will use four layers of high voltage tracking detectors and two different kinds of timing detectors, a scintillating fiber hodoscope and scintillator tiles for precise position and timing information on electrons and positrons in each event. The Mu3e experiment will proceed in two phases, reaching a final sensitivity of 1×10^{-16} after upgrading the beamline in 2025. Figure shows a diagram of the Phase I Mu3e detector apparatus.




Figure 1.11: Diagram of the Mu3e experiment detectors. A schematic view of the Mu3e detector apparatus with labels for the major components. A cross-sectional view is shown on the right. Two additional recurl stations will be added in phase II. The total length of the full detector is around 180 cm and the diameter is about 15 cm. The detector is inserted into a 100 cm diameter and 300 cm long superconducting solenoid providing a uniform magnetic field of 1 T [36].

1.4.3 $\mu N \rightarrow eN$ Experiments

The most stringent limits on the rate of $\mu \rightarrow e$ conversion come from the SINDRUM II experiment, which operated at PSI from 1993 to 2000. The SINDRUM II collaboration reported several results for $\mu \rightarrow e$ conversion searches using various stopping target materials. The final result from this experiment was published in 2006, where the final search was conducted using a gold muon stopping target and the reported upper limit was $R_{\mu e}(\mu^- \text{Au} \rightarrow e^- \text{Au}) < 7.0 \times 10^{-13}$ [38]. The data used for this analysis contained 4.4×10^{13} stopped muons on target.

The signal in $\mu \rightarrow e$ conversion searches is a monoenergetic CE, the energy of which is equal to

$$E_{CE} = m_{\mu}c^2 - E_B(Z) - \mathcal{R}_N(A) \quad (1.15)$$

where $m_{\mu}c^2$ is the rest energy of the muon, E_B is the atomic binding energy which depends on the number of protons in the stopping target nucleus Z , and \mathcal{R}_N is the nuclear recoil energy which depends on the number of nucleons in the nucleus A . For many elements, the atomic binding energy and nuclear recoil energy values are small, so the CE energy is approximately equal to the muon

rest mass. For aluminum, $E_{CE} = 104.97$ MeV [39]. The SINDRUM II experiment used several different stopping targets: lead, titanium, and gold.

Like the other rare muon searches that we have mentioned, $\mu \rightarrow e$ conversion searches begin with stopping muons in a target. However, unlike the two muon decay experiments which we just discussed, $\mu \rightarrow e\gamma$ and $\mu \rightarrow eee$, $\mu \rightarrow e$ conversion searches use a μ^- beam rather than a μ^+ beam. In SINDRUM II, the metal stopping target was placed in the center of a superconducting solenoid and detectors for observing CEs surrounded the stopping target. Two sets of drift chambers and two sets of hodoscopes were used to track particle trajectories and record timing, which allowed for CE track reconstruction. Using a beam of negatively charged μ^- instead of μ^+ enables the muons to be captured onto a metal stopping target.

$\mu \rightarrow e$ conversion searches only have one particle in the final state of a signal event, so there are no accidental backgrounds, but a number of other processes contribute final state electrons that can appear to originate from the stopping target with the right CE energy. Two types of backgrounds come from stopped muons: muon decay-in-orbit (DIO) and muon nuclear capture. Several other types of backgrounds come from beam contamination or the environment. A detailed discussion of the backgrounds that $\mu \rightarrow e$ conversion searches face and how the different types are mitigated will follow in the next chapter.

Currently, two next generation experiments are being constructed to search for $\mu \rightarrow e$ conversion on a nucleus: the Mu2e Experiment at Fermilab and the COMET Experiment at J-PARC. The Mu2e Experiment is the topic of this thesis, and as such, a much more detailed discussion of the experimental backgrounds, detector design, and expected experimental sensitivities will follow in the next chapters. The expected single event sensitivity (SES) of Mu2e is about 3×10^{-17} , about four orders of magnitude of improvement compared to SINDRUM II.

The COMET experiment is another next generation $\mu \rightarrow e$ conversion experiment, a sister experiment to Mu2e. COMET is being constructed at J-PARC in Osaka, Japan and will proceed in several phases of operation. The first phase, Phase- α , will be conducted primarily for beam commissioning. In Phase- α , the pion capture solenoid and the detector solenoid will be missing. The positioning of the proton beam and muon yield and profile measurements will be the only data collected in this phase [40]. The second phase, Phase-I, will include the pion capture solenoid, the detector solenoid, and half of the muon transport solenoid. Phase-I will produce a muon beam quality measurement and a limit on $R_{\mu e}$ with an intermediate sensitivity on the order of $\mathcal{O}(10^{-15})$ [41]. In the detector solenoid, the CyDET system will monitor the stopping target using a cylindrical drift chamber and a series of hodoscopes for tracking particle position and timing. The final phase, Phase-II, will use the full COMET detector and solenoid designs, reaching the final expected sensitivity of 2×10^{-17} . In Phase-II, the second half of the muon transport solenoid will be installed, the beam power is increased, and the StrECAL straw tube tracker and electron calorimeter

detector will be used at the end of the detector solenoid to measure $\mu \rightarrow e$ conversion [41]. Both the COMET and Mu2e experiments are expected to reach comparable sensitivities in the search for $\mu \rightarrow e$ conversion.

References

- [1] G. Aad, *et al.* (ATLAS Collaboration), *Physics Letters B* **716**, 1 (2012).
- [2] P. Zyla, *et al.* (Particle Data Group), *Prog. Theor. Exp. Phys.* **2020**, 083C01 (2020), and 2021 update.
- [3] B. Odom, D. Hanneke, B. D’Urso, and G. Gabrielse, *Phys. Rev. Lett.* **97**, 030801 (2006).
- [4] Y. Fukuda, *et al.* (Super-Kamiokande Collaboration), *Phys. Rev. Lett.* **81**, 1562 (1998).
- [5] T. Kajita, in *Proceedings of the Japan Academy, Series B*, Vol. 86 (2010) pp. 303–321.
- [6] S. Abe, *et al.* (KamLAND Collaboration), *Phys. Rev. Lett.* **100**, 221803 (2008).
- [7] R. Aaij, *et al.* (LHCb Collaboration), *Nature Physics* **18**, 277 (2022).
- [8] B. Abi, *et al.* (Muon $g - 2$ Collaboration), *Phys. Rev. Lett.* **126**, 141801 (2021).
- [9] T. Aaltonen, *et al.* (CDF Collaboration), *Science* **376**, 170 (2022).
- [10] Y. Kuno and Y. Okada, *Rev. Mod. Phys.* **73**, 151 (2001).
- [11] L. Calibbi and G. Signorelli, *Riv. Nuovo Cim.* **41**, 71 (2018).
- [12] M. Ardu and G. Pezzullo, *Universe* **8**, 299 (2022).
- [13] M. Kobayashi and T. Maskawa, *Prog. Theor. Exp. Phys.* **49**, 652 (1973).
- [14] N. Cabibbo, *Phys. Rev. Lett.* **10**, 531 (1963).
- [15] B. Pontecorvo, *Sov. Phys. JETP* **6**, 429 (1957).
- [16] R. Davis, D. S. Harmer, and K. C. Hoffman, *Phys. Rev. Lett.* **20**, 1205 (1968).

- [17] Y. Suzuki, in *11th Moriond Workshop: Tests of Fundamental Laws in Physics* (1991) pp. 427–435.
- [18] Q. R. Ahmad, *et al.* (SNO Collaboration), *Phys. Rev. Lett.* **87**, 071301 (2001).
- [19] S. Weinberg, *Phys. Rev. Lett.* **43**, 1566 (1979).
- [20] W. J. Marciano, T. Mori, and J. M. Roney, *Annual Review of Nuclear and Particle Science* **58**, 315 (2008).
- [21] A. de Gouvêa and P. Vogel, *Prog. Part. Nucl. Phys.* **71**, 75 (2013).
- [22] B. W. Lee and R. E. Shrock, *Phys. Rev. D* **16**, 1444 (1977).
- [23] A. de Gouvêa and N. Saoulidou, *Annual Review of Nuclear and Particle Science* **60**, 513 (2010).
- [24] J. Aysto, *et al.*, “Physics with low-energy muons at a neutrino factory complex,” (2001), arXiv:hep-ph/0109217 [hep-ph; hep-ex].
- [25] L. Calibbi, A. Faccia, A. Masiero, and S. K. Vempati, *Phys. Rev. D* **74**, 116002 (2006).
- [26] L. Calibbi, D. Chowdhury, A. Masiero, K. M. Patel, and S. K. Vempati, *Journal of High Energy Physics* **11**, 1 (2012).
- [27] J. M. Arnold, B. Fornal, and M. B. Wise, *Phys. Rev. D* **88**, 035009 (2013).
- [28] R. Harnik, J. Kopp, and J. Zupan, *Journal of High Energy Physics* **3**, 1 (2013).
- [29] M. Blanke, A. J. Buras, B. Duling, S. Recksiegel, and C. Tarantino, “FCNC Processes in the Littlest Higgs Model with T-Parity: an Update,” (2009), arXiv:0906.5454 [hep-ph].
- [30] H.-C. Cheng and I. Low, *Journal of High Energy Physics* **2003**, 051 (2003).
- [31] C.-H. Lee, P. S. B. Dev, and R. N. Mohapatra, *Phys. Rev. D* **88**, 093010 (2013).
- [32] E. C. Dukes, “Mu2e-II: The Mu2e Experiment in the PIP-II Era,” (NuFact 2021 2021-09-08), mu2e-doc-39318-v3.
- [33] A. M. Baldini, *et al.* (MEG Collaboration), “Search for the Lepton Flavour Violating Decay $\mu^+ \rightarrow e^+$ with the Full Dataset of the MEG Experiment,” (2016), arXiv:1605.05081 [hep-ex].
- [34] A. M. Baldini, *et al.* (MEG II Collaboration), *Eur. Phys. J. C* **78**, 380 (2018).
- [35] U. Bellgardt, *et al.* (SINDRUM Collaboration), *Nucl. Phys. B* **299**, 1 (1988).

- [36] A. Bravar, in *SciPost Physics Proceedings* (2019) doi:10.21468/SciPostPhysProc.1.037.
- [37] K. Arndt, *et al.* (Mu3e Collaboration), Nucl. Instrum. Methods Phys. Res. A: Accel., Spectrom., Detect. Assoc. Equip. **1014**, 165679 (2021).
- [38] W. Bertl, *et al.* (SINDRUM-II Collaboration), Eur. Phys. J. C **47**, 337 (2006).
- [39] L. Bartoszek and *et al.* (Mu2e Collaboration), *Mu2e Technical Design Report*, Tech. Rep. (Fermilab, 2015) arXiv:1501.05241 [ins-det] .
- [40] M. Moritsu, Universe **8**, 196 (2022).
- [41] R. Abramishvili, *et al* (COMET Collaboration), Prog. Theor. Exp. Phys. **2020** (2020), 033C01.

Chapter 2 The Mu2e Experiment

The Mu2e Experiment is currently being constructed on the muon campus at Fermilab in Batavia, Illinois. As the context of the previous chapter suggests, Mu2e is searching for evidence of CLFV by aiming to observe neutrinoless $\mu \rightarrow e$ conversion in the presence of a nucleus. The Mu2e experiment will set an upper limit on the $\mu \rightarrow e$ conversion process, with an expected sensitivity of $R_{\mu e} < 6.2 \times 10^{-16}$ at a 90% confidence level [1] after Run I. This chapter will introduce the Mu2e Experiment, starting with the historical design premise, and following through with a discussion of the sources of experimental backgrounds that the experiment will face. The next chapter will detail the experimental setup and detector design.

2.1 Conceptual Design

The concept for the Mu2e experiment is largely based on the proposed design for the MECO experiment at Brookhaven [2], and the earlier MELC experiment at Moscow Meson Factory. These designs detailed three main steps for performing a $\mu^- \rightarrow e^-$ conversion experiment. First, use a solenoid magnet system to produce and steer an intense muon beam. Second, bring the muons to rest on a stopping target. Last, use an annular-shaped detector apparatus around the stopping target to observe CEs. This procedure is also the basis for Mu2e.

Mu2e begins with the generation of a beam of negative muons, μ^- . The FNAL accelerator complex provides an 8 GeV proton beam, in which case muons can be produced by directing protons onto a production target. When protons interact with a production target, many particles are produced, a majority being negative pions, π^- . Pions are short-lived particles that rapidly decay to muons in flight, with over 99.9% of pions decaying as $\pi^- \rightarrow \mu^- \nu_\mu$ with a mean lifetime of 26 ns [3]. The solenoid system provides a graded magnetic field in the production region that guides negatively charged particles into a concentrated beamline while sweeping positively charged remnants away.

A secondary, low momentum muon beamline can now be directed onto a stopping target to form muonic atoms with the nuclei of the stopping target material. When a negatively charged muon stops in a target, it quickly cascades down to the ground state, within about 10^{-13} s [4]. Muonic atoms will decay, be captured, or convert after a short time which depends on the atomic number of the muonic nucleus. For aluminum, muonic atoms have a mean lifetime of 864 ns [4]. Muonic

atom lifetimes range from roughly 100 ns to 2 μ s, with lighter nuclei typically living longer than heavier nuclei.

As mentioned in the previous chapter, the signal of $\mu^- \rightarrow e^-$ conversion on a nucleus is a single, monoenergetic electron called a CE. The energy of the CE is generally close to the muon rest energy, but depends on the choice of the muon stopping target:

$$E_{CE} = m_\mu c^2 - E_B(Z) - \mathcal{R}_N(A) \quad (2.1)$$

where $m_\mu c^2$ is the muon rest mass, 105.66 MeV/ c^2 , $E_B(Z)$ is the atomic binding energy of the muon, and $\mathcal{R}_N(A)$ is the nuclear recoil energy. For aluminum, the conversion electron energy is 104.97 MeV [5]. Any process that may produce an electron around 105 MeV is a background for this experiment.

Instead of directly measuring the BR of $\mu N \rightarrow e N$, $\mu \rightarrow e$ conversion experiments measure $R_{\mu e}$, which is the ratio of neutrinoless $\mu \rightarrow e$ conversion to all muon nuclear captures [6]:

$$R_{\mu e} = \frac{\Gamma[\mu^- + A(Z, N) \rightarrow e^- + A(Z, N)]}{\Gamma[\mu^- + A(Z, N) \rightarrow \nu_\mu + A(Z - 1, N)]} = \frac{\Gamma[\mu \rightarrow e \text{ conversion}]}{\Gamma[\text{nuclear capture}]} \quad (2.2)$$

While $R_{\mu e}$ is not a BR, it is often referred to as such. From here, the calculated SES is defined as the conversion rate for which there is expected to be one signal event. In general, the rate for conversions on nuclei depends on many complicated factors relating to the wavefunctions of the nucleus, bound initial state particle, and final state particle. When a muon is captured by the nucleus in the stopping target, the wavefunction of the muon is much larger than that of the electrons that typically surround the nucleus, and thus the wavefunctions of the muon and the nucleus largely overlap. This creates theoretical uncertainty on the wavefunction on the muon, and taking a ratio as is done in the evaluation of $R_{\mu e}$ allows for the cancellation of most muon wavefunction contributions that have uncertainty [7].

The use of $R_{\mu e}$ has another feature that $\mu \rightarrow e$ conversion experiments can benefit from: dependence on atomic number Z of the stopping target material. The conversion rate may be calculated in great detail [7], where many of the contributing factors arise from the interactions within the nucleus, interactions relating to photon exchange with the nucleus upon conversion, and particle waveforms. Numerical evaluations of these factors require choosing a theory with CLFV, and so different conversion rates can be calculated for different BSM theories. Figure 2.1 shows the conversion rate versus atomic number for several different BSM physics scenarios. For low Z , different models generally result in a similar conversion rate, but for higher Z , the conversion rates that we find for different BSM scenarios become more discernible for each other. As a result, it may be beneficial for $\mu \rightarrow e$ conversion experiments to take several physics runs with different stopping




Figure 2.1: Dependence of $R_{\mu e}$ on Z . $R_{\mu e}$ conversion ratios as a function of the atomic number Z of the stopping target material. The different lines represent three different BSM scenarios where either photonic dipole, scalar, or vector terms dominate and lead to enhanced rates of CLFV. The conversion ratios are normalized by the conversion ratio in aluminum nuclei ($Z = 13$) [7].

targets. In the past, SINDRUM II used lead ($Z_{Pb} = 82$), titanium ($Z_{Ti} = 22$), and gold ($Z_{Au} = 79$). The most stringent limit set by SINDRUM II is the limit using the gold target, as was mentioned in Chapter 1 [8]. Mu2e and COMET will use an aluminum stopping target, ($Z_{Al} = 13$). Aluminum has been chosen as the stopping target material for Mu2e to balance having a long muonic atom lifetime to control backgrounds while maintaining a high energy CE signal as aluminum has a low cost in the binding and recoil energy exchanged with the nucleus upon capture on the stopping target. The expected sensitivity of Mu2e will improve on the limit set by SINDRUM II by 4 orders of magnitude.

2.2 Experimental Backgrounds

As promising as the plan for carrying out the Mu2e experiment is, there are many experimental backgrounds that $\mu \rightarrow e$ conversion experiments must mitigate before they have any chance of observing CLFV. The backgrounds can be categorized into several groups:

- Intrinsic backgrounds
- Prompt beam processes
- Delayed beam processes

- Cosmic ray induced backgrounds
- Reconstruction errors

This section will discuss the sources of each category of background and describe general methods that may be used to mitigate each category. Specific mitigation methods employed by Mu2e will be mentioned in the next chapter where detector design will be covered. The first three background categories scale with beam intensity, while the last two scale with time.

2.2.1 *Intrinsic Backgrounds*

Intrinsic backgrounds occur as a result of stopping muons on a target, and thus scale with beam intensity. The two processes that are considered to be intrinsic backgrounds to $\mu \rightarrow e$ conversion experiments are muon DIO and radiative muon capture (RMC). Muon DIO events occur when a muon is captured on the nucleus, but undergoes a Michel decay while it is captured. When a free muon undergoes a Michel decay, $\mu \rightarrow e\nu_\mu\bar{\nu}_e$, the resulting electron can only have at most half of the muon's rest energy. However, if a muon is bound in a nuclear orbit, the final state electron can exchange energy with the nucleus. The possibility for this energy exchange smears the endpoint of the Michel spectrum up to the rest energy of the muon, making muon DIO a possible background for $\mu \rightarrow e$ conversion. At the kinematic limit of the bound muon DIO decay, the neutrinos carry away no energy and the daughter electron appears with an energy equal to that of the stopped muon, simulating a CE.

RMC occurs when the bound muon is captured by the stopping target nucleus, changing the atomic number, $\mu^- N(A, Z) \rightarrow \gamma\nu_\mu N(A, Z - 1)$. For aluminum, this process is $\mu^- \text{Al} \rightarrow \gamma\nu_\mu \text{Mg}$. The high energy photon that is produced by this capture can convert to an e^-e^+ pair, where the electron can be near the energy of a CE. Photons can convert internally while virtual, or externally in the stopping target or in other surrounding materials. However, the endpoint energy of the RMC spectrum depends on the choice of stopping target material. If the stopping target material is chosen such that there is a significant mass difference between the stopping target nuclei and the daughter nuclei that is present after capture, the RMC endpoint energy can be reduced and the resulting electrons can be separated from CE. For aluminum, the RMC endpoint energy is 101.9 MeV, over 3 MeV below the CE energy [6].

2.2.2 *Prompt Beam Processes*

Backgrounds of this type originate from particles in the beam, and produce electrons that are in-time with the incoming muons that will be stopped on target. The background electrons from these processes cannot be temporally separated from CEs. Four of these processes are radiative pion capture (RPC), pion and muon decay-in-flight (DIF), and beam electrons.

Muon beams, like the one used for Mu2e, are created by channeling daughter muons from pion decays using magnetic fields. Since the muon beam is created using parent pions, there is a nonzero probability that pions produced at the production target can travel through the transport solenoids and produce decay products that mimic CEs. Most low energy muon beams have a significant amount of pion contamination. One way that these beam pions can be a source of background is RPC, where the pion is captured on the stopping target or surrounding material, $\pi^- N(A, Z) \rightarrow \gamma N(A, Z - 1)$. If an asymmetric $\gamma \rightarrow e^- e^+$ conversion follows, the electron can carry energy up to 130 MeV [1]. Mu2e will use a pulsed proton beam, which means that RPCs can occur either in-time with the beam or out-of-time. Since the pion lifetime is so short, in-time RPCs that arrive with the beam pulse can be mitigated by delaying the live-time CE search window relative to the arrival of the beam pulse. Out-of-time RPCs come from pions in between beam pulses, so they cannot be mitigated using time. Instead, requiring extinction of protons between subsequent beam pulses controls these pions.

Similar to RPCs, pion DIF events originate from pions contaminating the muon beamline. In the DIF case, the pions decay before they arrive at the stopping target or in the region of the stopping target, $\pi^- \rightarrow e^- \bar{\nu}_e$. Muons can also decay before reaching the stopping target, generating muon DIF backgrounds through Michel decays, $\mu^- \rightarrow e^- \nu_\mu \bar{\nu}_e$. When pions and muons at rest undergo these decays, the daughter electrons in both cases have a much lower energy than what can be concerning for CEs. However, since the pions and muons in the beamline are moving at relativistic speeds, the energies of daughter electrons are boosted. For example, a muon moving at $0.6c$ with a momentum of 79 MeV can produce electrons at 105.6 MeV [6]. Using a pulsed beam and a delayed live-time window helps to mitigate DIF backgrounds. See Figures 3.3 and 3.4 for more details on the pulsed beam structure.

Beam electrons are produced when pions or muons decay far upstream before the detector solenoid and travel with the beam down to the detectors. There are many processes that can produce electrons in this way, for example, early DIF events near the production target, or neutral pions decaying to photons, which then convert to $e^- e^+$ pairs. Again, using a pulsed beam and a delayed live-time window helps mitigate this source of background. In addition, the solenoids that carry the beam can be used to mitigate this background by selecting for only low-energy muons with collimators, magnetic field strength, and steering around bends in the magnets.

2.2.3 Delayed Beam Processes

Any slow moving particles that are produced in the beamline have the potential to eventually drift down to the stopping target. Daughter electrons from these particles can be produced and detected at virtually any time, not just in-time with beam pulses, and so they are called delayed beam processes. In muon beams, antiprotons produced in the production target via the process

$pp \rightarrow ppp\bar{p}$ are the main source of this background. The 8 GeV primary proton beam used to produce the muon beam at Fermilab is energetic enough to produce antiprotons that are concerning background sources. Antiprotons do not decay, so they are long-lived and can travel down to the stopping target. They are negatively charged, so they can also interact with the stopping target and annihilate, generating many possible secondary particles.

This type of background cannot be further reduced using the beam pulse structure and livetime window since they can reach the detectors at any time. Antiprotons with momenta less than 100 MeV/c traveling at $0.1c$ or slower can take microseconds to reach the detector region in Mu2e from the production region [6]. Instead, slow moving antiprotons are mitigated by preventing them from traveling to the stopping target. Thin absorbers in the beamline can be used to remove these particles from the beam. These absorbers are designed to minimize beam loss by optimizing the thickness of the absorber and appropriately selecting the absorber material.

There are some delayed beam processes that are mitigated by using a pulsed beam structure, namely delayed RPC. The stopped-pion stop-time distribution depends on the arrival-time distribution of the protons at the production target in initial Mu2e simulations, thus RPC backgrounds have both "in-time" and "out-of-time" contributions. In-time RPC backgrounds arise from protons that arrive at the production target as part of the 1695 ns beam pulse, while out-of-time contributions correspond to protons that arrive at the production target between the micropulses. [6] Mu2e requires extinction between beam pulses on the level of 10^{-10} to control out-of-time contributions to the RPC background. Recall that in-time contributions may be controlled by delaying the start of the live-time search window.

2.2.4 Cosmic Ray Induced Backgrounds

Cosmic rays present significant background to $\mu \rightarrow e$ conversion experiments; for Mu2e this is the dominant background. Since there is only one final state electron in the CE signal, any process that can produce an electron with 105 MeV/c is a potential background. Cosmic ray muons are energetic enough to produce electrons at the energy of a CE, can occur at any time, and can also produce tracks originating from the stopping target. Many different cosmic ray muon scenarios can generate an event that looks like a CE. This background is irreducible and scales with the live time of the experiment, as cosmic muons originate from the atmosphere. For Mu2e, cosmic ray background events are expected to happen about once per day if left unmitigated.

Some of the different cosmic ray scenarios that can mimic final-state CEs include [6]:

- Cosmic ray muon decays in the detector region or near the stopping target
- Cosmic ray muon interacts with the stopping target or surrounding materials

- Cosmic ray muon scatters in the detectors and is misidentified as an electron
- Cosmic ray muon enters the muon beamline with a special initial trajectory where it is misidentified

Cosmic ray identification can be improved by optimizing reconstruction algorithms, however, these events can be virtually indistinguishable from candidate CE events. Mitigation techniques against cosmic rays include shielding and active veto systems; Mu2e will use both. By using a combination of concrete shielding layers and an active, high-efficiency veto detector, this background can be kept to a reasonable level. A detailed description of the Mu2e Cosmic Ray Veto detector will follow in Chapter 4, as the main topic of this thesis is the aging of this detector.

2.2.5 Reconstruction Errors

Lastly, any time event reconstruction is performed, there is a risk for errors and misidentification. Though reconstruction errors are not an experimental background, they do limit the sensitivity reach of the experiment and can generate tracks that look like CEs. There are a number of different processes that can occur within the Detector Solenoid (DS) as a result of muon capture on the stopping target nucleus. Decaying excited nuclei can produce high-energy protons, neutrons, and photons. Ejected protons can leave large signals in the detectors, which can cause a reconstruction error. Ejected neutrons may be re-captured and result in additional photons. Any additional photons can undergo pair production or Compton scattering, which can generate electron tracks in the detectors and be reconstructed as conversion-like tracks.

False signals of this type can be mitigated by careful design of the tracking detectors and determination of track quality within reconstruction software. In addition, these particles can be prevented from reaching the tracking detectors altogether by using absorbers. The stopping target is surrounded by polyethylene absorbers to reduce the proton and neutron flux to the downstream detectors [9]. These absorbers also prevent damage to the downstream detectors that may be caused by highly ionizing protons.

2.3 Background Evaluation/Sensitivity Estimate

All of the abovementioned background processes may be simulated in the Mu2e environment to estimate the sensitivity of the experiment and how many background events are expected. Recently, a sensitivity estimation campaign was executed to estimate the sensitivity of Mu2e using only Run I data. There is also an estimated sensitivity using the full Mu2e dataset, which was used when determining the projected $R_{\mu e}$ discovery reach of the experiment.

The procedure that Mu2e employs to determine the optimal upper limit for $R_{\mu e}$ draws on Feldman-Cousins statistical analysis, which describes a rigorous method to determine confidence intervals for experiments with small signals [10]. In the Feldman-Cousins procedure, one considers an ensemble of identical experiments. In the Mu2e context, this ensemble is a large number of Mu2e experiments that all have the same number of protons on target, a fixed analysis procedure, and a fixed value of $R_{\mu e}$. Now, the ensemble of experiments collects data, signal and background events fluctuate between experiments. By analyzing the experimental outcomes of the ensemble, a two-sided limit for $R_{\mu e}$ is calculated based on the assumption that no signal events are observed. If the chosen initial value of $R_{\mu e}$ is small, then most experiments in the ensemble will not exclude $R_{\mu e} = 0$. However, if the chosen initial value of $R_{\mu e}$ is large, then most experiments in the ensemble will claim discovery. For some intermediate value of $R_{\mu e}$, half of the experiments in the ensemble claim discovery [11]. This value of $R_{\mu e}$ is the optimal discovery sensitivity of the analysis.

To determine the optimal $R_{\mu e}$ values for different Mu2e runs, Feldman-Cousins analysis is implemented alongside simulations of Mu2e data. For an experiment like Mu2e, discovery can be declared if the value $R_{\mu e} = 0$ is excluded from the limits that are determined at a 5σ level. To narrow in on the window of $R_{\mu e}$ that Mu2e has the power to explore, different choices of confidence levels can be used as assumptions to the Feldman-Cousins procedure. First, a 90% CL is assumed as the basis for the two-sided $R_{\mu e}$ limits, which gives a two-sided limit for $R_{\mu e}$ versus the number of observed signal events in the experiment. From here, the optimal signal window can be determined in momentum space. Next, the expected upper limit for $R_{\mu e}$ can be calculated for a 5σ discovery sensitivity using the optimized signal window. This analysis determines the SES, the upper limit on $R_{\mu e}$ in the absence of any observation, the discovery $R_{\mu e}$ value that may be claimed if signal events are observed, and the number of signal events that must be observed to claim a 5σ discovery [11–13].

2.3.1 Run I Sensitivity Estimate (SU2020 paper)

A detailed analysis was completed to determine the sensitivity for Run I and published in Ref. [1]. The expected 5σ discovery sensitivity is $R_{\mu e} = 1.2 \times 10^{-15}$, which requires observing five $\mu^- \rightarrow e^-$ conversion events within the signal region. In the absence of any signal, the upper limit from Run I data is $R_{\mu e} < 6.2 \times 10^{-16}$ at a 90% confidence level, assuming a total of 3.8×10^{19} protons on target. [1]. This is an improvement of over three orders of magnitude compared to the current experimental limit of $R_{\mu e}$ set by the SINDRUM II experiment.

The Run I data-taking plan for Mu2e begins with commissioning. During this time, the experiment will begin operating with an expected mean proton beam intensity of 1.6×10^7 protons per pulse. After commissioning, beam intensity will increase to deliver a mean of 3.9×10^7 protons per pulse [1]. About 75% of the total number of protons on target in Run I will be delivered in the

low intensity running mode, with the remaining 25% delivered in the high intensity running mode.

Overall, this estimate yields a total expected background of 0.11 ± 0.03 events [1]. All simulated events proceed through the trigger and are subject to event selection before they may qualify as $\mu^- \rightarrow e^-$ conversion events. First, pre-selections are applied to remove events that do not produce a sufficient number of hits in the tracker, do not have a trajectory that originates from the stopping target, or are not traveling in the correct direction. Next, artificial neural network-based algorithms are used to aid in particle identification and momentum reconstruction. Finally, only events within the optimized Run I signal momentum and time window $103.6 < p < 104.9$ MeV/c and $640 < T_0 < 1650$ ns are considered as conversion-like events and individual background contributions may be integrated over this window [1]. Table 2.1 shows the breakdown of contributions from different background modes to the total expected background count.

These background estimates use a recent branch of the Mu2e simulation framework, called SU2020, or Sensitivity Update 2020, named for the year in which the campaign began. The Mu2e simulation framework is based on Geant4, where an extensive simulation of the Mu2e environment and beamline has been created. Using these detector simulations, the timing response of the subdetectors and the effects of hit readout, digitization, and reconstruction can be evaluated. Specifically, the SU2020 branch uses Geant4 v10.15 [1]. A customized physics list, called "ShieldingM", is used to accurately consider cross sections and time dependence of physics processes in the Mu2e environment. The detector simulations and reconstruction algorithms assume that the detector has been perfectly aligned and that there are no dead channels in any subdetector.

Mu2e is currently under construction at Fermilab and will soon be entering installation and commissioning phases [14]. Many of the detector subsystems are currently undergoing vertical

Table 2.1: Run I background summary. Background summary and single event sensitivity for Mu2e Run I using the optimized signal momentum and time window, $103.6 < p < 104.9$ MeV/c and $640 < T_0 < 1650$ ns [1].

Channel	Mu2e Run I
Cosmic rays	0.046 ± 0.010 (stat) ± 0.009 (syst)
DIO	0.038 ± 0.002 (stat) $^{+0.025}_{-0.015}$ (syst)
Antiprotons	0.010 ± 0.003 (stat) ± 0.010 (syst)
RPC in-time	0.010 ± 0.002 (stat) $^{+0.001}_{-0.003}$ (syst)
RPC out-of-time	$(1.2 \pm 0.1$ (stat) $^{+0.1}_{-0.3}$ (syst)) $\times 10^{-3}$
RMC	$< 2.4 \times 10^{-3}$
Decays in flight	$< 2 \times 10^{-3}$
Beam electrons	$< 1 \times 10^{-3}$
Total	0.105 ± 0.032
SES	$< 2.4 \times 10^{-16}$

slice testing and final fabrication jobs are underway. The Mu2e Experiment plans to take two run periods of data: Run I and Run II. Run I will take place in 2025 and 2026. Run I will end in January 2027, when Fermilab will begin a two-year-long shutdown for a high-intensity beamline upgrade, the LBNF/PIP-II upgrade [15]. Run II will begin in 2029 and last for three or four years, using the upgraded beamline [16]. The estimated sensitivity reach for $R_{\mu e}$ after each run will be discussed later in this chapter after describing the experimental backgrounds.

2.3.2 Full Sensitivity Estimate

In addition to estimating the sensitivity using only Run I data, the sensitivity of the full Mu2e experiment may be estimated using a combination of both data runs. Across the full lifetime of the experiment, it is expected that Mu2e will be delivered 3.6×10^{20} protons on target [17]. This number is assumed for the rest of the full sensitivity estimations detailed here. The most recent full sensitivity estimate campaign is called the CD3 sensitivity estimate and was performed in 2017 [17]. At this time, the Mu2e data-taking plan was not yet finalized, so the simulated number of protons on target was calculated using a beam intensity of 3.9×10^7 protons per pulse.

After optimizing all simulations available at this point, the combined analysis of Mu2e Run I and Run II data is expected to be able to set a final upper limit of $R_{\mu e} < 8 \times 10^{-17}$ at 90% CL [17, 18]. The CD3 estimated breakdowns for each background mode are included in Table 2.2 to illustrate the magnitude of contribution from each respective background mode on the total number of background events that are expected. This sensitivity optimization analysis was performed using a version of Mu2e code called Mu2eCCFC, which includes the multi-step, Feldman-Cousins based approach described above [12] and Geant4-based Mu2e detector environments for simulations.

Table 2.2: Total estimated background summary. Background summary and single event sensitivity for the full Mu2e Run I and II datasets using the signal momentum and time window, $103.85 < p < 104.9$ MeV/c and $700 < T_0 < 1695$ ns [17].

Channel	Mu2e Run I + II
Cosmic rays	0.209 ± 0.022 (stat) ± 0.055 (syst)
DIO	0.144 ± 0.028 (stat) ± 0.11 (syst)
Antiprotons	0.040 ± 0.001 (stat) ± 0.020 (syst)
RPC	0.021 ± 0.001 (stat) ± 0.002 (syst)
RMC	$0.000^{+0.004}_{-0.000}$
Muon decay-in-flight	< 0.003
Pion decay-in-flight	$0.001 \pm < 0.001$
Beam electrons	$(2.1 \pm 1.0) \times 10^{-4}$
Total	0.41 ± 0.13 (stat+syst)
SES	$(3.01 \pm 0.03$ (stat) ± 0.41 (syst)) $\times 10^{-17}$

2.4 Outlook Beyond Mu2e

Mu2e already has a flourishing upgrade plan beyond its lifetime, where major detectors will be replaced or greatly enhanced to increase resolution. After Run II of Mu2e data taking concludes, there are plans for the detectors to undergo extensive redesign, ultimately returning to operation in the form of the Mu2e-II experiment. Though Mu2e-II has not officially been confirmed, a dedicated team has been running simulations to demonstrate the experimental power these upgrades could produce. [19] A major change between Mu2e and Mu2e-II will be replacing Mu2e’s aluminum stopping target. Currently, Mu2e-II has been exploring different possibilities for stopping target materials, namely using a titanium stopping target. Titanium has a higher atomic number than aluminum, $Z_{\text{Al}} = 13$ and $Z_{\text{Ti}} = 22$, so Mu2e-II data may be used with Mu2e data to distinguish between the mechanism of CLFV in the event of observing $\mu^- \rightarrow e^-$ conversion. The upgrades that are proposed for Mu2e-II result in a sensitivity improvement of about one order of magnitude, calculated as $R_{\mu e} < 6.4 \times 10^{-18}$ at 90% CL [19].

References

- [1] F. Abdi, *et al.* (Mu2e Collaboration), “Mu2e Run I Sensitivity Projections for the Neutrinoless $\mu^- \rightarrow e^-$ Conversion Search in Aluminum,” (2022), arXiv:2210.11380 [hep-ex].
- [2] Y. K. Semertzidis (MECO Collaboration), Nucl. Phys. B, Proc. Suppl. **149**, 372 (2005).
- [3] P. Zyla, *et al.* (Particle Data Group), Prog. Theor. Exp. Phys. **2020**, 083C01 (2020), and 2021 update.
- [4] T. Suzuki, D. F. Measday, and J. P. Roalsvig, Phys. Rev. C **35**, 2212 (1987).
- [5] A. Czarnecki, X. G. i Tormo, and W. J. Marciano, Phys. Rev. D **84**, 013006 (2011).
- [6] L. Bartoszek and *et al.* (Mu2e Collaboration), *Mu2e Technical Design Report*, Tech. Rep. (Fermilab, 2015) arXiv:1501.05241 [ins-det] .
- [7] R. Kitano, M. Koike, and Y. Okada, Phys. Rev. D **66**, 096002 (2002).
- [8] W. Bertl, *et al.* (SINDRUM-II Collaboration), Eur. Phys. J. C **47**, 337 (2006).
- [9] R. H. Bernstein, Front. Phys. **7**, 1 (2019).

- [10] G. J. Feldman and R. D. Cousins, *Phys. Rev. D* **57**, 3873 (1998).
- [11] A. Gaponenko, “Optimizing for Sensitivity,” (Mu2e Backgrounds Meeting 2017-06-08), mu2e-doc-10709-v3.
- [12] R. Bernstein, “Mu2eCCFC Update: Discovery Window,” (Mu2e Backgrounds Meeting 2016-06-07), mu2e-doc-7549-v12.
- [13] R. Bernstein, “Comparison of Optimization in Momentum Since CD3,” (Mu2e Backgrounds Meeting 2017-05-04), mu2e-doc-9999-v7.
- [14] P. Derwent, “Director’s Review: Cost, Schedule & Processes,” (Director’s Review 2023-02-21), mu2e-doc-44593-v7.
- [15] M. Ball, *et al.* (PIP-II Collaboration), *The PIP-II Conceptual Design Report*, Tech. Rep. (Fermilab, 2017) [FERMILAB-DESIGN-2017-01;FERMILAB-TM-2649-AD-APC].
- [16] J. Miller, “Mu2e Mu2e-II: Search for Charged Lepton Flavor Violation,” (FNAL P5 Meeting 2023-03-16), mu2e-doc-44957-v3.
- [17] A. Gaponenko, “CD3 Backgrounds,” (DOE CD-3c Review 2017-07-13), mu2e-doc-7464-v13.
- [18] A. Edmonds, “The Mu2e Experiment: A Search for Charged Lepton Flavor Violation in Muons,” (University of Birmingham Seminar 2022-06-01), mu2e-doc-42307-v3.
- [19] K. Byrum, *et al.* (Mu2e-II Collaboration), in *2022 Snowmass Summer Study* (2022) arXiv:2203.07569 [hep-ex] .

Chapter 3 Mu2e Detector Overview

The detectors of the Mu2e Experiment have been carefully designed to maximize the observation potential for muon to electron conversion on the stopping target. Below, Figure 3.1 shows an overview of the Mu2e experimental setup. A large, 'S' shaped solenoid system produces a graded magnetic field to guide muons toward the stopping target. Beginning at the upstream end of the experiment, the proton beam enters the Production Solenoid (PS) incident onto the production target. The Transport Solenoid (TS) contains collimators and proton absorbers to direct the muon beamline onto the stopping target. The stopping target is located at the front of the Detector Solenoid (DS) and is surrounded by absorbers. The collimators and absorbers in the TS and DS are not depicted in Figure 3.1. Moving further downstream, the tracker and the calorimeter are also housed in the DS to detect conversion electrons that are ejected from the stopping target. At the far end of the DS, there is a Stopping Target Monitor (STM) to estimate the number of muons that were captured on the stopping target. Lastly, the Mu2e CRV surrounds the entire DS and half of the TS to provide an active shield against cosmic rays. The CRV is not shown in the figure below, and will be described in detail in Chapter 4. This chapter will provide brief overviews of the various detectors of the Mu2e experiment, with emphasis on how the design of each detector system aids in mitigating the experimental backgrounds that were discussed in Chapter 2.


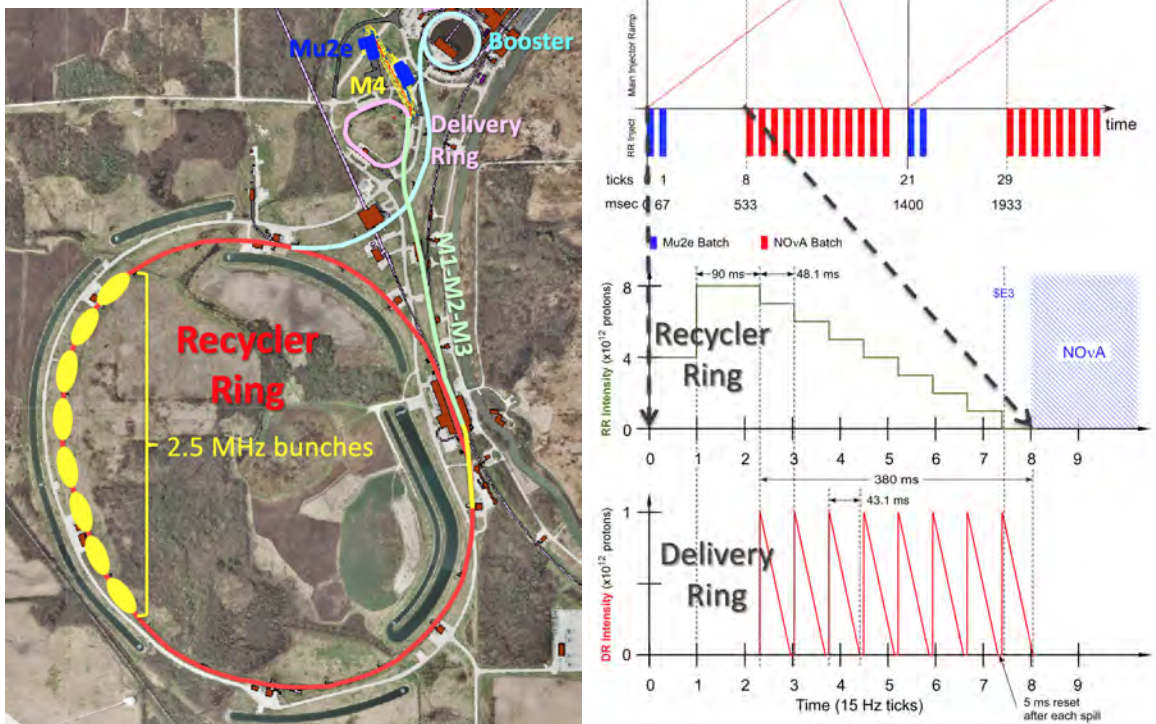


Figure 3.1: Mu2e detector overview. Schematic overview of the Mu2e apparatus with labels in blue for the magnetic field strengths generated by the solenoids at particular points [1].

3.1 Mu2e Beamline at Fermilab

Fermilab provides Mu2e with one of the most intense muon beamlines in the world. The Mu2e experimental hall is located in the muon campus, southwest of Fermilab’s Wilson Hall. The accelerator facility at Fermilab generates an 8 GeV, 8 kW proton beam to be delivered to the Mu2e production target in pulsed batches with a 1695 ns period between batches [2]. The batch structure has been chosen to reduce prompt background processes. In order to satisfy the beam requirements needed for Mu2e, the proton beam must undergo several manipulations to achieve the desired intensity and frequency of batches, as well as extinction between subsequent beam pulses.




(a) Fermilab accelerator facilities

(b) Proton pulse timing before the Mu2e beamline

Figure 3.2: Fermilab accelerator overview. Figure (a) shows a labeled map of the Fermilab accelerator facilities with the Booster, Recycler, and Delivery rings shown. The muon campus is represented by blue buildings, with the Mu2e building labeled. The M4 beamline connecting the Delivery ring to the Mu2e building is also labeled. Cartoon bubbles depicting proton pulse bunches in the Recycler ring are shown in yellow. Figure (b) gives an overview of the timing of proton pulses as they move through the Fermilab accelerator system. The three plots show the timing structure of how proton pulses are extracted from the Booster ring, rebunched in the Recycler ring, and then transferred to the Delivery ring, respectively [3].

The Fermilab proton beamline begins with accelerating H^- ions through a linac, foil-stripping them, and injecting the resulting protons into the Booster synchrotron. In the Booster, the protons are raised to a kinetic energy of 8 GeV. Booster protons are then extracted in batches and injected into the Recycler ring. During each accelerator cycle, two Booster proton batches

containing 4×10^{12} protons per batch are transported to the Recycler ring via the MI-8 beamline at an injection rate of 15 Hz. At this rate, the basic unit of time in the accelerator cycle is that 1 tick is equal to 66.7 msec. In the Recycler ring, the two Booster batches are rebunched to fit Mu2e’s beamline characteristics. As protons circulate around the Recycler ring, each Booster batch of protons separates into four bunches of approximately 100ns duration using a 2.5 MHz RF system [4]. This technique naturally creates narrow groups of protons with high levels of extinction between batches. However, extraction and injection kickers are also used to transfer beam from the Recycler ring to the Delivery ring, preserving the level of extinction from rebunching and ensuring that no out-of-time beam is transferred to the Delivery ring [2]. From the Recycler ring, one bunch at a time is transported to the Delivery ring, at a rate of one bunch every 48.1 msec. Once in the Delivery ring, the bunches contain 1×10^{12} protons per bunch, which can then be resonantly extracted into the beamline for Mu2e, called the M4 beamline. A single bunch circulating in the Delivery ring takes 43.1 msec to be resonantly extracted, ultimately forming a train of 25,000 proton pulses that are spaced in time by 1695 ns, the period of the Delivery ring. On average, the proton intensity is 3.9×10^7 protons per pulse. This process of extracting protons from the Booster ring, rebunching in the Recycler ring, and transferring from the Delivery ring to the M4 beamline repeats every 1.4 s, where Mu2e utilizes the part of the accelerator cycle between NOvA extraction periods, as shown in Figure 3.2(b) [3] To achieve the desired experimental sensitivity, a total of 4.7×10^{20} protons on target are needed over the lifetime of the experiment.



(a) Proton pulse timing within Mu2e beamline

Figure 3.3: Proton pulse timing overview. Timing profile between two subsequent proton pulses within the M4 beamline as they are delivered to the Mu2e production target [3].

The M4 beamline is a new facility that transports the proton beam to the production target in the Mu2e PS. Instrumentation within the M4 beamline measures the beam intensity, beam loss, and spatial profiles in the transverse planes to characterize the proton beam before it arrives at the tungsten production target. Additionally, the M4 beamline is equipped with deflecting AC

dipole magnets and a collimation system to further achieve high levels of extinction between proton pulses [2]. To satisfy the requirements of suppressing prompt background modes, an extinction level of 10^{-10} is required between pulses, illustrated in Figure 3.3(b) [3]. Here, extinction is defined as the ratio of the number of out-of-time protons to the number of in-time protons.

Separating subsequent narrow proton batches by a period of 1695 ns is a convenient, beneficial choice for Mu2e. This period between batches is much longer than the 864 ns lifetime of the muonic aluminum nuclei that form when muons are captured on the stopping target. In Figure 3.4, a complete timeline between two proton pulses is illustrated [5]. The live time window opens 700 ns after the arrival of a proton pulse, allowing prompt pion backgrounds like RPC time to decay. The window is then open for 900 ns, during the time which muons will decay from or be captured by the aluminum stopping target. Though most pions arrive early, within 200 ns, and have a short lifetime, the pion arrival and decay distribution has a long tail. This choice of the delayed start time and long live window length minimizes pion-induced backgrounds while optimizing the detection of muon decays. Simulations show that the pion background is suppressed by $\mathcal{O}(10^{11})$ by opening the live window at 700 ns [6]. Mu2e’s 1695 ns proton pulse repetition rate is much longer than the beam structure used in the past by SINDRUM II, improving the temporal separation between background and signal. This is one of many ways that Mu2e is improving upon the previous generation of $\mu^- \rightarrow e^-$ conversion experiments.




Figure 3.4: Live window and background timing overview. Timeline of two subsequent proton beam pulses with the live search window shown in between. The simulated pion flux is shown in pink, depicting prompt backgrounds. Muon arrival at the stopping target is shown in solid dark blue. The muon decay or capture rate is also shown over time in dashed light blue. The live search window is shown in orange, beginning at 700 ns and lasting for 900 ns [5].

3.2 Mu2e Solenoids

The Mu2e solenoid system steers and transports the muon beamline using a gradient magnetic field. There are three solenoids to the system, which are magnetically coupled: the PS,

the TS, and the DS. Each part of the system consists of multiple solenoid module sections, wound from superconducting, aluminum-stabilized Nb-Ti Rutherford cables, shown in Figure 3.5. The inner bore of the solenoid system will be evacuated to 1×10^{-4} Torr to minimize backgrounds from multiple scattering and prevent electrical sparks in the tracker straws. The solenoids generate a strong magnetic field of $\mathcal{O}(1\text{T})$ that decreases in strength along the beamline. The specific magnitude of the magnetic field at different locations along the solenoids combined with the 'S' shape of the entire solenoid system directs only low energy muons toward the stopping target.



Figure 3.5: Superconducting solenoid cable cross-section. A cross-section look at the superconducting, aluminum-stabilized Nb-Ti cable used to wind the Mu2e solenoids [7].

The Mu2e solenoid system also requires ancillary systems to maintain operation. A cryogenic cooling system, vacuum system, magnet power converters, magnetic field mapping system, and quench protection system have all been designed to support the Mu2e solenoids. For details on ancillary systems, refer to the Mu2e Technical Design Report [2].

3.2.1 Production Solenoid

The PS is a challenging magnet that must withstand a harsh radiation environment from the incoming proton beam and produce a strong magnetic field. In total, the PS consists of three superconducting coils that are bolted together to form a single cold-mass assembly. In total, the PS is 4 m long [8]. The inner bore of the PS is lined with a heat and radiation shield made of bronze and water to prevent radiation damage from incoming protons [2]. The PS and TS are mechanically connected by bellows that are welded between the two to form a complete vacuum volume. The bellows are flexible to allow for any movement that may occur due to cool-down and magnetic

forces between solenoids.

The PS has been designed to generate a steep gradient in the magnetic field to capture and steer negative pions from the production target down the TS. At the most upstream end of the PS, the magnetic field strength is 4.6 T. The magnetic field gets progressively weaker moving downstream; by the time the PS meets the TS, the magnetic field strength has decreased to 2.5 T. One requirement on the magnetic field in the PS is that the axial field monotonically decreases with no more than $\pm 5\%$ non-linearity over the 2.8m length of the solenoid. This avoids any local extrema that may form in the magnetic field and trap particles.

3.2.1.1 Tungsten Production Target

The production target is housed within the PS, protons incident onto the production target generate the pions that will create the secondary muon beamline. The production target is made of tungsten, a material which has a large enough pion production cross section to produce the required number of stopped muons for Mu2e. The target is a 160 mm long rod and 6.3 mm in diameter, shaped like a pencil [2]. The production target is suspended in the center of the PS inner bore by a wheel-shaped support structure. Conic flares at the ends of the cylindrical target serve as attachment points for the support structure.

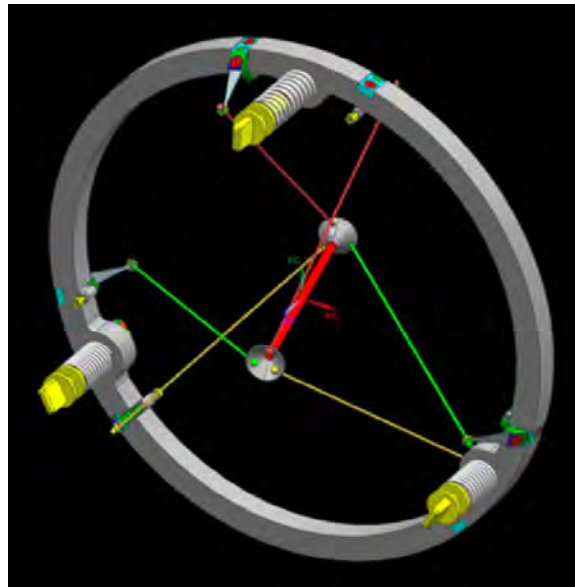


Figure 3.6: Production target simulation. A simulated 3D view of the Mu2e production target and the support structure [9].

3.2.2 Transport Solenoid

The TS is a long, curved magnet that directs muons to the stopping target. The distinct shape of the magnet along with a series of absorbers and collimators selects only low-energy, negative muons to arrive at the stopping target. Any high-energy, neutral, or positively charged particles traveling down the TS will hit the wall of the inner bore or interact with the absorbers. In total, there are 52 short solenoid coils in the TS, giving it a full length of 13.2 m along the transport center [7, 8]. A diagram of the TS is shown in Figure 3.7. The TS is split into five sections: three

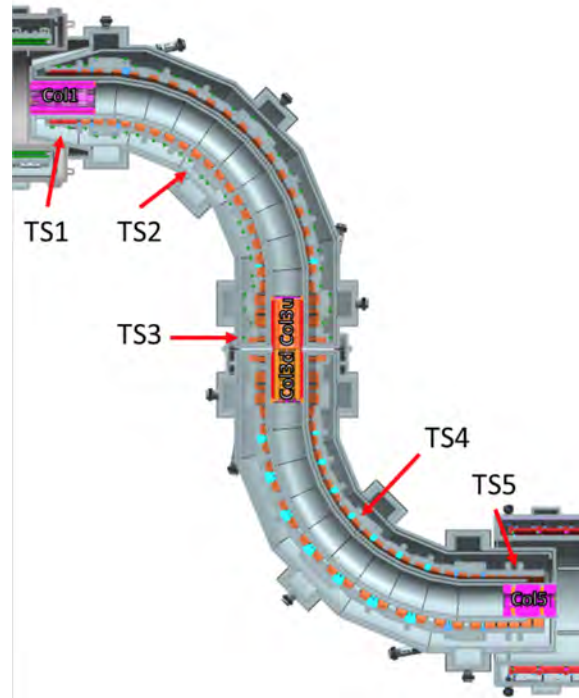


Figure 3.7: Transport Solenoid diagram. A detailed diagram of the TS showing the five different sections of the TS, starting with TS1 at the top of the diagram and TS5 at the bottom of the diagram. The collimators are shown as pink and orange boxes [2].

straight sections (TS1, TS3, and TS5) and two 90° toroidal sections (TS2 and TS4). In the straight sections of the TS, the axial magnetic field gradients are constant, preventing magnetic traps. In the toroidal sections of the TS, the magnetic field varies as $1/r$ [2]. Spiraling particles drift vertically in the toroidal sections depending on the sign their charge and the magnitude of their momentum and pitch. Particles with the wrong sign or momentum are displaced in these regions and drift into the walls, absorbers, or central collimators. The pair of collimators in the middle of the TS are offset to allow the passage of negatively charged particles and block positively charged particles. An antiproton absorber window is also used in the middle of the TS to block slow-moving antiprotons from reaching the DS.

The magnetic field across the entire TS has a small gradient that aids in particle transport;

the upstream end begins at 2.5 T and the downstream end ends at 2 T. The direction of the magnetic field gradient accelerates negatively charged μ^- towards the DS and separates positively charged μ^+ with forces in the opposite direction. The 'S' shape of the TS ensures that there is no line-of-sight between the production target and the muon stopping target. Neutral particles like photons and neutrons that are produced at the production target cannot enter the DS around this shape.

3.2.3 Detector Solenoid

The DS houses the muon stopping target, the tracker, and the calorimeter. The muon stopping target is located at the upstream end of the DS, and the tracker and calorimeter are located further downstream. The Muon Beam Stop (MBS) and STM detectors sit at the very far downstream end of the DS. In total, the DS is 11.8 m long and is made from 11 coil segments [8].

The magnetic field in the DS can be divided into two regions: the 'gradient section', which is about 4 m long, and the 'spectrometer section', which is about 8 m long. The magnetic field in the gradient section begins at 2 T and decreases linearly to 1 T at the entry of the spectrometer section. The spectrometer section has a uniform 1 T field, which is achieved by using long, single layer solenoid coils for this section instead of short, double layer segments [8]. The muon stopping target is located in the gradient section; the magnetic field gradient ensures that all CEs are directed towards the tracker and calorimeter. In the gradient section, CEs that are ejected from the stopping target in the upstream direction are reflected back downstream. The spectrometer section contains the straw tube tracker and crystal calorimeter. The uniform magnetic field in the spectrometer section creates (periodic, stable) particle trajectories and allows for accurate momentum measurements.

3.2.3.1 Muon Stopping Target

The muon stopping target is located at the upstream end of the DS. Muons are directed onto this target to be captured, from which they may convert to electrons. The stopping target design has been optimized to maximize the number of muons captured, while also minimizing the amount of material that particles must cross as they are ejected from the target. Interactions in the stopping target could cause conversion electrons to lose energy, so material choice and design should minimize this effect. The stopping target is made of 37 thin, annular aluminum foils. The foils are 100 μm thick with an outer radius of 75 mm and an inner radius of 21 mm. The entire stopping target is 800 mm in length; the foils are spaced by 22.2 mm to maximize muon capture [6]. Since muons move in helical trajectories in a magnetic field, the hole in the center of the foils is not detrimental muon capture on the target. Muons that pass through the center of an upstream foil will stop in a downstream foil. In addition, the annular design reduces unwanted interactions from the beam electron flash. Figure 3.8 shows a simulated view of the stopping target along with an actual

picture of the target.

The stopping target material was chosen to optimize CE signals in the detector region. Muonic atom lifetimes generally get longer as the atomic number of the nucleus decreases. Aluminum has an atomic number $Z = 27$ and muonic aluminum has a lifetime of 864 ns, longer than the lifetime of many of the background modes from the primary beam pulse. At the same time, sensitivity of μ^- to e^- conversion experiments benefits from using elements with low atomic numbers as the stopping target material. The expected conversion rate increases with atomic number, reaching a maximum at selenium and antimony [10]. The long lifetime of muons captured on aluminum balanced with its relatively low atomic number makes this element a good balance between optimizing for temporal background separation and high sensitivity.

The stopping target is also surrounded by two absorbers: the Inner Proton Absorber (IPA) and the Outer Proton Absorber (OPA). These two absorbers protect detector elements from damage by absorbing neutrons and protons that can be ejected by muons as they stop. If unmitigated, the neutrons and protons could damage detector elements and increase the amount of dead-time from the Cosmic Ray Veto, which will be discussed in the next chapter. The IPA and the OPA are made of polyethylene and have a conic shape with a hollow center [6].




Figure 3.8: Muon stopping target. Figure a) shows a simulated 3D view of the muon stopping target and its support structure. Each disk in the stopping target is supported by thin wires that are tensioned by bronze weights, which are visible on the lower left-hand side of the image [10]. Figure b) is a photograph of the actual aluminum stopping target and support structure. As in the simulation, the disks are supported by thin wires, though they are hard to see in the photograph [11].

3.2.3.2 *Muon Beam Stop*

The Muon Beam Stop (MBS) is a detector at the far downstream end of the DS that absorbs muons and any other beam particles which are not stopped in the stopping target. In addition, backgrounds from any particles that are absorbed in the beam stop, called 'backsplash', must be minimized [2]. Most of the particles that will be absorbed by the MBS will be muons from the beamline that were not stopped in the target, thus a material with a high atomic number where muons have a short lifetime is an optimal design choice. At the same time, neutrons may be ejected as a result of muon capture, so a material with low atomic number should also be incorporated to reduce neutron flux to the tracker and calorimeter. To meet these two goals, the MBS consists of both stainless steel and high-density polyethylene [2].

3.2.3.3 *Stopping Target Monitor*

The Stopping Target Monitor (STM) is also located at the far downstream end of the DS. This detector serves as the primary system for counting the number of stopped muons on the stopping target. Since the total number of all muon captures is necessary in the calculation of the $R_{\mu e}$ rate of $\mu \rightarrow e$ conversion, the STM is the primary means of normalization for the experiment. Part of the detector requirements is that determining the number of ordinary muon captures should not have significant effect on the uncertainty on the measured $R_{\mu e}$ conversion rate [12]. The goal for the STM is to measure the number of stopped muons to an accuracy of 10% over the course of the experiment. To accomplish this, the STM is aimed at the stopping target from the end of a long offshoot tunnel with collimators to protect the STM electronics.

To count the number of muons stopped on target, the STM measures photons that are emitted when captured muons transition between orbital states or when muonic nuclei decay after having captured a muon. The spectrum of x-rays and gamma rays emitted from these processes is dependent on the choice of target material, but is well known for a given element [12]. The number of captured muons can be determined by analyzing this spectrum; the rate of secondary photon emission from the target is correlated to the muon capture rate, with exact rates published in literature or measured by experiments like AICap [13].

The STM employs a standard method to measure x-rays from muon transitions: high purity Germanium solid-state detection high purity Germanium (HPGE). However, HPGE detectors are not well-suited for the intense Mu2e environment. Germanium is susceptible to radiation damage and commercial off-the-shelf HPGEs and their associated electronics can only handle event rates $\mathcal{O}(\text{MeV}/\text{sec})$ [6]. To mitigate radiation damage and manage event rates, the STM is heavily shielded and is placed far from the stopping target. The STM is 35 m from the stopping target and views the target through a long pipe. The pipe is connected to the downstream end of the DS and a

series of collimators allows the STM to have a full view of the stopping target while blocking the view of surrounding materials [2].

3.3 Straw Tube Tracker

The main purpose of the straw tube tracker is to accurately determine the momentum of particles coming from the stopping target. As particles move through the DS, they are moving in helical trajectories. Under the constant 1 T magnetic field in the tracker region, the momentum of a particle is directly related to the curvature of its trajectory. Thus, by tracking the path of a particle through the tracker, the particle's momentum may be precisely reconstructed. The Mu2e tracker has been designed to maximize momentum resolution, minimize potential contribution to backgrounds, and survive the intense radiation environment in the DS.

The Mu2e Experiment tracks particles as they are ejected from the stopping target using gas-filled tubes, called straws, that contain high voltage wires. As charged particles travel close to these sensitive wires through the gas within the straws, the gas becomes ionized and the ions in the gas drift toward the wires. The high voltage wires transmit the waveform to electronics at the end of the straw, where timing and position information can be extracted from the signal. A detector containing many layers of thin straws can track long particle trajectories at high rates with minimal energy loss, so Mu2e is constructing the tracker in modular sections within a region of constant magnetic field to exploit many properties of CEs. Additionally, an ideal particle tracking detector is massless to completely avoid effects from multiple scattering and energy loss. The Mu2e collaboration has taken great care to design a tracker with high momentum resolution that also fulfills the other experimental requirements.

The Mu2e tracker is just over 3 m in length and consists of 36 layers of straws. The most basic component, the straw drift tubes, consist of a thin straw wall, an inner sense wire, and the drift gas that fills the straw volume. The Mu2e tracker uses 15 μm thick aluminized Mylar for the straw walls. Each straw measures 5 mm in diameter and range from about 40 cm to 120 cm in length. Straws are formed by two layers of spiral wound Mylar, coated on the inner surface with aluminum overlaid with gold and on the outside with aluminum. The inner gold coating serves as the cathode layer, while the outer coating serves as additional electrostatic shielding and leak prevention [6]. The sense wires are 25 μm gold-plated tungsten wires which will operate around 1500 V. The drift gas is 80:20 Ar:CO₂; since the bore of the DS is evacuated to 10^{-4} Torr, the gas in the tracker helps the straws retain their shape.

The straws connect like a harp on both ends to 120° arcs in arrangements called panels. Front end electronics are housed within the outer ring of each panel. A photograph of a panel is shown in Figure 3.10. Each panel contains 96 straws, with 2 staggered layers of 48 straws. Three



Figure 3.9: Straw tube details. A photograph of a Mu2e tracker straw next to a standard sized pencil. Each straw in the tracker is thinner than a pencil. Aluminum coating in a wound pattern is visible on the outside of the straw [14].



Figure 3.10: Photograph of a Mu2e tracker panel. Each panel contains 2 layers of 48 straws, staggered to reduce gaps. The outer ring of the panel in this photograph does not contain any electronics [15].

panels connect together to form a full layer of the tracker. Next, two layers of three connected panels are rotated by 30° relative to each other to form a plane. There are 36 planes in the Mu2e tracker. Two planes are further connected to form a station. Neighboring stations in the tracker system are rotated by 60° relative to each other to ensure full coverage of incoming particles. Along the beam axis, the tracker provides full coverage for particles traveling in a path of radii r , $380 \text{ mm} < r < 700 \text{ mm}$. The inner portion of the tracker is completely uninstrumented, reducing the background from low-momentum electrons produced by muon DIO. This also lets the tracker avoid beam particles and associated backgrounds like pion and muon DIF, beam electrons, and other particles generated by muon capture on target. Examples on the radii of different types of tracks are shown in Figure 3.11.

The straws are read out by electronics at both ends. As mentioned above, tracker electronics are housed inside of the outer ring of each panel. The momentum resolution requirement on the tracker is $\sigma < 180 \text{ keV}/c$, which is fulfilled through both the geometric and electronic design [2]. To determine the relative position of passing electrons, the tracker uses 'time division,' where the




Figure 3.11: Straw tube tracker overview. The upper right shows 6 panels and how they are arranged to form a station. In the bottom portion of the figure, the full tracker assembly is shown, containing 18 stations. In the upper left, the tracker is shown from the point of view of the beam. The three circles represent projections of the helical trajectories of Michel electrons in black, intermediate momentum electrons in blue, and CEs in green [6, 14].

signal pulse timing is measured at each end of the straw to determine the position along the wire. Each straw is equipped with readout electronics to record signal times and use them to reconstruct the position of passing particles over time. Each side of a straw is connected to a preamplifier board which amplifies the raw straw signals. Then, the signals are digitized with custom digitizer boards, which can each read 16 straws, before transferring tracker data to the Mu2e data acquisition system.

3.4 Crystal Calorimeter

After particles pass through the tracker, the Mu2e electromagnetic calorimeter precisely measures their energy and final positions in the DS downstream from the tracker. The calorimeter consists of two annular disks separated by a distance of 70 cm. The separation distance between

the two disks of the calorimeter was chosen to optimize acceptance for CEs. This length is half of the distance of one period of a 105 MeV electron trajectory, so a CE that passes through the center of the first disk will interact in the second disk. Much like the tracker, the calorimeter is uninstrumented in the central gap to avoid background from the muon beam and low-momentum electrons. Figure 3.12 shows a simulation of both calorimeter disks, with labeling for the inner and outer disk radii and disk spacing.





Figure 3.12: CsI crystal calorimeter overview. A simulation of the Mu2e calorimeter. The crystals are depicted in a dark teal blue, electronics enclosures are depicted in gray and green, and cooling pipes are shown in blue and red. The inner and outer diameters of the disks are labeled, as well as the separation distance between the two disks [1].

The building blocks of the calorimeter are undoped cesium iodide (CsI) crystals that measure $34 \times 34 \times 200 \text{ mm}^3$. Each disk contains 674 undoped CsI crystals, stacked with the square ends on the face of the annular disks and the length parallel to the beamline, as shown in Figure 3.13. When charged particles pass through the CsI crystals, the scintillation light that is produced is proportional to the incident particle's energy. Each crystal is read out by two custom array large area UV-extended silicon photomultiplier (SiPM)s. The dimensions of the readout SiPMs are 2×3 arrays of $6 \times 6 \text{ mm}^2$ cells [1]. Each SiPM is connected to a front-end electronics board to amplify

and shape signals. Then, a custom digitizer module groups signals from the SiPMs and transfers them to the Mu2e data acquisition system. These electronics reside in the enclosures around the outer circumference of the calorimeter.



(a) Cesium iodide crystal



(b) Crystal stacking in the calorimeter

Figure 3.13: Calorimeter crystal details. Figure (a) shows a photograph of a CsI crystal sitting in protective paper wrapping [16]. Figure (b) is a photograph of the start of the first stacking test of the CsI crystals. The crystals are being stacked into the calorimeter frame. Crystals are wrapped with reflective cladding and protective outer layer [17].

In addition to measuring particle energies and positions, the Mu2e calorimeter also has a few other important functions. The calorimeter serves as an independent confirmation for CE signal events by functioning as a trigger and measuring the kinetic energy of incident particles independently from the tracker. The calorimeter also improves the quality of track reconstruction through the tracker by using the final position of a particle in the calorimeter as the "track seed". This additional information greatly improves the efficiency of reconstruction when used for consistency checking between the final calorimeter position and trajectories through the tracker. Finally, the calorimeter is a powerful tool in particle identification. Energy deposition in the calorimeter is

used to separate between electrons and muons of the same momentum using the ratio E/p . The calorimeter has been designed to have energy resolution of $\sigma_E/E < 10\%$, timing resolution of $\sigma_t < 500$ ps, and a position resolution of < 1 cm. The e/μ separation endows the calorimeter with a muon rejection factor around 200 [1].

References

- [1] S. Giovannella, *J. Instrum.* **15** (06), C06022.
- [2] L. Bartoszek and *et al.* (Mu2e Collaboration), *Mu2e Technical Design Report*, Tech. Rep. (Fermilab, 2015) arXiv:1501.05241 [ins-det] .
- [3] S. Werkema, “CD-3c Review: Mu2e Accelerator Systems,” (DOE CD-3c Review 2016-06-14), mu2e-doc-7530-v6.
- [4] K. Lynch, in *Proceedings of The 19th International Workshop on Neutrinos from Accelerators NUFAC2017 — PoS(NuFact2017)*, Vol. 295 (2018) p. 082.
- [5] R. Ray and J. Whitmore, “CD-3c Review: Mu2e Project Overview,” (DOE CD-3c Review 2016-06-14), mu2e-doc-7084-v23.
- [6] R. H. Bernstein, *Front. Phys.* **7**, 1 (2019).
- [7] M. Lamm, “CD-3c Review: Mu2e Solenoids,” (DOE CD-3c Review 2016-06-14), mu2e-doc-7461-v9.
- [8] M. J. Lamm, *et al.*, *IEEE Transactions on Applied Superconductivity* **22**, 4100304 (2012).
- [9] M. Campbell, “Mu2e Production Target and EOAT,” (Mu2e Target Station Meeting 2016-06-20), mu2e-doc-7660-v1.
- [10] J. Miller and S. Krave, “Requirements and Specifications: WBS 5.05 Muon Stopping Target,” (Stopping Target and DS Internal Shielding CRR 2019-06-05), mu2e-doc-1437-v8.
- [11] J. Miller, “L3 WBS 5.05 Stopping Target,” (Director’s BCP Review 2022-05-10), mu2e-doc-42066-v2.
- [12] J. Miller, “L3 WBS 5.06 Stopping Target Monitor,” (Director’s BCP Review 2022-05-10), mu2e-doc-42069-v2.

- [13] A. Edmonds, “Latest Updates from the AICap Experiment,” (2018), arXiv:1809.10122 [ins.det].
- [14] T. Miyashita, “The Mu2e Experiment,” (51st Annual Fermilab User’s Meeting 2018-06-20), <https://indico.fnal.gov/event/16332/contributions/36956/attachments/22954/28458/Mu2e.pdf>.
- [15] Y. Wu, “The Mu2e Experiment,” (DPF 2021 2021-07-13), <https://indico.cern.ch/event/1034469/contributions/4427349/>.
- [16] D. Lin, L. Borrel, and J. Ojang, “Optical Properties of Crystal Measurement,” (Mu2e EMC Meeting 2020-07-16), mu2e-doc-34413-v1.
- [17] S. Giovannella, “Organization for Crystal Stacking at SiDet,” (Mu2e EMC Meeting 2022-06-23), mu2e-doc-42577-v2.
- [18] R. Blaustein, “Digging begins for muon g-2 and mu2e beamlines,” (2014), https://www.fnal.gov/pub/today/archive/archive_2014/today14-12-09.html.

Chapter 4 Cosmic Ray Veto

The rest of this dissertation will primarily document work performed on the Mu2e Cosmic Ray Veto (CRV). This chapter will describe the CRV detector in detail, from its basic components, to the overall organization and design, electronics, and outer shielding layers. Later, Chapters 6 and 7 will discuss scintillator response over time and work that has been performed to understand the aging of the Mu2e CRV.

The Mu2e CRV serves as an active shield to detect and reject cosmic rays, the experiment's largest source of background. The bulk of the CRV is comprised of scintillator, a plastic material that emits light proportional to the energy deposited by charged particles as they pass through. The CRV will surround the entire DS and part of the TS to protect against cosmic rays from any incoming angle. To achieve the desired experimental sensitivity, the Mu2e CRV must be 99.99% efficient at detecting and rejecting cosmic rays. Such a high level of efficiency is achieved through using multiple layers of scintillator and looking for coincidences between layers to certify a cosmic signal. The Mu2e CRV contains four layers of long, scintillating counters. When a signal is detected in adjacent CRV counters that span at least three out of four layers at the same time, the signal is identified as an incoming cosmic ray. Due to the unpredictable trajectories and energies of cosmic rays, the best way to mitigate against this background is to employ a dead-time window where all data is rejected until the cosmic ray has exited the DS. High efficiency in the CRV is required to minimize the dead-time necessary to mitigate against cosmic ray backgrounds. An experimental requirement is that the CRV produces less than 10% dead-time.

In addition to requiring the CRV to operate at a high efficiency, the detector also needs to be compact enough to surround the DS within the confines of the Mu2e hall. The CRV is a modular detector that optimizes spatial constraints by using multiple, overlapping layers of scintillator to detect charged particles passing through. The smallest modular pieces of the CRV are called counters, which are long bars of scintillator outfitted with read out fibers. These bars get stacked in four staggered layers into larger structures, called modules, which will surround the outside of the DS to shield against cosmic rays. We will begin the discussion on the CRV by first describing the smallest parts, the components of a counter, then describe the overall CRV structure, and finish by discussing the read out electronics and data acquisition.

4.1 CRV Counters and Dicounters

The basic component of the CRV is what is referred to as a counter. Each counter has cross-sectional dimensions of $2 \times 5 \text{ cm}^2$ and measures in length between 1 m and 7 m. Different counter lengths are necessary to surround the different areas along the surfaces of the TS and DS. In total, the CRV contains 5504 counters, organized into four layers. A CRV counter has three parts: the scintillator bulk, inner readout fibers, and an outer reflective cladding. A cross-sectional view of a counter is shown in Figure 4.1.




Figure 4.1: Photograph of a CRV counter cross-section. The clear plastic scintillator is in the center of the counter, filling the bulk, and is surrounded by a thin, white layer of titanium dioxide cladding. Inside the bulk, the two green dots are the ends of the wavelength shifting fibers.

All processes associated with counter fabrication are performed in-house at Fermilab or University of Virginia. The scintillator bars are produced at Fermilab. The bars are then shipped to University of Virginia where they are assembled into counters. Let's begin the discussion of CRV counters by describing the three individual components that make a counter and finish the discussion with how the counters are organized and read out.

4.1.1 Scintillator Bars

The scintillator bars for Mu2e are produced in-house at Fermilab's state-of-the-art production system for extruded scintillator. The facility at FNAL is a collaboration between FNAL and Northern

Illinois University's Northern Illinois Center for Accelerator and Detector Development (NICADD), and thus the scintillator produced at this facility is called FNAL-NICADD scintillator. The plastic is a general purpose polystyrene (DOW Styron 665W) doped with 1% 2,5-diphenyloxazole (PPO) and 0.03% 1,4-bis(5-phenyloxazol-2-yl) benzene (POPOP) fluorescent additives [1]. A great deal of research and development went into choosing the best polystyrene base to use for the scintillator, selecting and optimizing concentration of the dopants, and maintaining quality of raw material sources over time. This specific formulation of polystyrene and dopants was developed for use in MINOS, MINERvA, and ALICE collaborations and has proven to be a reliable scintillator with high light yield [2, 3]. Upon delivery, the polystyrene and dopants arrive in large boxes of pellets. The FNAL-NICADD facility uses a continuous in-line compounding and extrusion process for the scintillator, which allows for minimal handling of the polystyrene and dopant pellets and computerized control over speed through the extrusion line [1]. The two fiber channels in the center of the scintillator bars are co-extruded by injecting dry nitrogen into the semi-molten plastic during the extrusion process [4].




Figure 4.2: Photograph of a CRV scintillator extrusion. The clear plastic scintillator is in the center of the counter, filling the bulk, and is surrounded by a thin, white layer of titanium dioxide cladding. Inside the bulk, the two holes that are co-extruded for the inner fibers are visible [5].

Scintillator production began on January 16, 2018 and concluded on May 16, 2018. A total of 31.3 km of scintillator was extruded in this span. After the scintillator was produced at FNAL, it was sent in batches to University of Virginia to be assembled into counters, dicounters, and then modules. Various quality control techniques were employed at University of Virginia to evaluate the scintillator, including checking scintillator samples for light yield consistency, size, and shape. During quality control testing, it was discovered during the manufacturing run that the use of a new polystyrene batch resulted in a 13% decrease in light yield. The reason for the decrease in light yield was never determined [4]. As a result, extrusions with higher light yield are used in critical areas of the CRV.

4.1.2 Titanium Dioxide Cladding

To increase the light yield that can be detected at the ends of the bars, a reflective cladding layer is added to the scintillator surface to improve reflection of light from the surface of the bars. This cladding is added during the scintillator extrusion process where each bar is co-extruded with a 0.25 mm thick 30% titanium dioxide (TiO₂) and polystyrene cladding layer. The co-extrusion process ensures that the interface between the cladding and scintillator is whetted and that no air gaps are present that could impact light yield. The cladding is visible above in Figure 4.2, where the white outer coating of cladding is visible around the perimeter of the clear scintillator bulk. In addition to its reflective properties, the outer cladding layer aids in the assembly of counters into dicounters and modules because the durability of the coating enables direct gluing of counters to each other and to the absorber layers that will be in between layers of scintillator [5].

Similar to the polystyrene used in the scintillator bar extrusions, the titanium dioxide arrives in large shipments of pellets that is integrated into the automated extrusion process. Each shipment of titanium oxide pellets arrives with small, manufacturer-extruded samples of the cladding material, called coupons. These coupons can be used for quality control of the cladding and a later chapter of this dissertation will discuss studying titanium dioxide coupons for signs of aging.

4.1.3 Wavelength Shifting Fibers

To transmit light from incoming cosmic rays to electronics that will be mounted to the end of the counters, fibers are necessary inside the scintillator bulk. The counters are long, but the attenuation length of light inside the scintillator bulk is short, so fibers absorb the light from the scintillator and transmit the light signals down to the ends of the counters.

The specific fiber that will be used in the majority of the Mu2e CRV is 1.4 mm Kuraray Y11 wavelength shifting fibers. Studies were performed using a fiber scanner to optimize the diameter of the fibers and study the light properties and quality of the fibers that were chosen [6]. The results of light yield for fibers of different diameters are shown in Figure 4.4 As mentioned above, each scintillator bar has two co-extruded holes in the bulk. These holes are where the wavelength shifting fibers are inserted. There are no adhesives, fillers, or epoxies used to fix the fiber into the hole to prevent any loss of light. Kuraray Y11 fibers contain a fluorescent dye, K27, that absorbs blue light from the scintillating counters between 375 and 475 nm and re-emits the light in the green spectral region between 450 and 600 nm [7]. The green emission spectrum of the Kuraray Y11 fibers is well matched to the spectral photon detection efficiency response of the readout SiPMs, which will be discussed later. The fibers that were chosen for the CRV are double-clad and non-S type to enhance the light yield and attenuation length properties that the fiber provides [6].

The most important area to monitor with the CRV is the surface on the top of the DS since




Figure 4.3: Emission and absorption spectra for Kuraray Y11 wavelength shifting fibers. The fibers absorb in the blue and re-emit light in the green [5].

cosmic muons generally have trajectories traveling downward from the sky to the surface of the Earth. Simulations have also been performed to determine the expected number of cosmic ray muons that produce CE-like background events and where each cosmic crosses the CRV. A vast majority of these simulated events only pass through the modules on the top of the DS, called CRV-T or CRV-Top modules, without passing through any other side of the CRV [5]. Additionally, one known characteristic of using detectors made from plastic scintillator is that the scintillator becomes less efficient over time. In general, the light yield as a result of cosmic ray energy deposition decreases over time, resulting in lower efficiency. As a result of cosmic background simulations and knowledge of detector aging, the counters that will sit on top of the DS were the last counters to be produced. This provided the collaboration with more time to study the light yield characteristics of the counters and decide to increase the fiber diameter used in the CRV-T modules from 1.4 mm to 1.8 mm. The larger diameter increases the initial light yield of the CRV channels by 20% where it is in use, which will hopefully prevent the CRV-T modules from losing efficiency in critical areas [8]. As of 2023, all counters have been produced and CRV-T modules include the larger diameter fiber. A discussion about sectors and the specific organization of the other sides of the CRV will follow in the next section.

4.1.4 Dicounters

At this point in the fabrication process, the scintillator extrusions, or counters, are then organized into pairs, called dicounters. Pairs of counters are placed side-by-side to form flat dicounters with a 10×2 cm cross-section. A dicounter is made by scoring and cleaning a long




Figure 4.4: Light yield vs fiber diameter. Light yield measurements versus attenuation length in meters for Kuraray Y11 wavelength shifting fibers with 1.0, 1.4, and 1.8 mm fibers. These measurements were made using a custom-made fiber scanner setup [6].

edge of two counters, then bonding the two scored edges together using 3M DP420 epoxy [5]. The ends of each counter are made flush with one another and a custom-made fixture is used to apply pressure to four counters at a time to ensure a durable bond as the epoxy cures. While the pairs are bonding, the wavelength shifting fibers can be inserted into the channels in the counters.

On each end of a dicounter, a fiber guide bar (FGB) is mounted onto the end to hold the wavelength shifting fibers into a fixed position, form a mechanical interface for readout electronics, and prevent light leaks. FGBs are attached to the end of each dicounter using 3M DP100 epoxy and #4-20 thread-forming screws. The use of these screws has no measurable effect on light yield [4]. The fibers that are embedded within the scintillator are fixed to the FGBs using fast-curing 3M DP100 epoxy within four counterbored holes. The position of the fibers in the scintillator is not well constrained, so excess fiber length and chamfers within each counterbore funnel the fiber into the fiber hole to ease FGB installation. After the epoxy has cured, excess fiber is cut off using a nichrome-wire hot knife to prevent cracking and damage to the fiber [4]. Then, the face of each fiber is polished using a repurposed flycutter from the NOvA experiment. Quality of the flycutting is checked through visual inspection, photographing the fiber face surface, and analyzing the fiber face using a surface roughness tester.

With the fiber ends polished, readout electronics boards may be mounted to each FGB. This board is called the Counter Motherboard (CMB) and holds the readout SiPMs, calibration LEDs, and temperature sensor. The readout electronics will be described in detail in the next section. The entire assembly at the end of a dicounter, including the FGB and the CMB, is known as a manifold. Figure 4.5 shows two labeled diagrams of the manifold at the end of a dicounter. Finally, after

the manifold has been attached, dicounter surfaces along the edges are covered with light-tight aluminum tape. After fabrication, the aluminum tape is protected using a low-friction, Teflon coated, fiberglass tape [4].




Figure 4.5: Fiber guide bar and counter motherboard details. Top: Exploded diagram of the mechanical design of a dicounter end and the readout electronics, known as a manifold. Bottom: X-ray view of the manifold assembly. Two types of manifolds are used: those with 90° HDMI headers, as shown in the figures, and those with straight HDMI headers (not shown).

4.2 CRV Modules

With dicounters in hand, CRV modules may finally be produced. The full CRV contains a total of 83 modules in seven different sizes, each containing four layers of dicounters and three layers of absorber. Within a module, dicounter layers are staggered to prevent the formation of

any gaps that cosmic particles may pass through undetected. Most modules have layers offset by 42 mm to minimize the effects of projective gaps. Figures 4.6 and 4.7 show how dicounters are stacked within a module. Figure 4.6 shows how dicounters are staggered to prevent gaps that allow cosmic particles to pass through the CRV undetected. Figure 4.7 shows the cross-section of an entire module containing four layers of eight dicounters with many different dimensions labeled.




Figure 4.6: CRV dicounter stacking diagram. Cross-sectional diagram illustrating how dicounters are stacked. Each layer is staggered relative to the last by about 42 mm. Internal gaps from epoxy between two counters in a single dicounter are represented, as well as gaps between neighboring dicounters to show how maximal coverage is achieved.

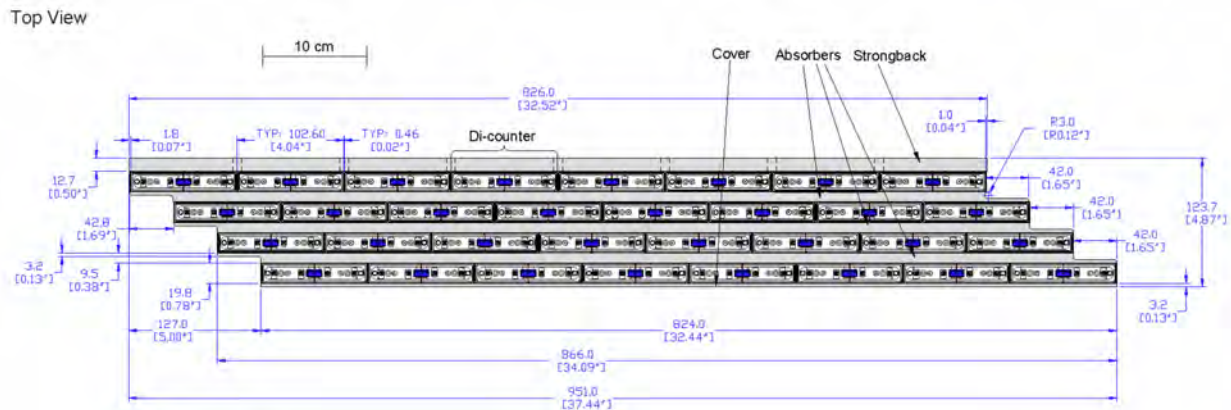


Figure 4.7: CRV module schematic. Diagram of the cross-section of the end of an entire module. The length of the module runs into the page in this diagram. All dicounters are shown. The dimensions of the module and length of the offsets between dicounters are labeled in blue. The absorber layers, the strongback, and the cover are labeled and represented by gray boxes.

To construct a module, first the four layers of dicounters are assembled. Adjacent dicounters within each layer are shimmed (typically with four 0.51 mm thick shims) to keep the overall layer width uniform and the spacing between the dicounters consistent [4]. The outer perimeter of each

layer is made light-tight using black RTV. Then, the absorber layers are bonded to the dicounter layers using Devcon HP-250 epoxy. The aluminum absorber layers are 9.925 mm thick, providing both mechanical rigidity and reducing the number of electrons that are produced from gamma interactions from traversing multiple dicounter layers, which would create spurious coincidences. Finally, the supporting strongback and outer cover may be attached to the module. The strongback is 12.7 mm thick and will interface with the CRV support structure. The outer aluminum cover is 3.175 mm thick.

Nominally, all modules contain 64 scintillating counters arranged in four layers of eight dicounters and three layers of aluminium absorber. There are also few narrow modules that only contain 32 counters, otherwise all modules are the same width, but different lengths. The length of a module depends on its placement surrounding the TS or DS.

4.2.1 Sectors

Because the CRV needs to cover the entire DS and half of the TS, CRV modules vary in length between 1m and 7m long. The CRV is divided into different areas, called sectors, corresponding to which side of the experiment the particular modules will cover (see Figure 4.8). Beginning at the upstream end of the experiment, we can list the different sectors of the CRV in order. At the most upstream end of the experiment, the PS is not covered by the CRV. Moving down the TS, the TS is about halfway covered by the CRV (see Figure 4.8). The upstream sector, CRV-Upstream (CRV-U), runs perpendicular to the DS and forms the wall that is beside the section of the TS which is covered by the CRV. The top of this section of the TS is covered by the CRV-Transport Solenoid (CRV-TS) and CRV-Transport Solenoid-Extension (CRV-TS-Ext) sectors. Moving now to the DS, the left and right sides of the DS are flanked by the left and right sectors of the CRV, CRV-Left (CRV-L) and CRV-Right (CRV-R) respectively. The top of the DS is covered by the top sector, CRV-Top (CRV-T). The right side of the DS includes the cryogenic feedthrough, which is surrounded by the CRV-Cryostat (CRV-Cryo) sector. Finally, at the downstream end of the DS, the final downstream wall is covered by the downstream sector, CRV-Downstream (CRV-D). There is not a bottom sector or any modules instrumented beneath the DS because cosmic rays only travel downwards to the Earth's surface from the sky.

In total, the CRV contains 83 modules of 13 different types. Most modules contain the nominal 32 dicounters, arranged in four staggered layers with dual-ended readout. Some areas of the CRV have special conditions that allows for different module designs. In the CRV-U and CRV-D sectors of the CRV, an offset between layers is not necessary as the projective gaps point to the horizon where the muon flux is effectively zero [4]. Some modules are thin and only contain 16 dicounters to form the end of a sector. Modules in the CRV-U, CRV-TS, and CRV-TS-Ext sectors feature single-ended readout, with reflectors or absorbers on the opposite end. This is due to the




Figure 4.8: CRV sector diagram. Rendering of the full CRV detector, with sectors labeled. The PS and part of the TS are also shown in the top image. The DS is completely covered by the CRV, shown in gray, and shielding, shown in red. For scale, a cartoon of a person is also shown.

large radiation doses near the PS and the inability to access manifolds after installation. Reflectors on the end of long modules are used to improve light collection. Conversely, some short modules in the CRV-D and CRV-Cryo sectors have black absorbers on the far end, as these modules do not benefit from additional light collection.

The module support structures are attached to shielding blocks, except for a set of guide rails for the side modules that are affixed to the floor. The module support structure has been designed to accommodate changes in attachment positions in the event the shielding blocks are removed and replaced. The positions of the shielding blocks are not well defined, so any servicing of the detectors and replacing of the blocks could move the module support structures. The shielding endcap is designed to move, so the support structures for the downstream sections of the CRV-T, CRV-L, CRV-R, and CRV-D support structures are also movable with the endcap [4]. Each module is supported by rails. There is a wheel truck on the bottom of each strongback that supports the modules, where each truck rides on the inverted V-shaped rails attached to the shielding blocks that will be placed above the TS and DS. The rails allow the modules to roll along the length of the shielding for installation. Careful ordering of the modules during installation will allow for the modules to roll tightly together. Side modules are brought together also by the use of turnbuckles at the top and bottom of the strongbacks.

The top modules in the CRV-T, CRV-TS, and CRV-TS-Ext sectors lie flat on a series of frames. Low friction plastic on the frames allow for the modules to be slid into place. The modules are held in position with the use of fasteners between the modules and frames. Like the side modules, turnbuckles also pull adjacent top modules together. The CRV-U and CRV-D modules are attached to frame assemblies. For CRV-U, the frame sits on pedestal blocks separate from the shielding. The top of the CRV-U frame assembly is connected to shielding blocks by tracks that allow horizontal movement. The base of the frame assembly rests on tracks and a leadscrew provides horizontal movement for the entire CRV-U frame assembly. Since the CRV-U is closest to the PS, the radiation dose will be high. If the dose rates are higher than expected, the entire CRV-U assembly may be moved away from the PS using the turnscrew [4]. For CRV-D, the frame assembly is mounted on brackets that attach to the movable endcap. Part of the CRV-Cryo sector uses a track which is attached to the wall of the building.

4.3 Readout Electronics and Data Acquisition

The electronics readout system for the Mu2e CRV has been designed to digitize the light signals from the scintillator bars in amplitude and time and read the data out at a high rate. Due to the high rates, the system must also be able to buffer an entire supercycle (2.56×10^5 microbunches) worth of data. Many components on the readout electronics boards are off-the-shelf parts to reduce

cost. The primary sensor to detect light from the scintillators are SiPMs mounted at the ends of each dicounter. There are several layers of readout electronics: SiPMs mounted on the Counter Motherboards (CMBs), Front End Boards (FEBs), and the Readout Controllers (ROCs). A block diagram for data flow through the entire CRV system is shown in Figure 4.9.

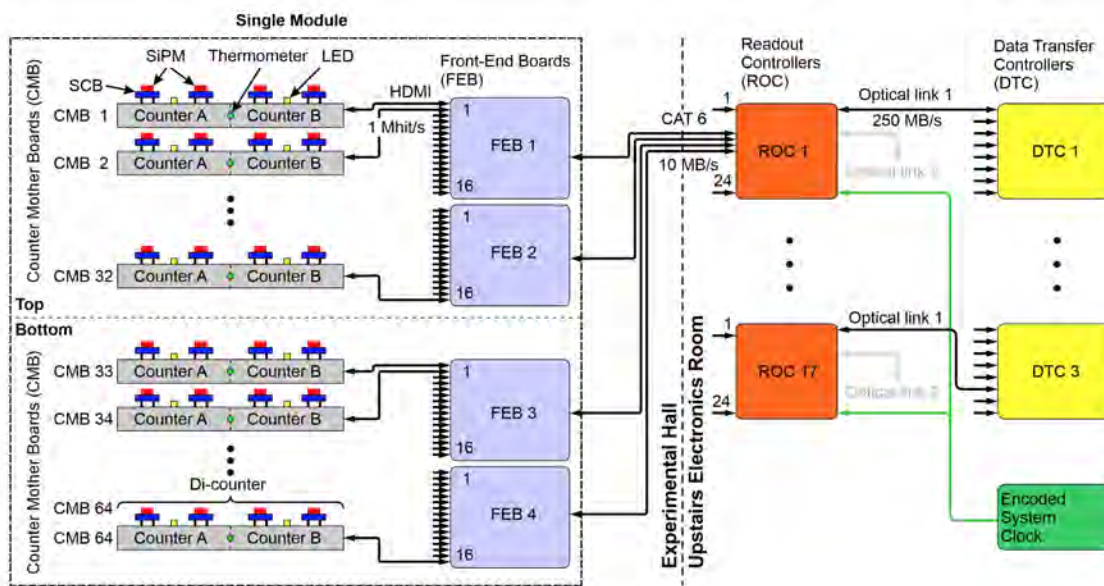


Figure 4.9: Block diagram of CRV electronics chain. This block diagram shows the flow of data through the entire Mu2e CRV system as signals come from dicounters and go through several layers of readout electronics to eventually get transferred to the Data Acquisition (DAQ) system [9]. The dashed line near the middle of the diagram represents which electronics will be housed in the experimental hall versus upstairs in the shielded electronics room. CMBs are attached directly to the ends of dicounters while the Front End Board (FEB)s are mounted in shielded enclosures on the outside of each module. The Readout Controller (ROC)s and other DAQ electronics will be housed in racks upstairs in the electronics room.

The SiPMs convert light from the inner wavelength shifting fibers into electrical signals. Each CMB has four SiPMs mounted onto SiPM carrier boards inside. The CMBs are responsible for several duties, such as providing bias to the SiPMs and shaping the output signals. The CMBs then connect to FEBs. FEBs receive and digitize the signals from the SiPMs in the CMBs. FEBs also have a number of other functions: supplying power to the CMBs and SiPMs, buffering all of the data taken in a supercycle, aggregating and delivering zero-suppressed event data, and sending status information to the ROCs [5]. Each FEB serves 16 CMBs, so two FEBs are required to read out one end of a module. The ROCs aggregate data from multiple FEBs and act as an intermediary between the FEBs and the experiment's Data Transfer Controllers (DTCs). Each ROC can serve up to 24 FEBs and can also send status, control, and timing information to the FEBs that it serves. Events from the ROC are then sent to the DTCs, the frontend hardware for DAQ. In total, 19,456 SiPMs, 4,864 CMBs, 316 FEBs, and 16 ROCs will serve the entire CRV.

The Mu2e DAQ is designed to operate with software-only triggers run on a farm of online processors, with a latency of about one second, or about one spill cycle length. The trigger signals from the DAQ are sent to the FEBs. Triggers are sent based on the satisfaction of track-finding criteria in software. When a trigger is received, the FEB sends the entire micro-pulse of data to the ROC and then the DAQ through the DTCs.

4.3.1 Counter Motherboards

The CMB is mounted directly onto the end of a dicounter and provides direct electrical interface for the photodetectors, the SiPMs. Each CMB serves four SiPM channels, one for each fiber in the dicounter. In addition to providing bias to the SiPMs and shaping the output signals as mentioned above, the CMBs also monitor the temperature of the manifolds with an on-board temperature sensor and are able to use flasher LEDs for calibration. Each CMB has a unique serial number that is logged in a database containing information about quality control of the board and its position on the CRV. Each CMB is outfitted with an HDMI connector to receive the bias voltages for the SiPMs, trigger signals for the LEDs, and power for the temperature sensor from the FEB.

The components of a CMB are: four SiPMs, passive shaping components, two LED flashers, four CMOS switches, and a temperature sensor. The two 405 nm LEDs on board are used for calibration purposes. By flashing the LEDs at one end of a dicounter, the illumination of SiPMs on the other end can provide timing information and a heartbeat test. The duration and intensity of LED flashes may be controlled by the FEB. The set of four CMOS switches forms a gate that can quickly lower the SiPM bias voltage by changing the ground reference. This feature is called the flash gate and is used for lowering the SiPM bias voltage during proton beam arrival. If a SiPMs bias voltage is below the breakdown voltage during beam arrival, it can escape the prompt radiation from the beam flash. Just prior to the opening of the signal window, the SiPM bias voltage is restored above the breakdown voltage. The temperature from the temperature sensors on board each CMB are used to adjust the SiPM bias voltage in response to temperature fluctuations. The temperature sensors are the components with a unique serial number used for cataloging the CMBs. All CMBs underwent quality control testing at University of Virginia before being inserted into manifolds. These tests included operational checks of the temperature sensor serial number readback, flasher LEDs, and flash gate, and evaluations of a connected SiPM's response to light [4].

SiPM carrier boards are what hold the SiPMs in place and flush with the dicounter end surface. The inner connection to each SiPM carrier board is made using a set of spring-loaded contacts called pogo pins that gently press the SiPMs flush against the polished surface of the wavelength shifting fibers. These pogo pins assure that there is contact between the SiPMs and fibers without damaging the SiPMs with excessive force.

4.3.1.1 SiPM Readout

SiPMs, also known as multi-pixel photon counters (MPPCs), are off-the-shelf photodetectors with high detection efficiency and a small profile, which makes them ideal candidates as the CRV's primary photodetector. The most important factors when selecting a photodetector for the CRV are the need to maintain high light yield and good time resolution over the lifetime of the experiment, as well as a stable, well understood photoelectron (PE) threshold during operation. The specific SiPM chosen for the Mu2e CRV is the Hamamatsu S13360-2050VE [10]. This device has a large photosensitive area, $2 \times 2 \text{ mm}^2$, a high gain, a large number of photosensitive pixels, and a high photon detection efficiency, satisfying the Mu2e requirements. Device specifications are listed in Table 4.1 and a photo of one of these devices is shown in Figure 4.10.




Figure 4.10: Photograph of a Hamamatsu S13360-2050VE SiPM. The active area of pixels measures $2 \times 2 \text{ mm}^2$ and includes 1584, $50 \mu\text{m}$ photosensitive pixels. The hole in the center is the through-silicon via (TSV).

These multi-pixel photon counter (MPPC) devices operate using a high-density matrix of single-photon avalanche photodiodes (SPADs) connected in parallel. Each SPAD has a high internal gain which makes it sensitive to single photons. By operating the device in Geiger mode, a reverse-bias mode at a given voltage over the breakdown voltage, a single photon incident on a pixel is enough to generate electrical signals within the photodiode. photon detection efficiency (PDE) is the fraction of photons incident on the device that are detected. PDE is dependent on wavelength, bias, temperature, and the pixel packing fraction. When a photon is incident on a pixel, the electrons that are liberated generate a measurable current within the SPAD, in discrete quanta depending on the number of electrons that are liberated. The output of a SiPM is a superposition of the current pulses generated by each detected photon, also known as photoelectrons (PEs).

Each side of a dicounter is read out by two SiPMs, one for each fiber, with dual-ended readout for dicounters where spatial constraints allow. Using two photodetectors at both ends

provides redundancy and allows for spatial resolution along the length of a dicounter within 15 cm through using the timing information of when a signal reaches both ends of a dicounter. Each SiPM is coupled directly to the end of a single wavelength-shifting fiber through the SiPM carrier board inside the CMB. Each SiPM is surface-mounted by Hamamatsu onto one side of a SiPM carrier board, while the other side of the board makes contact with the spring-loaded pins on the CMBs. Using the SiPM carrier boards and the spring-loaded pogo pins ensures reliable contact between each SiPM and the surface of the wavelength-shifting fiber. All SiPMs and SiPM carrier boards undergo quality control upon receipt. For each SiPM, the breakdown voltage, gain, dark count rate, and cross-talk probability are measured and checked with the expected device specifications. These devices have also been irradiated and tested for radiation damage to ensure that they are compatible with operating in the Mu2e environment for the lifetime of the experiment [11]. A sufficient number of spare components have been ordered that may replace rejected or damaged channels.

4.3.2 Front End Boards

The FEB reads out and digitizes the electronic signals from the SiPMs in time and charge domains. PE signals are sent from the SiPMs to the FEBs via the HDMI connection on the CMBs. FEBs are also responsible for providing bias voltage to the SiPMs and controlling the flasher LEDs in the CMBs. Each FEB serves 64 SiPM channels, thus four FEBs are needed to read out an entire module on both ends. A photograph of a prototype FEB with the main components labeled is shown in Figure 4.11.

The main components of each FEB are the four field programmable gate array (FPGA) chips mounted to each board. The FPGA chips provide event-building logic necessary to produce

Table 4.1: SiPM device specifications. Specifications for the Hamamatsu S13360-2050VE MPPC used in the CRV [10]. The SiPMs in the CRV are operated at a bias that is 2.5 V higher than the breakdown voltage at 25°C.

Package type	Surface-mount, single channel
Photosensitive area	2 mm × 2 mm
Number of pixels	1584
Pixel size	50 μm
Photon detection efficiency	40%
Spectral response range	320 to 900 nm
Peak sensitivity wavelength	450 nm
Dark count rate	< 300 kHz @ 0.5 PE threshold
Terminal capacitance	140 pF
Gain	1.7×10^6
Measurement condition	Ta = 25°C




Figure 4.11: Photograph of a prototype FEB. The main components are labeled in yellow [9]. The HDMI connectors on the bottom of the board connect to the CMBs and an Ethernet connector on the top of the board connects to the ROC. The final FEB design will be similar to this prototype, however, a few components will have different models.

human-readable data. Each FEB has 16 HDMI ports, enough to read out the CMBs of two layers of a module to take in signals from the CMBs. There are eight ultrasound chips on board, specifically Texas Instruments AFE5807 octal ultrasound analog front end chips [12]. Each ultrasound chip has eight channels of low-noise preamplification, variable gain amplification, low pass filters, and samples each channel at 80 MHz using a 12-bit analog-to-digital converter (ADC) [4]. Each of the four FPGA chips on board serve 16 channels, and each has its own LPDDR RAM buffer for storing events until they can be read out. There are two CAT6 Ethernet connections on each FEB to provide power to the board and a link with the ROC for data transfer. An on-board bias generator supplies the SiPMs with appropriate bias voltages and 12-bit digital-to-analog converter (DAC) chips allow for the voltage of each SiPM channel to be adjusted relative to the common bias. A 24-bit ADC and a network of multiplexers are used to measure the current of each SiPM with 100 pA resolution. Control and readout of each FEB is managed by a single Texas Instruments RM48 arm-cortex microcontroller on each board [13]. A block diagram illustrating the flow of data from the ultrasound chips through the FPGAs is shown in Figure 4.12.

The initial design for the FEB included four Xilinx Spartan-6 FPGA chips on board [14]. In early 2022, Xilinx announced that the Spartan-6 FPGA was discontinued, driving a redesign

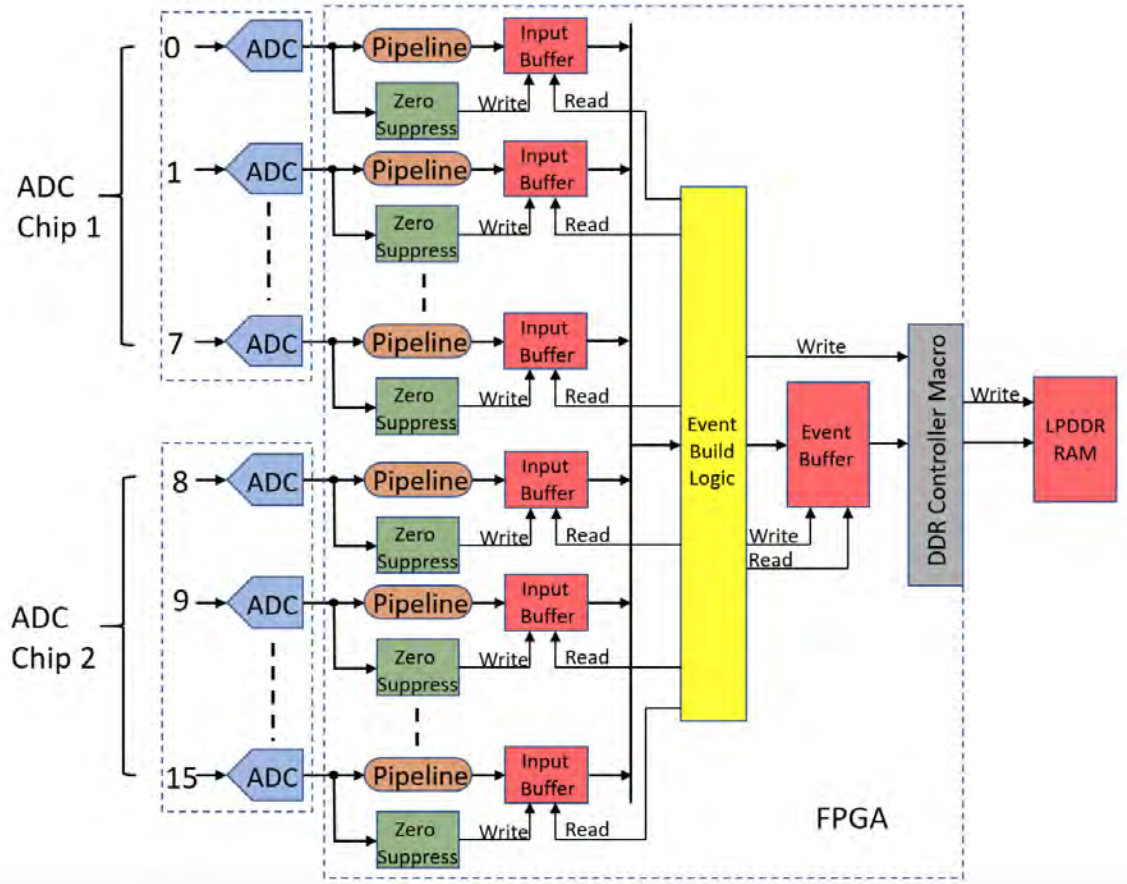


Figure 4.12: Block diagram of FEB data flow path. This block diagram shows the flow of data through the FEBs as ADC signals come in from the ultrasound chips and enter an FPGA [4].

campaign for the Mu2e FEB. The new board design has been named the FEB-II. After considering several alternate FPGAs, the component that was chosen for the FEB-II is the Xilinx Spartan-7 [15], the next generation FPGA chip from the previous Spartan-6. The part number for the 2G LPDDR RAM memory chip on board also needed to be replaced in order to be compatible with the new FPGA on board. As of May 2023, all Spartan-7 FPGA chips that were ordered for the FEB-II production run have been received.

FEBs are fabricated by a third-party vendor. Upon receipt at Kansas State University, all FEBs undergo quality control testing for functionality. Then, the boards are calibrated such that the voltage across each channel is 0V. These calibration values for each channel are saved, ensuring accurate voltage settings for each SiPM. FEBs have also been studied for radiation tolerance using a proton beam at the Northwestern Medicine proton therapy cyclotron [16]. Various tests have been performed using FEBs to read out dicounters in the Fermilab Test Beam Facility. Several dicounters have been put into the proton test beam with FEBs connected and successful event readout has been demonstrated with the first version of the board design. When the FEBs are installed, they will

be secured in aluminum enclosures on the outside of each module to mitigate the strength of the magnetic field in which they will operate.

4.3.3 Readout Controllers

Readout Controllers (ROCs) link the FEBs to the rest of the CRV DAQ system. Each ROC supplies power, timing, and triggers to up to 24 FEBs through the Cat6 Ethernet connectors on both boards. In return, the ROC receives data and slow control information from its FEBs, which will be transferred to the DAQ system. Each connection from a ROC to an FEB has 100 Mb/s bandwidth and can supply up to 24 W of power.



Figure 4.13: Photograph of a CRV ROC. The main components are labeled in yellow [17]. The 24 RJ45 Ethernet connectors on the bottom of the board connect to the FEBs. The final ROC design is essentially identical to the initial design.

As seen in Figure 4.13, the 24 input RJ45 ports are divided into three groups of 8 on the front panel of the board. Each group of 8 FEBs is serviced by a Xilinx Spartan6 FPGA that handles event building and temporary storage with a 2 Gb LPDDR memory chip [14]. The accumulated events from each of these groups are transferred to an FPGA-microcontroller pair that handles packaging the event data into packets and transmission to the rest of the DAQ through a Data Transfer Controller (DTC). Two 2.5 Gb/s optical fiber transceivers are present on-board; one of which interfaces with the DTC to receive triggers and status requests and to send FEB data and status information back and the other is to daisy chain multiple ROCs together. The FPGA-microcontroller pair also handles the receipt and processing of triggers, status requests, and other FEB controls [4]. The FPGA in the pair is a Xilinx Spartan6 and the microcontroller is a Texas

Instruments Stellaris [14, 18]. Unlike FEB production, all Spartan6 FPGAs for the ROC boards were received before the Spartan 6 became obsolete. The last main component on board the ROC is a 40 MHz timing link that establishes synchronization of the CRV readout electronics with a clock from the accelerator. A block diagram of the data path for events read from the FEBs is shown in Figure 4.14.

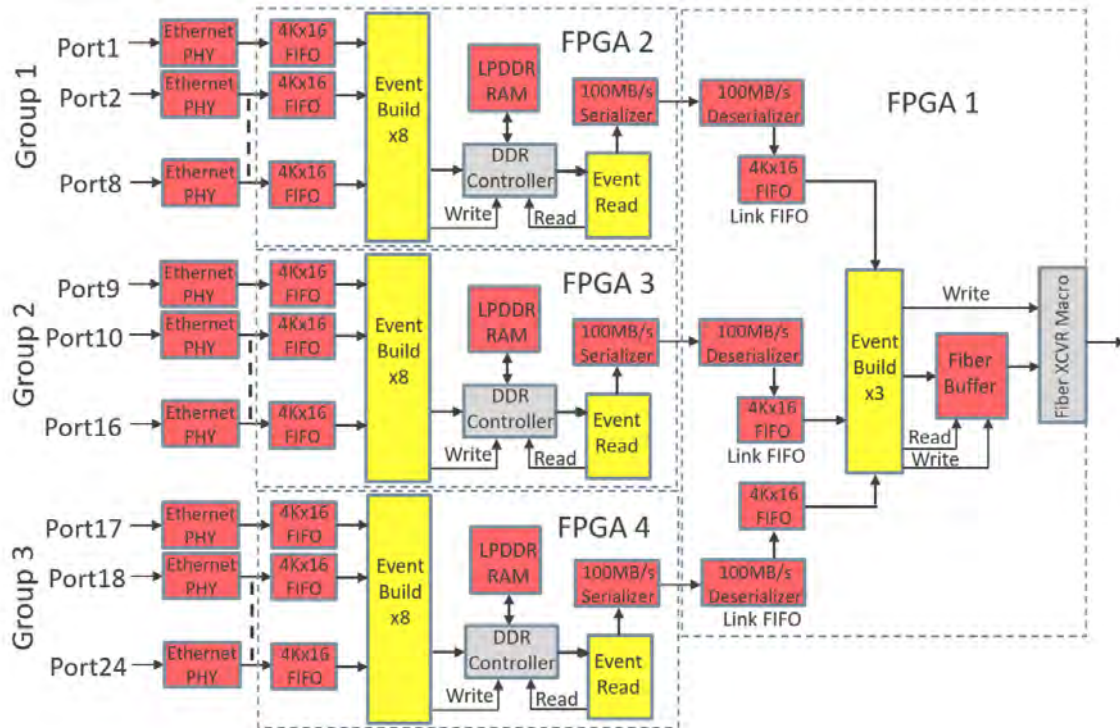


Figure 4.14: Block diagram of ROC data flow path. This block diagram shows the flow of data through the ROCs as data comes in from the FEBs [4].

Each ROC is programmed and tested at Fermilab to check the nominal power supply and consumption, communication interfaces, and operation of the microcontroller, FPGAs, and memory. Memory is particularly important in the ROC because fast and slow status information is sent and received back and forth from the FEBs to the ROC. There is a block of memory space in the ROC that is allocated to store detailed status information for up to 24 FEBs and the ROC itself. The DAQ accesses this memory block through the fiber link to read the status information as needed. Control of the CRV electronics parameters is also done through settings sent from the DAQ to the ROCs and then distributed to the FEBs. There is no immediate feedback when setting the readout system parameters, but the status blocks in the ROC memory are updated to reflect control changes [4].

4.3.4 Data Acquisition

As a whole, the Mu2e DAQ system is tasked with monitoring, selecting, and validating physics and calibration data from the Mu2e detector for final stewardship by the offline computing systems [5]. This system also must combine information from the many detector sources in the Mu2e experimental apparatus and apply filters to greatly reduce the average data volume before it can be transferred to offline storage. In addition, the DAQ must also provide a timing and control network for synchronization of the data sources and readout.

The central component of the Mu2e DAQ system is a commercial 3U server, which manages data collection from the ROCs, event building, and online processing. There are a total of 36 DAQ servers occupying four racks in the electronics room to serve the whole experiment [5]. The link between the ROCs and the DAQ is the DTC. The DTC collects data from multiple ROCs and can perform event building and data pre-processing. For Mu2e, the DTC is implemented using a commercial PCIe card located in the DAQ server [5].

References

- [1] D. Beznosko, A. Bross, A. Dyshkant, A. Pla-Dalmau, and V. Rykalin, *FNAL-NICADD Extruded Scintillator*, Tech. Rep. (Fermilab, 2005) [FERMILAB-PUB-05-344].
- [2] D. Michael, *et al.* (MINOS Collaboration), *Nucl. Instrum. Methods Phys. Res. A: Accel., Spectrom., Detect. Assoc. Equip.* **596**, 190 (2008).
- [3] L. Aliaga, *et al.* (The MINERvA Collaboration), *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.* **743**, 130 (2014).
- [4] S. Boi, Ph.D. thesis, University of Virginia (2021).
- [5] L. Bartoszek and *et al.* (Mu2e Collaboration), *Mu2e Technical Design Report*, Tech. Rep. (Fermilab, 2015) arXiv:1501.05241 [ins-det] .
- [6] E. C. Dukes, P. J. Farris, R. C. Group, T. Lam, D. Shooltz, and Y. Oksuzian, *JINST* **13**, P12028, (2018).
- [7] Kuraray, “Wavelength Shifting Fibers,” (2017), datasheet, Available at: <http://kuraraypsf.jp/psf/ws.html>.

- [8] M. Solt, S. Roberts, Y. Oksuzian, and D. Coveyou, “Mu2e 1.8mm Fiber,” (Mu2e Summer 2021 Collaboration Meeting 2021-06-23), mu2e-doc-38615-v1.
- [9] L. Xia, “Introduction to CRV Readout System and FEB,” (CRV FEB CRR 2021-05-25), mu2e-doc-38267-v4.
- [10] Hamamatsu, “MPPC S13360-2050VE,” (2016), datasheet, Available at: https://www.hamamatsu.com/us/en/product/optical-sensors/mppc/mppc_mppc-array/S13360-2050VE.html.
- [11] V. Zutshi, “Mu2e CRV: Photodetectors,” (DOE IPR 2019 2019-12-10), mu2e-doc-30529-v1.
- [12] Texas Instruments, “AFE5807,” (2013), datasheet, Available at: <https://www.ti.com/product/AFE5807>.
- [13] Texas Instruments, “RM48L952,” (2015), datasheet, Available at: <https://www.ti.com/product/RM48L952>.
- [14] Xilinx, “Spartan-6 Family Overview,” (2011), datasheet, Available at: <https://docs.xilinx.com/v/u/en-US/ds160>.
- [15] Xilinx, “7 Series FPGAs Data Sheet: Overview,” (2020), datasheet, Available at: https://docs.xilinx.com/v/u/en-US/ds180_7Series_Overview.
- [16] L. Xia, “Mu2e CRV: Electronics,” (DOE IPR 2019 2019-12-10), mu2e-doc-30124-v4.
- [17] S. Corrodi, “The High-Efficiency Cosmic Ray Veto Detector for the Mu2e Experiment at Fermilab,” (NuFact 2022 2022-08-05), <https://indico.fnal.gov/event/53004/contributions/244339/>.
- [18] Luminary Micro Inc., “LM3S5B91 Microcontroller,” (2009), datasheet, Available at: <https://www.ti.com/en/download/shared/product-brief-LM3S5B91.pdf>.

Chapter 5 Mu2e Software Framework

After each detector in the Mu2e apparatus collects data from particles that pass through the detector volume, software is necessary to aggregate the information coming from each Mu2e detector for every event in order to select and reconstruct events of interest. The Mu2e software framework contains code for simulating, calibrating, reconstructing, analyzing, and graphically viewing events from Mu2e. Much like the Mu2e detectors, the Mu2e software framework uses many off-the-shelf software packages with customized settings when possible rather than develop all necessary tools in-house [1]. When a particle interacts with the Mu2e detectors, each detector sends data to its DAQ to be filtered. These filtering decisions performed by the DAQ happen almost instantly during on-spill time with the beam and off-spill time between proton beam pulses. The software that performs these decisions is called the "online" software. The events that are selected by this online filtering process are then transferred to storage disks where more advanced event analysis may occur. The analysis and processing that occurs on filtered events in storage is not bound by a strict time limit set by the beam, and the software that performs this advanced reconstruction and analysis is called the "offline" software. The Mu2e detectors have already been designed to reduce backgrounds from undesirable muon decay modes or other beam particles. Similarly, the experimental software has been designed to reduce background and record only interesting events where conversion electrons may be present. Efficient and fast event filtering and selection, data transfer, data storage, and analysis are all performed using software and are critical to the success of the experiment.

At this stage in the Mu2e lifetime, the software is currently under development and being trained with simulated data. By simulating different kinds of events in the Mu2e detector environment, the event selection and reconstruction code can be developed and trained in parallel to the fabrication of the Mu2e detectors so that it will work on experimental data once it is available. Realistic simulations will ensure that our analysis framework can handle the output data rates of Mu2e and select meaningful events which may have a conversion electron present. Success of the Mu2e experiment hinges on the event selection and reconstruction framework to sort through the massive amount of data that the experiment will collect.

Once the Mu2e detectors are functioning, they will be constantly interacting with particles from the beamline and the environment. To achieve this, a fast "online" software takes data from the

tracker and calorimeter and uses pattern recognition to search for particle trajectories in the detector data. The layer of software that determines whether or not to record an event is called the trigger. If the event passes all filters of the trigger, then data from the CRV is aggregated into the tracker and calorimeter event data and the full event is written out to disk files for further analysis with the "offline" software. This chapter will describe the overall Mu2e software framework, contextualize the Mu2e online software, trigger, and offline software and describe the function of each of these tools, and then discuss three studies that were performed using this framework.

5.1 Mu2e Framework Architecture

Before discussing how the Mu2e online and offline software works, it is important to discuss two ideas: data products and base software packages. The first topic, data products, refers to the analog to digital conversion of electrical charge signals left by passing particles in the Mu2e subdetectors to objects in software-space that we can use to build Mu2e events. The second topic covers the two important off-the-shelf products that Mu2e uses to process and simulate events, effectively doing the "heavy lifting" when it comes to event building and generation. Then, we can describe the event building process and how data from the multiple Mu2e subdetectors is aggregated into a single event. After the events are built, we can discuss the online software framework and how pattern recognition and filtering are used to select the events for offline analysis.

5.1.1 Event Building

Event building is the process of sending data from the various different Mu2e subdetectors through the DAQ system and converting the event data into a format that is readable to an online executable. All Mu2e subdetector systems interface with the DAQ via a ROC board. The ROC sends data to DTC boards via optical fiber. As data is streaming through the DAQ system, each DTC is receiving fragments of events from different ROCs. To continue with event building, all fragments from a given event must be transferred to the same DTC. The DTCs can send and receive event fragments to and from each other and rearrange the fragments between them [2]. Lastly, with the event fragments collected, the event building process concludes by converting the data from binary to art input to make the event readable by the online executables that will be used for the next steps of pattern recognition and analysis.

5.1.1.1 Data Products

As we have been alluding to, Mu2e filters events using different criteria on event characteristics, such as the momentum, energy, and timing of particles that have been ejected from the stopping target. We use objects called "data products" to conceptualize the signals that particles make in

Mu2e subdetectors. The data in an event in a file is organized into data products. Executable files then access these data products and decide which data products to read, create, and write out. Each Mu2e subdetector has its own data products.

For the tracker, the simplest data product is called a "hit", which translates to a point on a straw where a particle interacted with the detector electronics at a particular time. The signals at both ends of a tracker straw are used to reconstruct the hit time and longitudinal position of the hit along the straw. To improve the robustness of hits in the tracker, neighboring hits from the same panel in the two layer of straws are then combined to form a "panel hit". This combination of hits from the two layers in a panel improves spatial resolution along the wire by a factor of about $\sqrt{2}$ [3]. During pattern recognition steps of reconstruction, full trajectories of particles moving through the tracker, called "tracks", are created by looking for helical patterns in panel hits. As tracks are built, new data products are created that become more complicated until the full track is encapsulated in a single data product.

For the calorimeter, each crystal is read out, so any active crystals are identified and groups of activated crystals form a data product called a "cluster". Calorimeter clusters have a well defined position and time based on the signals from crystals that were activated and also have a well defined energy based on the incoming particle's deposition in the CsI crystals.

For the CRV, the data product that signifies the presence of a cosmic ray is called a "stub". The requirement on cosmic rays is that a cosmic ray must interact with three of the four layers in the CRV. Thus, a CRV stub is defined as a data product where a cosmic ray was detected in three out of four CRV layers at the same time. Similarly to the tracker, each CRV dicounter is read out at both ends, so timing signals at both ends can be used to determine the longitudinal position where a cosmic ray passed along the length of the dicounter. Each dicounter is read out with four SiPMs, so it is trivial to determine which CRV bars detected a particular cosmic ray particle.

5.1.2 Key Base Software Packages

As mentioned in the introduction to this chapter, Mu2e software uses many off-the-shelf products in our software framework. In addition to many small third-party software packages, two major functions that are performed by off-the-shelf software are event processing and event simulation. The main event processing software is called art and the main simulation software is called Geant4.

5.1.2.1 art

The software for Mu2e is based on an event-processing framework called art. art is a C++ framework that evolved from software from the Compact Muon Solenoid (CMS) experiment at

CERN that was adapted and designed to be used by many next-generation particle physics experiments at Fermilab. Many of the authors of art were involved with the design and implementation of event processing frameworks for experiments like BTeV, CMS, DØ, and MiniBooNE. Each of these experiments have similar framework requirements, but the development efforts for infrastructure software were isolated, which resulted in duplication of many efforts [4]. art aims to solve this problem and avoid duplication by providing a generalized, but customizable framework with low maintenance needs. The art framework itself is developed and maintained by specialized, particle physics software engineers at Fermilab to provide a robust foundation for developing user code for different experiments.

When an art program is run, it reads in events from some user-specified input source, invokes user-specified *modules* on the events, and then writes the results out to output files. A *module* is a piece of typically user-written code that can implement algorithms on the event data [4]. A domain-specific language called FHiCL (pronounced "fickle") is used to write configuration files that define the input data source, set the values of any configurable settings, specify the sequence in which modules are carried out on the data, and write data out to output files. When art creates output files containing physics data, these files have a specific file structure and naming convention. This creates uniformity in data format and allows for different analysis modules and workflows to be written to expect the same input data format.

5.1.2.2 *Geant4*

The simulations that have been created to demonstrate that Mu2e is a viable experiment with obtainable physics goals were made using a software called Geant4. Geant4 is a toolkit to simulate the passage of particles or radiation through matter. This package can generate simulated events of interest in any environment and simulate the resulting particle interactions in response to that event. This toolkit includes a complete set of physics processes for electromagnetic, strong, and weak interactions of particles in matter over an energy range that starts from milli-eV, eV, or keV, up to hundreds of GeV or even TeV level [5]. Each type of interaction may be implemented through different models; some different sets of modeling approaches are available in coherent configurations called "physics lists". Mu2e uses a custom model via a physics list in Geant4 called "ShieldingM".

In conjunction with Geant4, simulations of the Mu2e detector apparatus and environment have been developed, including a simulated magnetic field map which aligns with the field that will be generated by the Mu2e solenoid system. Geant4 can simulate any setup, detector, and radiation source and then record any output that results from source particles and secondaries interacting with the material of the setup [5]. Geant4 also provides visualization power. Using Geant4, models may be created and viewed of a detector geometry with different materials and shapes. Tracks of

particles passing through the models may be visualized via a graphical user interface.

5.2 Mu2e Online Software

Now that we have described key pieces of the underlying architecture of the Mu2e software framework, it is finally time to begin discussing the Mu2e Online software. Online software is an extremely fast way to process events, beginning with data product preparation, running through pattern recognition algorithms for track finding in the tracker and calorimeter, and concluding with data storage. The online reconstruction and trigger systems are closely related, online reconstruction creates the data products necessary for the trigger to make a decision on each event. The following requirements have been set for the online reconstruction and trigger system [3]:

- Provide efficiency of at least 90%
- Keep the trigger rate below a few kHz - equivalent to ~ 7 Pb/year
- Achieve a processing time < 5 ms/event

The first step of the online reconstruction is searching for tracks by preparing data products for the tracker and the calorimeter. For the calorimeter, clusters are created using activated calorimeter crystals. For the tracker, panel hits are generated using signals at each end of the tracker straws. After creating all panel hits from the data in an event, a multi-variate analysis (MVA) algorithm identifies and flags "background hits". Any hits that are compatible with low-momentum tracks on the order of a few MeV are flagged, as they may come from particles like Compton electrons. The flagged background hits are not used in the next stage of pattern recognition. The online reconstruction algorithm does not yet aggregate CRV data into an event until a trigger decision has been made to optimize processing speed.

5.2.1 Pattern Recognition

With data products in hand for the tracker and calorimeter, the next step of online processing is pattern recognition and track finding. Pattern recognition is the process of building helical shaped tracks from the individual panel hits in the tracker. Since Mu2e is searching for a monoenergetic electron ejected from the stopping target which travels through a constant magnetic field in the tracker region, the expected path that a CE takes through the tracker is well defined. The high momentum of CEs in the DS dictates the radius of the helical track and the stopping target provides a range of initial positions where we expect CEs to originate. Mu2e uses two algorithms to search for helical tracks in prepared panel hits: Tracker-based pattern recognition (TPR) and Calorimeter-based pattern recognition (CPR). These algorithms are run in parallel to search for tracks from both TPR and CPR in the same event.

5.2.1.1 Tracker-based Pattern Recognition

The TPR algorithm only uses data from the tracker, ignoring calorimeter data, to search for tracks. Pattern recognition proceeds in two steps: hit time clustering and helix reconstruction. Time clustering involves using an MVA-based algorithm to identify peaks in the distribution of panel hit times. Since panel hits occur at different points along the axis of the beamline as the particle travels through the DS, all panel hit times are extrapolated to the center of the tracker, the point $z = 0$ in the tracker frame of reference. Extrapolation requires assumptions that $\beta = v/c = 1$ and that the track pitch angle is equal to the average value expected for a CE. This extrapolation of all panel hits to the same point along the beam axis decouples the particle's time of flight from the width of the peaks in panel hit times that the MVA algorithm is searching for [3]. Time clusters are used to identify how many particles are in an event and whether or not the timing of the particle is consistent with that of a CE.

Helix reconstruction has two parts: circular track reconstruction in the transverse X-Y plane perpendicular to the beam and line reconstruction in the ϕ -Z plane, where ϕ is the hit polar angle with respect to the helix axis, and Z is the coordinate along the tracker central axis. Figure 5.1 shows how hit positions of a helical track are projected into the X-Y and ϕ -Z planes. Circle reconstruction




Figure 5.1: Helix panel hits projected into X-Y and ϕ -Z planes. Cartoon projections of the panel hit positions of a helix through the tracker into the transverse X-Y (left) and ϕ -Z (right) planes [3].

in the X-Y plane begins with looping over all panel hits from different tracker planes to form groups of hit triplets. If a triplet covers a sufficient area in the X-Y plane and is within the span of a few tracker planes in the Z direction, the expected helix center is evaluated by finding (x_0, y_0) . After looping over all possible combinations of hit triplets, the best estimate for the helix center is determined by taking the median of all of the (x_0, y_0) values collected. After the helix center has been estimated, a loop is repeated over the panel hits to estimate the circle radius using again the median of all the single values [3].

Line reconstruction in the ϕ -Z plane begins with resolving what is called the "2 π ambiguity" in the panel hits. In Figure 5.1, the 2 π ambiguity arises from the hits in the second loop of the helix. When the hits of a helix are projected into the ϕ -Z plane, the presence of the hollow space in the center of the tracker creates periodic gaps between the hits in different loops of a helix. To

correct for this 2π ambiguity, a factor of $2\pi \times i$ should be added to the ϕ of the hits in the i^{th} loop. The correction depends on the particle's angular velocity $\frac{d\phi}{dz} = \frac{1}{\lambda}$, which is estimated using two histograms. First, the peak of the histogram $\frac{d\phi}{dz_{i,j,k}} = \frac{(\phi_j + 2\pi k) - \phi_i}{z_j - z_i}$, with $i, j \in (0, N - 1)$ and $k \in (0, 10)$ is used to resolve the 2π ambiguity and assign each hit to its corresponding helix loop. Then, the peak of the resulting histogram $\frac{d\phi}{dz_{i,j}} = \frac{\phi_j - \phi_i}{z_j - z_i}$ is used as the best estimate of the helix $\frac{d\phi}{dz}$ [3]. After time clustering and helix reconstruction by TPR, helix data products are created which encapsulate this information. These helix data products will continue to be refined as we move through the rest of the data processing chain.

5.2.1.2 Calorimeter-based Pattern Recognition

The CPR algorithm incorporates data from the calorimeter to search for tracks, so if no calorimeter cluster is present in a given event, CPR will not run. A fast calorimeter reconstruction algorithm runs before and track reconstruction algorithms begin, so energetic calorimeter clusters are available to seed the track finding pattern recognition [6]. The pattern recognition search of grouping tracker panel hits into triplets in CPR is similar to that of TPR, but in this case, information from the calorimeter selects only panel hits in time and space that are coincident with the calorimeter cluster. The panel hits that are looped through to search for a helix must be within a ± 40 ns time gate with respect to the calorimeter cluster time and also must lie in the same semi-plane in the X-Y transverse direction that the calorimeter cluster was located in. Figure 5.2 below illustrates the effects of hit selection using CPR.

After hit selection, the CPR algorithm creates a triplet of hits using the calorimeter cluster position as one hit, a panel hit close to the stopping target as another, and the solenoid center in the X-Y plane as the last. These three points define an area where the helix search is conducted. Including a panel hit close to the stopping target in this initial triplet assumes that the particle originates from the stopping target. All panel hits are looped through, picking one hit in each tracker plane within the search area defined by our initial triplet. If a second hit is found in a given tracker plane, then the solenoid center is dropped from the points in the triplet and the panel hit that was found replaces it and the search restarts with the new triplet as the search area. As new hits are looped over and added to the helix candidate, the search adjusts the selected triplet. The helix parameters are updated using two different reduced- χ^2 fits: one in the X-Y and one in the ϕ -Z planes [3]. The hits are weighted by the inverse of the square of the expected uncertainty and hit errors can be calculated using these uncertainties to account for the orientation of the straw tubes in the tracker with respect to the reconstructed helix. At this point, helix data products are created for the CPR search that are the same format as the helix data products of the TPR search. Using both algorithms in parallel improves the efficiency of track reconstruction for CE tracks.

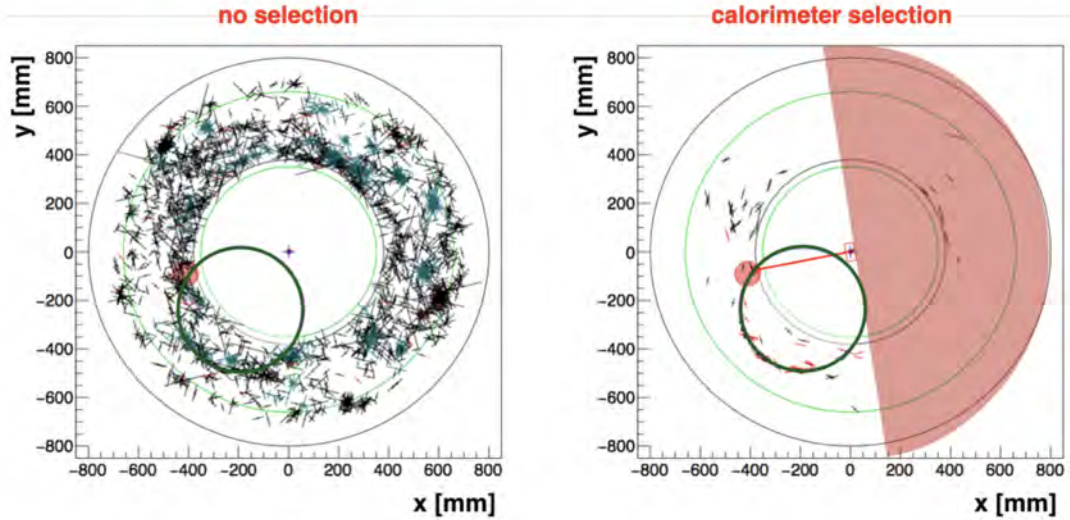


Figure 5.2: CPR selection of hits in X-Y plane. A 2D event display image of a typical CE event with background in the tracker in the X-Y plane. Each black dash represents a single hit, the red circle represents the calorimeter cluster in the X-Y plane, and the green circle represents a helical CE trajectory. The left image shows all panel hits before CPR selection and the right shows panel hits after CPR selection [3].

5.2.2 Track Reconstruction

The result of TPR and CPR pattern recognition algorithms are helix data products, which help make trigger decisions. However, a more robust data product than a helix is called a track. After a helix corresponding to a track candidate is found, a simplified Kalman fit is performed to improve the accuracy of reconstructed track parameters and background rejection [3].

At this stage, there is one unresolved problem with respect to panel hits that is worth mentioning: the "left-right ambiguity". When particles pass through the tracker and interact with the tracker straw tubes, the signal that is generated by a particle passing to the left of the sense wire of a given straw is identical to the signal that is generated if the particle was passing to the right side of the sense wire instead. This left-right ambiguity causes uncertainty on the hit positions. There is also an unresolved problem with the physics of particles moving through the tracker. Online reconstruction does not consider energy loss as particles travel through the DS. Resolution of these two issues will occur later during offline reconstruction.

5.3 Mu2e Trigger

The Mu2e trigger system is implemented through a series of filters at different steps along the data processing chain. The main physics triggers use information about reconstructed tracks from online track finding to make a final decision on whether or not an event is accepted [3]. Any

producer art modules have filter modules in between to stop execution of reconstruction if a filter flags an event. Using filters in this way reduces execution time and the amount of data that is sent to offline storage.

While online is running and proceeding through pattern recognition, the trigger is also running as a large set of filters. Multiple triggers are implemented as multiple independent reconstruction paths, each path running one or several reconstruction algorithms followed by a software filter to make the trigger decision [6]. Only on top of filtered events is CRV data aggregated into each event to search for the presence of a cosmic ray. After requests are sent for filtered events to obtain CRV data, the data is aggregated and sent to the data logger. In this way, event selection is complete at the end of online processing and the trigger is only run in parallel with online. The data logger sends the data to the output stream dispatcher and then the data is stored on-campus at Fermilab's Feynman Computing Center (FCC). At this time, the data logger performs some data quality monitoring and separates on-spill from off-spill events [2]. In conclusion, the Mu2e online processing chain runs in parallel with the trigger to do a preliminary reconstruction of events where filters are placed between subsequent reconstruction stages to screen out background events. The result of online processing and triggering is a stream of candidate CE events with information from all Mu2e subdetectors sent for storage at FNAL for further offline processing.

One metric to evaluate how well the trigger is working is by calculating the trigger efficiency. Trigger efficiency can be calculated as the efficiency of the total trigger system, or of individual trigger paths. For example, two independent trigger paths are triggers that use either the TPR or CPR track finding algorithm. The trigger system makes a final decision based on the OR of the two pattern recognition algorithms [6]. Figure 5.3 shows the trigger efficiency for the simulated conversion electron events which have a reconstructed track passing the offline selections. The total efficiency of using a combination of the CPR and TPR pattern recognition algorithms is well above 90%, satisfying one of the requirements of the trigger system. From Figure 5.3, the red markers indicate that the trigger performance is stable based on the OR of the output of running the TPR and CPR algorithms in parallel, emphasizing the importance of using both algorithms in online track finding. For CPR track finding, efficiency decreases with higher beam intensity because the algorithm is limited by the calorimeter acceptance and the trigger requirement on the seed cluster energy, $E > 50$ MeV [6]. However, the total trigger efficiency is almost independent of the beam intensity and ranges from 99% at zero beam intensity to 97% at 1.2×10^8 protons/pulse.

The Mu2e trigger is the most important tool in controlling the output data rate of the experiment. The Mu2e experiment is expected to generate 14 pB of data per year, so a high rejection factor is needed to reduce this rate to reasonable storage rates, determined as < 7 pB per year [2]. This is the reason that as soon as an event hits a filter and is flagged as a background, the event is not stored for offline analysis. In this way, Mu2e does process *every* event in the DS, but




Figure 5.3: Trigger efficiency for CE events. Trigger efficiency of CE events relative to offline reconstruction efficiency as a function of proton pulse intensity. The red markers show the efficiency of the final trigger, whereas the green markers represent the efficiency of only the TPR pattern recognition algorithm and the pink markers represent the efficiency of only the CPR pattern recognition algorithm [6].

not every event is stored. Some trigger-level histograms are populated for *every* event, so some trigger-level analysis studies with large statistics are possible for the future. For example, some metric of muon stop intensity could be recorded even for rejected events. This information could be used for normalization and calibration purposes in addition to the STM since the metric would be proportional to beam intensity.

5.4 Mu2e Offline Software

As a result of online processing and triggering, only interesting events are selected from the raw Mu2e data stream and passed to offline. The Mu2e offline software is the most robust way to reconstruct Mu2e events and search for CEs. Offline uses the same reconstruction algorithms as online, such as TPR and CPR, but using more stringent reconstruction conditions and less filtering, since offline is not restricted to running in-time with the experiment. Offline reconstruction also resolves standing problems after online processing, like the "left-right ambiguity", consideration of a particle's energy loss as it travels through the DS, and effects of multiple scattering. Calorimeter data is also processed differently in the offline framework; it is treated more simply during online processing due to computing time constraints. During online processing, only the energy and the transverse position of the calorimeter are known. During offline processing, the full Kalman fit determines the Z position of the cluster, its timing, and coordinate residuals [6].

As soon as data is stored on FNAL disks, offline data processing begins. The default reconstruction chain for data processing is called "Pass 1 Reconstruction". Pass 1 is where TPR and CPR are run and the online-standing issues in the data are resolved. In addition to Pass 1, any Mu2e user may write user code within offline to run independent analyses. This is done through the use of customized FHiCL files where users may specify or write custom modules for analysis.

5.5 Studies Using Offline

When I began research with Mu2e, some of my first studies involved self-contained investigations into different background modes and how we could use software to improve our understanding and rejection of these background modes. The first study that I will discuss here is estimating the background rate from slow-moving antiprotons that travel down the TS into the DS and happen to stop in the stopping target. For this study, I begin by designing and running my own simulations to create the input files and desired simulated events and continue to estimate how many of these simulated events can mimic a conversion electron. After this, I was involved in two studies using large-scale simulated cosmic ray data. In one study, I increased the trigger efficiency for certain types of cosmic events by tuning some parameters associated with track fitting. In the last study, I optimized the length of the CRV dead-time window that is employed when a cosmic particle is detected inside the DS using the data products on event timing. This optimization benefits the overall experiment by increasing the amount of live-time during which conversion electron events may be recorded.

5.5.1 *Estimating the Antiproton Background*

Slow, low-energy antiprotons traveling down the beamline from the production target are a dangerous type of background because these particles have no temporal correlation with beam pulses. Estimations for the number of fake CEs per stopped antiproton annihilation have fluctuated over the years as versions of GEANT4 have upgraded. This stems from the inability to model $\bar{p}p$ annihilations in our kinematic range of interest. The goal of this study was to independently create stopped antiprotons in the stopping target and simulate their annihilation to validate past estimation results.

To accomplish this goal, a few separate studies were completed to both estimate the expected number of signal electrons from antiprotons stopped on the muon stopping target and validate the use of GEANT4 simulations to accurately model antiproton annihilation. Two of the studies estimate the expected number of signal electrons per stopped antiproton on the muon stopping target using different kinds of stopped antiprotons. The last study seeks to validate GEANT4 simulations by comparing plots that GEANT4 produces today to historical, experimentally measured data for

antiproton annihilations.

5.5.1.1 Simulating Stopped Antiprotons in the Stopping Target at Random Positions

For this study, the goal was to generate randomly positioned stopped antiprotons in the stopping target, and then allow these antiprotons to annihilate and evolve as they travel through the Mu2e simulated detectors. These simulations result in a set of reconstructed events where the final state particles may be analyzed for CE signals, for example simulated spectra of final momentum may contain electron events near 105 MeV.

Generally, studies that seek to estimate the number of signal electrons per stopped antiproton on the stopping target proceed in two stages. The first stage generates the actual antiprotons stopping on the stopping target, a FHiCL file that contains appropriate generate blocks and corresponding art module to generate the desired antiprotons. The second stage takes care of the interaction and evolution of particles after they have been stopped, which can be achieved by using TrkAna FHiCL files that have already been developed for this purpose. TrkAna, meaning track analysis, is a simulation analysis package developed within Mu2e. This package takes .art files as an input and generates ROOT files containing a standardized ROOT TTree where each entry in the TTree corresponds to a single fitted track and contains reconstructed information from the tracker, calorimeter and CRV [7]. The following sections will describe each stage of simulating randomly stopped antiprotons, and there will be corresponding sections later in the discussion of using the SU2020 stopped antiproton files as input.

First Stage FHiCL Design

First, two files are required to create antiprotons in the stopping target: a FHiCL file and a corresponding art module. The FHiCL file holds the generate block and is responsible for the creation of .art files which contain the stopped antiprotons on the stopping target. The module fills data into histograms as it determines various characteristics of each particle: the particle's position, energy, momentum, and time of ejection from the stopping target. To create the FHiCL and corresponding module which produce randomly distributed antiprotons in the stopping target, pre-existing code which does something similar was used as a template, then different pieces were customized to generate the desired particles with specific characteristics.

The FHiCL file `pbartest_v2.fcl` is executed on Mu2e virtual machine nodes. The major block in this file is the generate block, which calls `pbarStopTarg_module.cc`. In the generate block, we may explicitly state to only make antiprotons, particles with PDG ID -2112, in a small kinetic energy range, so the antiprotons stop on target but are ejected moving downstream. The next sections of the FHiCL file configure Geant4 reconstruction settings and set the analyzer of the module appropriately for generated particles. Lastly, the reconstruction paths and filenames for the output .art and .root

files are set.

The module `pbarStopTarg_module.cc` explicitly chooses characteristics of the antiprotons like position, energy, and time of ejection. A 4-component struct holds the coordinates (x , y , z , t) for each generated antiproton in the Mu2e detector coordinate system. The first section of the module chooses the position of the antiproton in the stopping target. There are many geometry handles in the Offline world, this module requires geometry handles for the stopping target foils, which are already simulated in Offline. The position of each antiproton is chosen with a foil number, a radius, and an angle ρ in the stopping target. For the first set of tests, a random position is chosen by selecting a random foil number from $[1, n_{\text{foils}}]$, a random radius between the min and max of the chosen foil, and a random angle in 2π . The time for all ejections is set to 1000ns. The energy for the generated antiprotons is chosen at random using the kinetic energy range set in the FHiCL file, between $[0.1, 0.2]$ MeV. Then, the module calculates the total energy and momentum from the randomly chosen kinetic energy. Each component of the momentum is given a randomly selected weight; the z -component of the momentum is always positive such that all particles after annihilation travel downstream, but the x - and y -components of the momentum are also given random signs. The module then writes the generated antiprotons out, completely described by their position in the stopping target, momentum, time, PDG ID, and GenId, and generates ROOT histograms to show these distributions.

At the end of the first stage of stopped antiproton generation, the FHiCL module returns a `.root` file containing histograms of characteristics of the stopped antiprotons, as included in the module, and an `.art` file containing all of the stopped antiprotons which is used as an input for the second stage `TrkAnaDigisReco` jobs.

Second Stage TrkAnaDigisReco

For the second stage of antiproton annihilation, the stopped antiprotons in the stopping target must interact and evolve in the Mu2e virtual detector solenoid. Conveniently, the Mu2e Offline infrastructure is already equipped with modules which take generated particles as an input and complete the evolution of the particles, which ultimately results in electrons being detected in the tracker and calorimeter. One of the Offline directories that allows you to take any `.art` file which contains generated particles and evolve them in the Mu2e virtual detector system is called `TrkDiag`.

Within `TrkDiag/fcl/`, the FHiCL file called `TrkAnaDigisReco.fcl` is a script which runs `TrackAnalysisReco` on digi inputs, which is appropriate for these second stage jobs from stopped antiprotons on target. This script runs the official reconstruction sequence inline and returns a `.root` file with an extensive set of histograms, separated into two collections depending on the charge of the final state particles in each event (`TrkAnaNeg` and `TrkAnaPos`). Events in the `TrkAnaNeg` distributions are events with negatively charged particles in the final state. Events in the `TrkAnaPos` distributions are events with positively charged particles in the final state.

Results

From the large number of stopped antiproton decays that have been simulated in the stopping target, we only expect a small fraction of these decays to have an electron in the final state that is reconstructed with characteristics resembling that of a CE. Previous efforts to calculate the estimated number of signal electrons per stopped antiproton have been developed and these estimates are in the neighborhood of 10^{-5} to 10^{-6} signal electrons per stopped antiproton on target. The latest effort to calculate this background estimation was done in 2018 [8].

To estimate the number of signal electrons per stopped antiproton on target, we apply conditions, called cuts, to the final state electrons. The cuts that we apply select any events that have a final state electron with the characteristics of a conversion electron. Since the only process occurring here is the annihilation of many stopped antiprotons, the resulting electrons from this simulation represent a spectrum of background that could be misconstrued as a signal electron. Here, a set of cuts were applied that are widely used throughout the collaboration, 'Cut Set C', and a few additional cuts suggested to me by a collaborator who is an expert in simulations [9]. These cuts include information about whether or not a track was successfully reconstructed for an event and what the trajectory of the particle looks like as it moves through the DS. Any particles that travel too close to the beam axis, have tracks with very large or small radii, or have extreme impact parameters with respect to the beam axis are vetoed. Particles that are not identified as electrons or particles that are reconstructed to be traveling upstream are vetoed. Any particles outside of the delayed selection window of the beam are vetoed by their arrival time. All events that could be a result of cosmic rays entering the detector solenoid by vetoing and track that is close in time to a signal in the CRV. And lastly, any event with a poor quality track reconstruction is vetoed using a parameter that calculates track quality [9].

True conversion electrons also have a well-defined momentum as a result of the decay at rest from the muonic aluminum nucleus. In addition to applying 'Cut Set C' with these extras, we also apply a cut on momentum, selecting only final state electrons with a momentum between 104 and 105 MeV/c. After applying the cuts described above and a momentum cut, we may find the number of signal-mimicking electrons in each simulation. The number of antiprotons on target is set by the number of antiprotons generated with Geant4. Using these numbers as the numerator and denominator, we may estimate the background rate of conversion-like events from stopped antiproton annihilation. We decided to analyze the momentum spectrum of downstream-reconstructed electron events at the entrance of the tracker.

Random Stop Positions

A straightforward way to estimate the number of signal electrons per stopped antiproton on target is to generate a bunch of randomly stopped antiprotons in the Mu2e stopping target and only count the final state electrons with nearly the same momentum that we expect to see from a

conversion electron event. This approach is exactly how the first round of this study was completed. Mu2e's stopping target is located inside of the detector solenoid roughly 2 meters in front of the tracker. The aluminum stopping target contains 34 thin foils, each 0.1056 mm thick with a radius of 7.5 cm and 2.15 cm holes in the center [10]. Refer to Figure 3.8 in Chapter 3 for a diagram and photograph of the stopping target. In the first stage of generation, a randomly chosen integer sets the antiproton's foil and randomly fired radii and phases determine its' position on a foil. The distribution of stopped antiprotons can be verified by looking at histograms of the generated antiprotons' positions in x, y, and z. In these spatial distributions, one can verify that the antiprotons are only stopping in the foils by inspecting the x and y distributions for a dip around the center of the distributions since the foils have a hole in the center where particles cannot be stopped. One can also verify that antiprotons are stopped in every foil of the target by counting the quantized bunches of stopped antiproton positions along the z-axis.

In Figure 5.4, there are a wide range of possible momenta for final state electrons from antiproton annihilation in the stopping target. Remember, a signal CE in Mu2e has energy that is slightly below the muon rest mass and the conversion electron energy depends on the stopping target material. In the Mu2e aluminum target, the conversion energy is 104.973 MeV [10]. During analysis, a small momentum window constitutes the acceptable momentum range for a conversion electron. In this investigation, the number of electrons in the 104-105 MeV bin are counted as conversion electrons to estimate for number of conversion electrons per antiproton on target. In Figure 5.4, there are 940 electrons in the bin of interest, these are the mimic signal electron candidates.




Figure 5.4: Randomly stopped final state momentum distribution. This figure is a momentum distribution of final state electrons at the tracker entrance as a result of stopped antiproton annihilation. All electrons in this plot have passed the signal selection cuts that we applied to this data. There are a significant number of electrons in the 104-105 MeV/c region of interest for signal conversion electrons.

Choosing Specific Stopping Target Foils

One may wonder if there could be different rates of conversion-like electrons produced depending on where exactly antiprotons annihilate in the stopping target... Do antiprotons that get stopped in the first foil have the same probability to produce a signal-mimicking electron as antiprotons that get stopped in the last foil? How about antiprotons that get stopped on the very outer edge of the stopping target compared to antiprotons that get stopped near the central hole? This section and the next aim to answer these questions, as it is not apparently obvious that all points on the stopping target may result in the same antiproton annihilation behavior.

In `pbarStopTarg_module.cc`, we may now use an explicit foil index to choose which foil we would like the antiprotons to be stopped in instead of firing a random integer in the range of foil indices. First, a sample was generated where all antiprotons were stopped in the first foil of the stopping target. Then, a sample was generated where all antiprotons were stopped in the last foil of the stopping target. We wanted to compare these two samples to determine whether or not antiprotons stopped in the first foil had the same probability to produce a mimic signal electron as antiprotons stopped in the last foil. One may wonder whether or not starting in the first foil and passing through all 34 foils to reach the stopping target has any affect on the momenta of final state electrons from annihilations in the stopping target.

For these samples where antiprotons were generated in different foils of the stopping target or different radii, a total of 2 million antiprotons were generated in the first stage. The positions of the stopped antiprotons was verified by checking histograms of the position of the stopped particles in z and verifying that the positions coincided with the position of the first foil in the stopping target. Looking at the histogram, all of the stopped antiprotons are located at $Z = -4678\text{mm}$. This is the position of the first of 34 foils in the stopping target, exactly where the antiprotons should be stopped. In Figure 5.5, is the momentum distribution for the resulting electrons at the entrance of the tracker. The number of mimic signal electron candidates in the 104-105 MeV/c bin for the first foil sample in Figure 5.5a is 82. The number of mimic signal electron candidates for the last foil sample in Figure 5.5b is 95. These two numbers are reasonably close such that there is no difference in estimated number of signal electrons per stopped antiproton on target from different foils.

Choosing Specific Stopping Target Radii

Similar to investigating the rates from antiprotons stopped in the first foil of the stopping target compared to those stopped in the last foil, we may also check for differences in rates from antiprotons stopped on the very outer radius of the target compared to those stopped on the very inner radius of the target. This process was carried out in a similar fashion to looking at first foil versus last foil, but this time we may use an explicit radius to determine the position of stopped antiprotons instead of randomly firing a radius. The stopped antiproton positions were checked in histograms and the resulting momentum distributions are shown in Figure 5.6. The number of

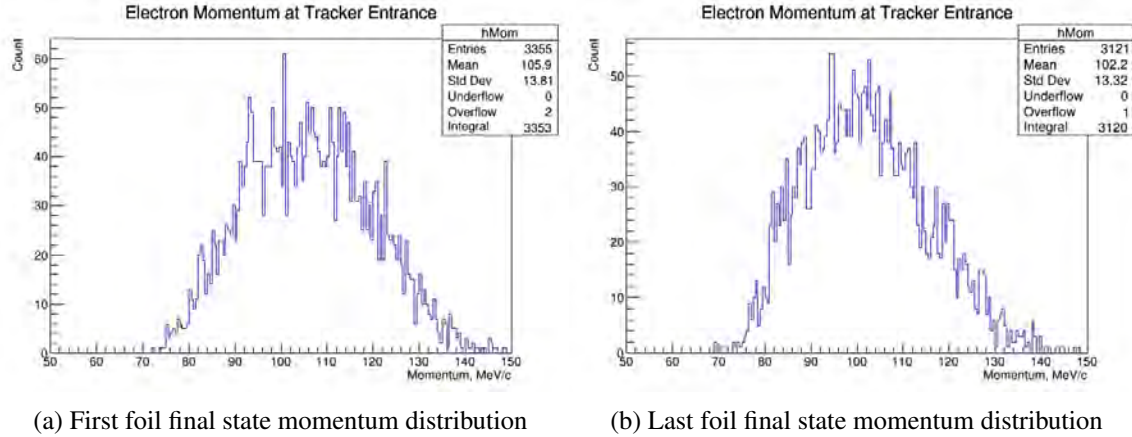


Figure 5.5: Final state momentum distributions in different foils. These figures are momentum distributions of final state electrons at the tracker entrance as a result of stopped antiproton annihilation. All electrons in these plots have passed the signal selection cuts that we applied to this data. Figure a) shows the final state momentum distributions as a result of antiprotons annihilating in the first foil of the stopping target and Figure b) shows the final state momentum distributions as a result of antiprotons annihilating in the last foil.




Figure 5.6: Final state momentum distributions at different radii. These figures show the momentum distribution of final state electrons at the tracker entrance as a result of stopped antiproton annihilation at different radii in the stopping target. All electrons in this plot have passed the signal selection cuts that we applied to this data. Figure a) shows the final state momentum distributions as a result of antiprotons annihilating in the inside edge of the stopping target at a small radius and Figure b) shows the final state momentum distributions as a result of antiprotons annihilating on the outside edge of the stopping target.

mimic signal electron candidates in the 104-105 MeV/c bin for this sample is 106 for the inside edge and 75 for the outside edge. These numbers are also consistent with one another and there is no significant change in estimated numbers of signal electrons per stopped antiproton on target from different radii on the target.

With the momentum distributions presented above, an estimation for the number of mimic signal electrons per antiproton on target may be calculated. The number of mimic signal electrons

Sample	# of pbar	# of signal e-	Estimated e-/pbar
Random Stops	20M	940	4.70×10^{-5}
First Foil	2M	82	4.10×10^{-5}
Last Foil	2M	95	4.75×10^{-5}
Inner Radius	2M	106	5.30×10^{-5}
Outer Radius	2M	75	3.75×10^{-5}

Table 5.1: Estimated stopped antiprotons for different foil positions. Estimations for number of signal electrons per stopped pbar on target for the first set of data samples. The number of electrons reported in the table above is the number of electrons between 104 and 105 MeV/c after the base cuts were applied.

is the numerator of this calculation, while the number of antiprotons generated in the first stage is the denominator. In summary, Table 5.1 reports the estimated number of signal electrons per stopped antiproton on target for the samples shown in this section. These numbers are all consistent with one another and demonstrate that different antiproton stop positions do not have a large affect on the estimated number of signal electrons per antiproton annihilation.

5.5.1.2 Using SU2020 Antiproton Stops as Input

To make this study more robust, we decided to use a simulated stopped antiproton file which was developed for use in the SU2020 sensitivity update campaign and contains stopped antiprotons in the stopping target. The SU2020 campaign was described in Chapter 2 when discussing the Run I sensitivity estimate; refer back to Section 2.3.1 for more details. This set of simulations is called pbar0s41b0.

First Stage FHiCL Design

The FHiCL file from the first part of this study was edited in order to take the SU2020 stopped antiproton file as an input. There was already some infrastructure developed by Pasha Murat [11] and others involved with SU2020 to make these stop files easy to import and use, so the necessary changes were minor.

Second Stage TrkAnaDigisReco

The same second stage FHiCL, TrkDiag/fcl/TrkAnaDigisReco.fcl, was used for evolving the stopped antiprotons out of the stopping target as was used for the first investigations of generating stopped antiprotons on target. This maintains consistency between the generated input and SU2020 input studies.

Choosing Different Physics Lists

Geant4 includes a collection of different reference physics lists which each implement different theoretical models for modeling the interaction of particles in Geant4 simulations. At this

point, we were curious about whether or not using different models would affect the estimation of number of signal electrons per stopped antiproton. Different Geant4 reference models can be easily chosen using a FHiCL parameter by adding a line to the end of the first stage and second stage FHiCL files. The line to add to a FHiCL file to choose a physics list looks like:

```
physics.producers.g4run.physics.physicsListName : "QGSP_BIC"
```

The default model that is used throughout Mu2e Offline is called ShieldingM. ShieldingM is a customized physics list for Mu2e that incorporates details about the Mu2e experimental environment [12]. The other models that were compared to the default Mu2e model are QGSP_BERT [13], QGSP_BIC [14], and FTFP_BERT [15]. Here is a brief description of each model used for comparison:

- Mu2e's default ShieldingM : Mu2e uses a combination of the reference physics lists at low energies to precisely model pion production and hadronic interactions [12]
- QGSP_BERT : this reference physics list applies the quark gluon string model for high energy interactions of protons, neutrons, pions, kaons, and nuclei and then uses the Geant4 Bertini cascade for primary protons, neutrons, pions and kaons below $\sim 10\text{GeV}$ [13]
- QGSP_BIC : similar to QGSP_BERT, this physics list also applies the quark gluon string model for high energy interactions, but for low energy interactions, it uses the Geant4 Binary cascade for primary protons and neutrons with energies below $\sim 10\text{GeV}$ [14]
- FTFP_BERT : this reference physics list applies the FTF model based on the FRITIOF description of string excitation and fragmentation for high energy interactions, then uses the Geant4 Bertini cascade for primary protons, neutrons, pions and kaons below $\sim 10\text{GeV}$ [15]

According to the Geant4 webpages, the best models to use for high energy physics calorimetry are FTFP_BERT [15] and QGSP_BERT [13], while the best models to use for high energy physics tracking are FTFP_BERT [15], QGSP_BERT [13], and QGSP_BIC [14].

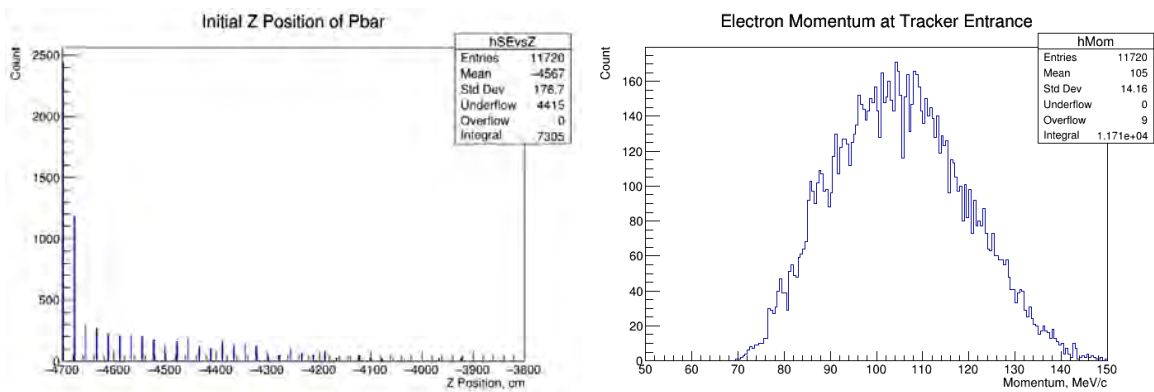
Results

The following plots were made with 10 million antiprotons initially stopped on target, using the default physics list, ShieldingM. For this part of the investigation, the same calculation technique was used as in the previous stage of generating antiprotons at random:

- The first stage FHiCL, pbarstops.fcl, generates .art and .root files which contain the initially stopped antiprotons and information about their distribution

- The second stage FHiCL, TrkAnaDigisReco.fcl, takes the first stage .art files as an input and allows the stopped antiprotons to interact and evolve through the Mu2e virtual detectors, generating .root files for the resulting daughter particles of these antiproton annihilation interactions
- The .root files from TrkAnaDigisReco are used in a ROOT macro, TrkAnaMacro.C, which applies Cut Set C to select events which satisfy criteria for passing event selection
- TrkAnaMacro.C saves histograms of events after passing selection cuts, these are the events which have the potential to be misconstrued as signal; events in the 104 - 105 MeV/c bin of the resulting momentum histogram are taken to be imposter 'signal electrons'
- The number of imposter 'signal electrons' is then divided by the number of antiprotons initially stopped on target, resulting in an estimation of number of signal electrons per stopped antiproton on target

After running the different steps of these simulation jobs, various histograms are produced that describe the simulated events. In Figure 5.7a, the stopped antiproton position can be checked to



(a) ShieldingM stopped antiproton position in z (b) ShieldingM final state momentum distribution

Figure 5.7: ShieldingM stopped antiproton position and final state momentum distribution. Figure (a) shows the stopped antiproton positions along the Z axis in mm with respect to the tracker coordinate system of the parent antiprotons which generated a final state electron. Figure (b) shows the momentum distribution of final state electrons at the tracker entrance as a result of stopped antiproton annihilation. All electrons in this plot have passed the signal selection cuts that we applied to this data. This dataset was processed using the default Mu2e physics list, ShieldingM.

ensure that they are distributed across all foils of the target. Most particles in the beamline stop in the first foil of the target, since particles lose energy as they pass through more foils of the target material. The shape of this distribution matches with investigations of stopped antiprotons from others, namely Giovanni DeFelice [16]. The momentum distribution shows that there are indeed

final state electrons in the range of 104 - 105 MeV/c. These events have the right momentum and PDG ID to be misconstrued as CEs from the stopping target in-time with the beam.

Comparing Different Physics Lists

Now we can take a look at the distributions of final state electron momenta from samples where antiprotons were evolved using different Geant4 physics lists. ShieldingM is largely based on combinations of other standard Geant4 physics models, so we do not expect large differences and this study serves as a check for ShieldingM's performance. Let's look at samples generated using three models that are different than ShieldingM, shown in Figure 5.8. Table 5.2 is a summary table

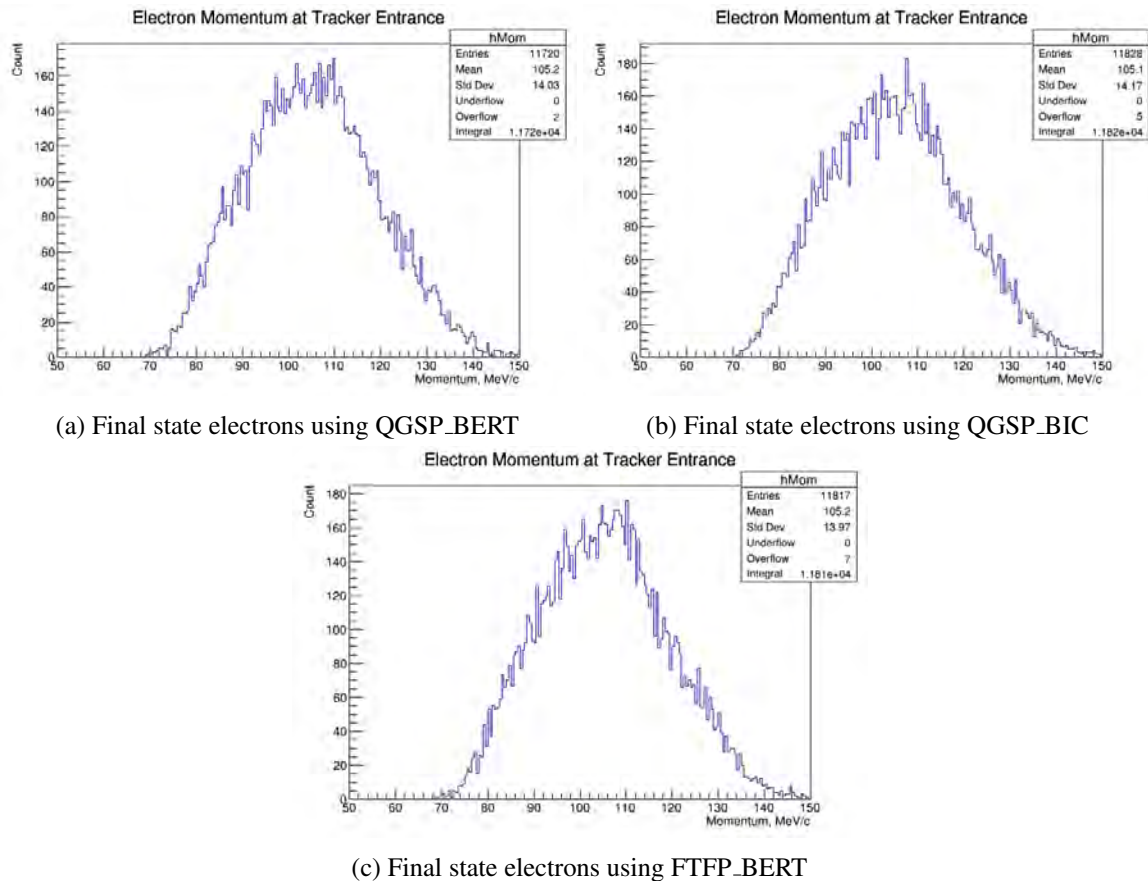


Figure 5.8: Final state electrons using different physics lists. These figures are momentum distributions of final state electrons at the tracker entrance as a result of stopped antiproton annihilation. All electrons in these plots have passed the signal selection cuts that we applied to this data. This dataset was processed using the physics list (a) QGSP_BERT, (b) QGSP_BIC, and (c) FTFP_BERT.

of the results of how many signal electrons per stopped antiproton were simulated in the samples from the four different models. Each sample contained 10 million generated antiprotons on the stopping target.

Sample	# of pbar	# of signal e-	Estimated SE/pbar	Error
ShieldingM	10M	337	3.37×10^{-5}	1.17×10^{-3}
QGSP_BERT	10M	299	2.99×10^{-5}	1.17×10^{-3}
QGSP_BIC	10M	318	3.18×10^{-5}	1.18×10^{-3}
FTFP_BERT	10M	335	3.35×10^{-5}	1.18×10^{-3}

Table 5.2: Estimated stopped antiprotons for different physics lists. Estimations for number of signal electrons per stopped antiproton on target for the set of data samples which use SU2020 stops as input and were run using four different physics lists.

5.5.1.3 Comparing Pion Momentum Spectra between the 1970's and Today

As was briefly described in the introduction to this study, estimating the amount of background from antiproton annihilation has been a source of tension for the Mu2e collaboration for years now. The goal of this part of the investigation is to compare today's Geant4 simulations with real experimental data of antiproton annihilations and what species are created after annihilations occur. If the Mu2e Geant4 simulations of antiproton annihilations are accurate, we would expect that any physics results from Mu2e simulations would closely match experimental data in a similar environment. Geant4 simulations in this study use ShieldingM Mu2e physics model.

Much of this part of the investigation was done with help from Bob Bernstein, who worked diligently to track down historical points of reference for antiproton annihilation data collected by other experiments. Bernstein also generated current distributions using features of Geant4 to mimic the experimental setup from the old data. We also know from Geant4 experts within Mu2e, namely Krzysztof Genser, that different versions of Geant4 may have very different behavior. Historically, Geant4 has been used for antiproton annihilation background estimations for decades, so a large part of this pion momentum spectra investigation also includes looking into older theses and experimental papers which report pion momentum spectra collected from old experiments to compare to the current spectra that we are generating today.

Roy Thesis Data

The only experimental data that is available to compare with in this case is from the 1970s in a thesis by Dr. Jaya Roy [17]. In the thesis, Roy studies data from the Brookhaven 15 ft D2 Bubble Chamber and shows plots of pion distributions from proton-antiproton annihilation on Deuterium in the bubble chamber. When protons and antiprotons annihilate, a majority of the annihilation products are pions. Both neutral and charged pions may be produced in these annihilations. In Mu2e, the pions that are generated as a result of proton-antiproton annihilation could further decay into electrons since they are short-lived or neutral pions can cause backgrounds that are difficult to detect in the tracker. For this reason, it is important to study the direct simulated pion products of annihilations from Geant4 and ensure that the distribution of pions is accurate with real annihilation

results. The Roy thesis data the only real data in the world with information on neutral pion production in antiproton annihilation. While there is no other data on antiproton annihilation available, this data is valuable because it is easy to see when annihilations occur in bubble chamber data because you can count total charge in the bubble chamber lines. Otherwise, it is difficult to study antiproton annihilation. When Geant4 and other models tune their internal parameters for simulating antiproton annihilations, they turn to this data. Papers which validate Geant4 check distributions for the charged pion populations, which are both given in Roy’s thesis. With Geant4 tools, we can simulate pion data using Mu2e Geant4 models that may be directly compared to the Roy thesis data.

For some plots in Roy’s thesis, there is no explanation of how the plots are normalized or what species are being shown; one of the biggest efforts for this study has been to calculate the normalizations of these plots and to discern the specific pion populations that are shown in each distribution. In this plot of $\bar{p}p$ data before being extracted with DataThief, it was clear that




Figure 5.9: Roy thesis extracted annihilations. This plot is data taken from the Roy thesis of $\bar{p}p$ data scanned with DataThief and put into ROOT. A normalization calculation was attempted. The x axis is pion momentum in units of MeV/c and the y axis is the normalized $(1/N_{ann})(dN/dp)$ of charged pions. The uncertainty in chi squared per degree of freedom of this distribution is around 4.

the binning was not consistent, and so we may only extract approximate data from this plot. The normalization of the plot is also not straightforward; the data in Roy’s thesis is only given in plots with raw numbers of events marked, not tables for each data point, and the axes throughout the paper sometimes have inconsistent binning. There is not sufficient information given in the thesis to compute a rigorous normalization. Instead, an approximate normalization was computed.

The plot in the Roy thesis is a plot of $\bar{p}n \rightarrow \pi^- + \text{anything}$. The plot clearly has 20 MeV bins, but there is no given normalization in the thesis. Unfortunately, a G4Study environment could not be replicated to simulate $\bar{p}n$ because the G4Study tool would not allow the generation of a

neutron material.

Current Geant4 Simulations

Now that we have settled some business related to the normalization and specific distributions being shown in Roy's thesis, we may move on to comparing this old, real pion spectrum data from antiproton annihilation in a bubble chamber with Geant4's current simulations of the same type of pions. Using the tool G4Study, we are able to create specific environments to mimic the experimental setup of Roy's bubble chamber. A G4Study was developed with an aluminum target and simulated the annihilation of 40 MeV/c antiprotons. We can make an analogous plot to the one shown in Figure 5.9 using the G4 study tool to compare to Roy's data. Let's look at what this distribution would look like using current G4 simulations below.




Figure 5.10: Geant4 pion spectrum. This plot shows the normalized charged pion spectrum produced using the G4Study tool on deuterium to make a distribution that can be compared to that from Roy's thesis that is shown in Figure 5.9 for real antiproton data on deuterium. The x axis is pion momentum in units of MeV/c and the y axis is the normalized $(1/N_{\text{ann}})(dN/dp)$ of charged pions. The uncertainty in chi squared per degree of freedom of this distribution is around 2.

The plot below in Figure 5.10 shows the current Geant4 normalized pion spectrum in blue and the normalized $\bar{p}p$ data from Roy's thesis in red. If we compare figures 5.9 and 5.10, the plots look like they have approximately the same shape. Each distribution for normalized charged pion momentum shown in figures 5.9 and 5.10 is fit with a fit function. Let's now put both of the points on the same graph to compare as well as just isolate the fits and put them on the same graph to analyze for differences between Roy's data and our current G4 simulations.

The normalizations for these plots are calculated by taking the integral under the curve, which gives you the number of charged pions per annihilation. For the current Geant4 spectrum, the normalization is 2.998(3) charged pions per annihilation. If we look at the Roy data and calculate




Figure 5.11: Roy vs Bernstein pion spectra. A comparison plot of data from Roy's thesis in red and the G4Study simulation in blue. The normalizations from the two plots generally agree, the G4Study normalization is 2.998 while Roy's normalization is 2.7.

the normalization using our scanned points, we find that the plot has a normalization of 2.7 charged pions per annihilation. This normalization is roughly what we would expect, so this is a good sign. The main takeaway for this plot is that the two curves generally agree with each other, especially in the region of interest for Mu2e. We must also consider that there are systematic uncertainties at play here since we have had to infer the bin size and other things from the plot, so it is not unimaginable that these effects have an impact on the shape of the normalized pion momentum spectra. Let's move to just looking at a comparison of the fits in Figure 5.12.




Figure 5.12: Roy vs Bernstein fit comparison. A comparison plot of the fits applied to data from Roy's thesis in red and the G4Study simulation in blue. The fits from the two sets of data generally agree, but diverge at both high and low momentum.

Visually comparing the two fit lines in Figure 5.12 shows that the blue fit line from Geant4 is almost always on top of the orange/red fit line from Roy's data. Taking the log scale of the plot into account, it appears as though the Geant4 fit is roughly 20% larger than the Roy fit, on average. What this really means is that our current simulations of Geant4 overestimate the number of pions that we can expect to be produced from antiproton annihilation. This is good news for Mu2e, and should make the collaboration feel better about the validity of our current simulations suite. Because G4Study could not create a neutron environment for antiproton annihilation, we cannot check to see that Geant4 is simulating $\bar{p}n$ correctly. From the $\bar{p}n$ data given in Roy's thesis, we cannot calculate any neutron cross sections. Thus, we may only make statements about Geant4's ability to simulate specifically $p\bar{p}$ with this study.

Overall, our conclusion is that Geant4 agrees with the Roy data, so we can be confident that it is simulating proton-antiproton annihilation correctly. However, we must be mindful that other Monte Carlo suites, like FLUKA, MARS, and MCNP, produce different results and continue monitoring for changes to Geant4 infrastructure that could alter the simulation of rare background processes.

Summary

In this study, we have analyzed our simulation of antiproton annihilation with Geant4 to make a few conclusions about simulating this type of background in Mu2e. The two conclusions that we can make are 1. about how Mu2e's specific use of Geant4 through our custom physics list compares to default Geant4 simulations and past Geant4 simulations and 2. about how current Geant4 simulations compare to real antiproton annihilation data.

In the first part of this note, we walked through comparing simulations of antiprotons annihilating on the Mu2e muon stopping target. In these studies, we compared both allowing the antiprotons to annihilate at different positions in the stopping target and also using different Geant4 physics lists to see if either would have an effect on our final distribution of final state electrons from the interactions. There was no significant difference observed in simulating antiproton annihilations at different positions in the stopping target. There were also no significant differences observed in simulating antiproton annihilations using different Geant4 physics lists. From this, we can conclude that Geant4 does not have any strange, unexpected dependence on stopping target geometry and also that our custom physics list, ShieldingM, produces reasonable results and is valid to use for our simulations.

In the second part of this note, we took a dive into the only real data for antiproton annihilations in the world and used Geant4 tools to simulate this environment to compare how well Geant4 simulations of antiproton annihilations approximate real data. When it comes to comparing Geant4 with real data from Roy's 1970's dissertation thesis, it appears that current Geant4 simulations roughly agree with the spectra that is published in the thesis using real data. If anything,

Geant4 overestimates the population of pions produced from antiproton annihilation in the same environment as the real data was taken. From this, we conclude that Geant4 is indeed appropriately simulating antiproton annihilation processes, and thus the calculations that have been done within the collaboration using Geant4 simulations are valid. We should also note that since Geant4 appears to overestimate the spectra, that our calculations for antiproton background could be on the upper end of what we should expect from real data-taking.

5.5.2 *Improving Cosmic Trigger Efficiency*

Because cosmic rays are the largest source of background for the Mu2e experiment, it is crucial to efficiently reject cosmic events from being conversion electron candidates. The next two studies involve optimizing the conditions that the Mu2e reconstruction algorithms use to identify and reject cosmic ray events. Accurate reconstruction of cosmic ray events is important for Mu2e because these events can be collected at any time, not just when the experiment has beam. When the experiment is not actively operating and rejecting cosmic events, these cosmic rays can be used for spatial calibration and momentum resolution measurements, as well as for characterization during off-spill data.

As we discussed earlier in this chapter, there are two pattern recognition algorithms within the Mu2e code base that search for particle helices within each event: calorimeter-based CPR and tracker-based TPR. For simulated CE events, both algorithms perform very well and have been fine-tuned to search for these CE events. Both algorithms have similar performance for simulated samples of purely CE events. However, during simulations of non-signal events, like cosmic rays, it was observed that the CPR and TPR algorithms can produce very different results. Ultimately, the TPR algorithm should be able to accurately reconstruct all charged particles that pass through the tracker, since TPR does not explicitly need calorimeter data for event reconstruction.

This study began by inspecting two simulations of cosmic ray-induced events, one run using only the CPR algorithm and the other using only the TPR algorithm for reconstruction. For clarification, cosmic ray-induced events are any events which started as cosmic rays that interacted and produced descendants in the detector volume that are reconstructed by the tracker. Both simulations contained the same number of initial cosmic events, but the output of reconstruction resulted in many more final state events being reconstructed by the CPR algorithm than the TPR algorithm. The large difference in the number of events reconstructed by the two pattern recognition algorithms was identified as a problem. This investigation was done to recover the events that were not reconstructed by TPR by improving the performance of the TPR algorithm. The difference between the number of events in each simulation can be seen when comparing the momentum distribution of events from the CPR dataset to that of the TPR dataset.

Comparing the two distributions in Figure 5.13 shows the difference between the outputs of




Figure 5.13: Initial reconstructed helix momentum distributions comparing CPR and TPR. Figure (a) shows the reconstructed helix momentum distribution for final state events in the simulation using only the CPR reconstruction algorithm [18]. On the Y-axis is event counts, this is a histogram. Figure (b) shows the same reconstructed helix momentum distribution, with event counts on the Y-axis, but for events reconstructed with only the TPR algorithm [18]. The different colored lines in each plot indicate different final state particle types: the solid blue lines represent all reconstructed events in each set, the pink lines represent events reconstructed as antimuons, or positive muons MuPlus, and the green lines represent events reconstructed as muons which are negatively charged, or MuMinus.

running the two reconstruction algorithms independently. The difference in the number of events between the CPR distribution and the TPR distribution is 362,403 events. My goal was to determine why TPR was reconstructing fewer cosmic events than CPR and then try to tune the TPR algorithm such that it is able to reconstruct the missing events. Ideally, the two reconstruction algorithms should have comparable independent performance for all event types, as is the case with CEs. Of course, the best course of action is to use both algorithms in conjunction for full reconstruction, but this study allows us to find and fix the shortcomings that one algorithm may have that the other does not.

Another distribution that can also give clues to this investigation is the momentum resolution of each algorithm, or how close each algorithm gets to the true simulated momentum value for each event. Using simulated data here allows us to know what the true values of the reconstructed helix momentum should be. Ideally, events should be reconstructed to perfectly match the simulated values, and thus the momentum resolution should be small and symmetric about zero. Symmetry about zero indicates that the algorithm is free of systematic errors that would cause consistent misreconstruction in one direction.

The momentum resolution distributions shown in Figure 5.14 look very different depending on which reconstruction algorithm was used on the events. The momentum resolution distribution for the CPR algorithm is relatively symmetric about 0 and has a narrow width. This indicates that the CPR algorithm gets very close to reconstructing the true momentum of each event without any systematic issues. On the other hand, the momentum resolution distribution for the TPR




Figure 5.14: Initial momentum resolution distributions comparing CPR and TPR. Figure (a) shows the momentum resolution distribution for final state events in the simulation using only the CPR reconstruction algorithm [18]. On the Y-axis is event counts, this is a histogram. The momentum resolution is calculated as the difference between the reconstructed helix momentum and the Monte Carlo truth momentum. Figure (b) shows the same momentum resolution distribution, with event counts on the Y-axis, but for events reconstructed with only the TPR algorithm [18]. Again, the different colored lines in each plot indicate different final state particle types: the solid blue lines represent all reconstructed events in each set, the pink lines represent events reconstructed as antimuons, or positive muons MuPlus, and the green lines represent events reconstructed as muons which are negatively charged, or MuMinus.

algorithm shows just the opposite: the TPR momentum resolution distribution is drastically skewed to the left. This indicates that something within the TPR algorithm is resulting in a consistent misreconstruction of events to a lower momentum than they should have. This has been an artifact of the TPR algorithm for some time in different event types and is an area of active improvement for the collaboration [18].

Next, the subset of 362,403 events that were reconstructed by CPR and not by TPR were isolated into one .art file. In this way, tools like event print statements and event displays could be used to deeply investigate the different data products that these events contained. When printing the data products in an event, the product names and different information associated with a given product is displayed in a table. Using a command line allows all event information to be accessed and printed. Another tool that was used during this investigation is the debugging software called GDB [19]. GDB is the GNU [20] standard text-only debugger that allows a user to set break points in a code which the user to print different code parameters and objects to get a deeper understanding of how the code is behaving on an event-by-event basis. For example, here we can use GDB and print out information after each step of reconstruction occurs, such as whether or not the necessary data products are created to progress to the next step of the algorithm. Upon further analysis of these events that were reconstructed by CPR but not by TPR, it was discovered that the events had a sufficient number of hits in the tracker to allow for track reconstruction, but the hits in the tracker were not grouped into triplets as the TPR algorithm requires. The absence of hit triplets prevented TPR from building a track for these events. Since the CPR algorithm does not use tracker hit triplets,

it makes sense that these events were reconstructed by CPR.

After the observation that the events in question had no tracker hit triplets, the parts of the TPR algorithm that place conditions on triplet grouping were under speculation. The piece of the algorithm that uses triplets is the piece of the algorithm that forms the transverse circular profile of the track, called CircleFit. When TPR attempts to form a triplet, the group of three hits must satisfy certain conditions to become a triplet. Some of these conditions involve restrictions on the radius and starting point of the reconstructed 2D helix shape and the distance between any two points in a triplet. Using GDB, it was determined that the conditions that were causing triplets to fail formation were parameters involving the distance between two points in a triplet, the initial starting point of the reconstructed helix coinciding with the stopping target, and the allowable radii that a reconstructed helix may have. Next, the values of these parameters that trigger triplet rejection were changed to try and increase TPR acceptance for these events. The names of each parameter that were found to be violated within TPR, its definition, changes in value, and the effect of each change are listed below.

- targetradius: radius of the stopping target
 - increased from 100mm to 600mm
 - effectively removes the requirement that cosmic events begin at the stopping target
- maxR: maximum radius that the reconstructed helix cross section can have
 - increased from 320mm to 600mm
 - allows any size of reconstructed helix for cosmic events
- maxdist: upper limit of allowed distance between two points in a hit triplet
 - increased from 500mm to 1000mm
 - allows for larger distances between any two points in hit triplets, increasing triplet acceptance
- mindist: lower limit of allowed distance between two points in a hit triplet
 - decreased from 100mm to 50mm
 - allows for smaller distances between any two points in hit triplets, increasing triplet acceptance

The initial values of each of these parameters were optimized for CE event reconstruction. Within CE trigger paths, many assumptions about the shape of the track and initial position of the particle

CPR	362,403 events	100%
TPR targetradius	31,750 events	8.8%
TPR targetradius + maxR	95,834 events	26.4%
TPR targetradius + maxR + maxdist	112,720 events	31.1%
TPR targetradius + maxR + maxdist + mindist	115,656 events	31.9%

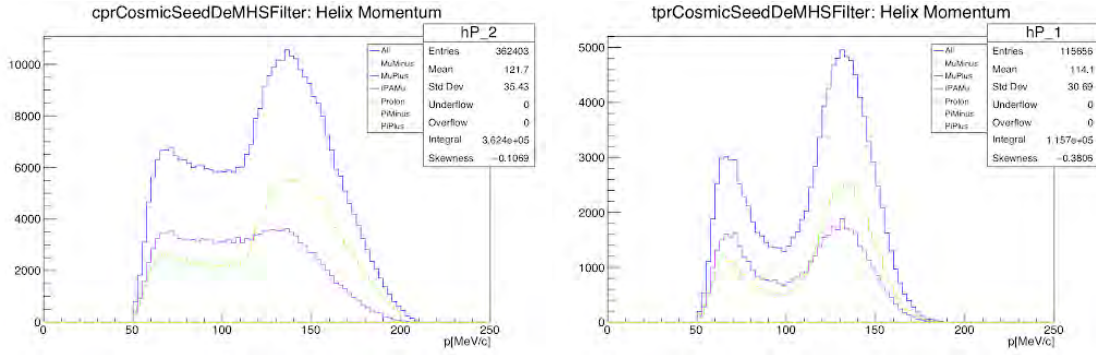
Table 5.3: TPR improvements from CircleFit changes. Improvements in cosmic-induced reconstruction as a result of changing the value of each of the listed CircleFit parameters within the TPR algorithm. The 362,403 events that were initially reconstructed by CPR, but not by TPR were taken to be 100% of the events that we desire to reconstruct as comsics [21].

are made since any signal CEs will come from the muon beam and originate from the stopping target with a specific momentum. The monoenergetic nature of CEs implies that CE event tracks have a narrow range of allowable radii and initial positions. Instead, the parameters found in this study indicate that the same conditions should not be placed on cosmic events. Cosmic events have a wide range of initial positions from the atmosphere, as well as a wide initial momentum range. Cosmic tracks may be any shape and size within the DS depending on the initial properties of the incoming cosmic particle. The changes made to the parameters listed above reflect the wider range of possible track shapes that cosmic-induced events may have and can be incorporated into cosmic trigger paths to increase acceptance for these events without affecting acceptance for CEs.

To measure the increase in cosmic acceptance as a result of changing the CircleFit parameters listed above, the simulation was relaunched between changing each of the parameter values. The cumulative result of changing the four parameters is given in Table 5.3, along with the total number and fraction of the total number of CPR-only events that were recovered. After applying all four of these changes to the CircleFit parameters, reconstructed momentum distributions were made for the subset of events that were initially reconstructed by CPR, but not by TPR. These distributions are shown in Figure 5.15 and give a clue about what kind of events are still missing from the TPR reconstruction.


The CPR distribution contains the full subset of 362,403 events, whereas the TPR distribution contains 115,656 events. At this point in the investigation, the TPR algorithm is now able to reconstruct about 1/3 of the events that were previously missing from this subset. The shapes of the distributions indicate that the TPR algorithm is not accurately reconstructing events with a momentum between 80 and 120 MeV/c or events with a momentum above 180 MeV/c; this is where the TPR distribution has fewer events than the CPR distribution. The momentum resolution may also be inspected for these events to look for whether or not the TPR momentum resolution distribution is still skewed.

Unfortunately, the distribution is still skewed to the left in roughly the same way as the initial momentum resolution distribution in Figure 5.14. Another distribution that may give clues



(a) Momentum distribution for CPR-reconstructed events (b) Momentum distribution for TPR-reconstructed events

Figure 5.15: Final state reconstructed helix momentum distributions comparing CPR and TPR. Figure (a) shows the reconstructed helix momentum distribution for final state events in the simulation using only the CPR reconstruction algorithm. On the Y-axis is event counts, this is a histogram. Figure (b) shows the same reconstructed helix momentum distribution, with event counts on the Y-axis, but for events reconstructed with only the TPR algorithm. The different colored lines in each plot indicate different final state particle types: the solid blue lines represent all reconstructed events in each set, the pink lines represent events reconstructed as antimuons, or positive muons MuPlus, and the green lines represent events reconstructed as muons which are negatively charged, or MuMinus.



(a) Momentum resolution distribution for CPR-reconstructed events (b) Momentum resolution distribution for TPR-reconstructed events

Figure 5.16: Momentum resolution distributions comparing CPR and TPR. Figure (a) shows the momentum resolution distribution for final state events in the simulation using only the CPR reconstruction algorithm. On the Y-axis is event counts, this is a histogram. The momentum resolution is calculated as the difference between the reconstructed helix momentum and the Monte Carlo truth momentum. Figure (b) shows the same momentum resolution distribution, with event counts on the Y-axis, but for events reconstructed with only the TPR algorithm. Again, the different colored lines in each plot indicate different final state particle types: the solid blue lines represent all reconstructed events in each set, the pink lines represent events reconstructed as antimuons, or positive muons MuPlus, and the green lines represent events reconstructed as muons which are negatively charged, or MuMinus.

about what kind of events are still missing is the reconstructed helix impact parameter. In Mu2e, a positive reconstructed helix impact parameter indicates that the track does not encompass the

detector center and a negative impact parameter indicates that the track does encompass the detector center.




Figure 5.17: Reconstructed helix impact parameter distributions comparing CPR and TPR. Figure (a) shows the reconstructed helix impact parameter distribution for final state events in the simulation using only the CPR reconstruction algorithm. On the Y-axis is event counts, this is a histogram. Figure (b) shows the same reconstructed helix impact parameter distribution, with event counts on the Y-axis, but for events reconstructed with only the TPR algorithm. Again, the different colored lines in each plot indicate different final state particle types: the solid blue lines represent all reconstructed events in each set, the pink lines represent events reconstructed as antimuons, or positive muons MuPlus, and the green lines represent events reconstructed as muons which are negatively charged, or MuMinus.

Comparing the two distributions in Figure 5.17 shows that the two distributions have very different shapes. In the CPR sample on the left, events have a wide range of impact parameters, from -400mm to +500mm. The CPR distribution is also relatively flat, with events covering all possible impact parameters. Conversely, the TPR sample on the right shows a large dip around 100mm, far fewer events were reconstructed with positive impact parameters than negative impact parameters, and the extreme ends of the distribution are also missing. The absence of events with positive impact parameter in the TPR distribution indicates that this algorithm is not reconstructing events that circle around the stopping target.

The reconstructed helix impact parameter is not only determined by the transverse CircleFit pieces of the TPR algorithm, but is also affected by the shape of the helix along the beamline. This piece of the reconstruction algorithm fits the angular ϕ helical rotation axis to the Z axis along the beamline and is called the FZFit. ϕ is the angle of the signal particle, measured with respect to the helix center and Z is the distance along the beamline. The two parameters that describe the dynamics of this fit are the initial angle, ϕ_0 , and λ , which is defined as $\lambda = \frac{dz}{d\phi}$. As a particle moves forward along the beamline, in Z, its rotation angle, ϕ , loops from 0 to 2π as the particle rotates around the helix center. λ is the slope between ϕ and Z of one revolution of the particle about the helix center. This parameter allows for correction of the 2π ambiguity and can be used for calculation of how many turns a particle takes in a reconstructed helix. The distributions of λ for the

events reconstructed by CPR and also by TPR after the CircleFit modifications may be compared in Figure 5.18.

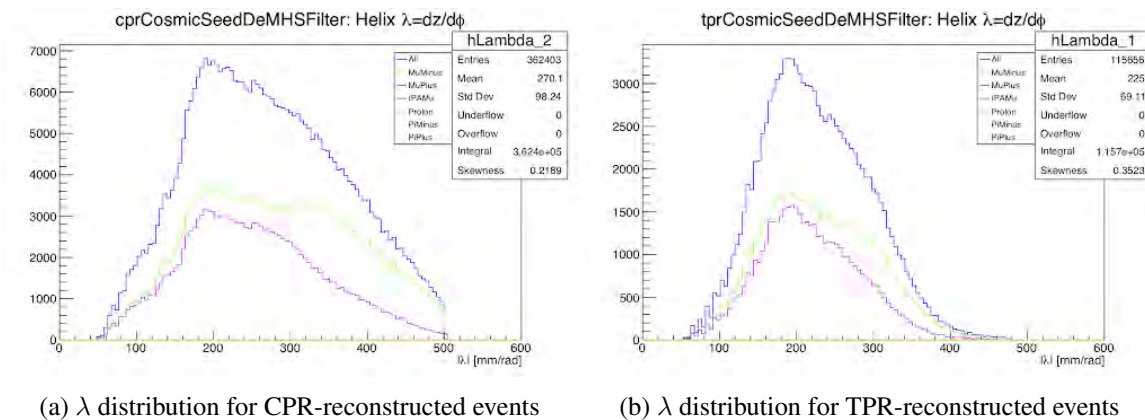


Figure 5.18: λ distributions comparing CPR and TPR. Figure (a) shows the λ distribution for final state events in the simulation using only the CPR reconstruction algorithm. On the Y-axis is event counts, this is a histogram. Figure (b) shows the same λ distribution, with event counts on the Y-axis, but for events reconstructed with only the TPR algorithm. Again, the different colored lines in each plot indicate different final state particle types: the solid blue lines represent all reconstructed events in each set, the pink lines represent events reconstructed as antimuons, or positive muons MuPlus, and the green lines represent events reconstructed as muons which are negatively charged, or MuMinus.

Comparing the two lambda distributions in Figure 5.18 reveals that a wide range of events are missing from the high side of the TPR λ distribution. This observation led to the idea that some events were being rejected based on the TPR algorithm's acceptable range of λ . When λ is calculated within TPR, the algorithm chooses pairs of hits and uses the difference in the ϕ ($d\phi$) and Z (dz) coordinates of the hits to calculate a lambda value. If $d\phi$ and dz were within an acceptable range defined in the code, the resulting lambda is acceptable and is filled into a histogram like the ones seen in Figure 5.18. Using GDB, it was discovered that within the subset of events reconstructed by CPR, but not by TPR, high momentum events with large helices were not using many hit pairs to calculate λ because ϕ changed slowly around the large track. This violated the condition on minimum $d\phi$ value to pass as an acceptable hit pair. In some cases, the dz value was large enough to also violate a condition on maximum allowed dz . In addition, the allowed maximum $d\phi$ was set to an arbitrarily high value, causing calculations for λ in some events to result in erroneously high values as a result of a large $d\phi$ in the denominator. Based on these observations, the allowed maximum dz was increased and the allowed maximum $d\phi$ was decreased to a reasonable value consistent with the limit set by the minimum allowed dz and minimum allowed λ value. After making these changes, an additional 130,065 events were reconstructed by TPR.

Figure 5.19 shows the final λ distribution for the events reconstructed by TPR. This distribution has many more events at the high λ end that were missing before the FZFit modifications.




Figure 5.19: λ distribution for TPR after ϕZ modifications. This figure shows the λ distribution for final state events in the simulation using only the TPR reconstruction algorithm. On the Y-axis is event counts, this is a histogram. Again, the different colored lines in each plot indicate different final state particle types: the solid blue lines represent all reconstructed events in each set, the pink lines represent events reconstructed as antimuons, or positive muons MuPlus, and the green lines represent events reconstructed as muons which are negatively charged, or MuMinus.

CPR	362,403 events	100%
TPR + CircleFit	115,656 events	31.9%
TPR + CircleFit + maxdphi	208,597 events	57.6%
TPR + CircleFit + maxdphi + maxzsep	245,721 events	67.8%

Table 5.4: TPR improvements from FZFit changes. Improvements in cosmic-induced reconstruction as a result of changing the value of each of the listed FZFit parameters within the TPR algorithm. The 362,403 events that were initially reconstructed by CPR, but not by TPR were taken to be 100% of the events that we desire to reconstruct as comsics [22].

The improvements can again be cumulatively be listed in Table 5.4. In the end, the CircleFit and FZFit modifications enabled the TPR algorithm to reconstruct about $\frac{2}{3}$ of the events that were initially reconstructed by CPR, but not by TPR. The momentum distribution for the events reconstructed by TPR after these modifications is shown in Figure 5.20.

The momentum distribution in Figure 5.20 shows that the FZFit modifications allowed TPR to reconstruct many more events than with just the CircleFit modifications. Comparing Figure 5.20 to the TPR distribution in Figure 5.15 reveals that many of the events that are able to be reconstructed after the FZFit modifications fall into the 80 - 120 MeV/ c^2 momentum range, which is our signal momentum range of interest.

To conclude this study, the last metric that can be used to evaluate the benefits of any parameter changes is efficiency. Specifically, in this case, we want to be sure that the modifications for increasing acceptance of cosmic events has no adverse reactions in acceptance of signal CE events. To evaluate efficiency, two more simulations were launched, both simulations including a mix of CE events and various backgrounds, including cosmic-induced events. The two simulations were processed with the default reconstruction algorithms before the modifications to CircleFit




Figure 5.20: Reconstructed helix momentum distribution for TPR after ϕZ modifications. This figure shows the reconstructed helix momentum distribution for final state events in the simulation using only the TPR reconstruction algorithm. On the Y-axis is event counts, this is a histogram. Again, the different colored lines in each plot indicate different final state particle types: the solid blue lines represent all reconstructed events in each set, the pink lines represent events reconstructed as antimuons, or positive muons MuPlus, and the green lines represent events reconstructed as muons which are negatively charged, or MuMinus.

	Default Algorithm	Modified Algorithm
Passed Events	10980/12798	11913/12798
Efficiency	85.79%	93.08%


Table 5.5: Overall TPR algorithm improvements in CEMix. Improvements in cosmic-induced reconstruction as a result of changing CircleFit and FZFit parameters within the TPR algorithm for the CEMix simulation sample. This simulation contains CE events as well as various backgrounds mixed-in [23].

and FZFit parameters and also processed with the modified parameters that we have discussed above. By counting the number of reconstructed events and comparing the number of reconstructed accepted events to the number of events within the simulation, the efficiency of each pathway may be evaluated. The efficiencies for the default and modified algorithms are given for the CEMix sample in Table 5.5 and for the CEPlusMix sample in Table 5.6.

The results in Tables 5.5 and 5.6 show that these modifications have a net positive effect on the efficiency of reconstructing samples containing CE events, a positive outcome of this endeavor. One more distribution that can be inspected for positive effects is the momentum resolutions distribution. Recall that the momentum resolution distribution for the TPR algorithm was skewed

	Default Algorithm	Modified Algorithm
Passed Events	5977/7641	6334/7641
Efficiency	80.11%	84.89%

Table 5.6: Overall TPR algorithm improvements in CEPlusMix. Improvements in cosmic-induced reconstruction as a result of changing CircleFit and FZFit parameters within the TPR algorithm for the CEPlusMix simulation sample. This simulation contains CE events as well as various backgrounds mixed-in, with higher momentum events than the CEMix simulations [23].



(a) Momentum resolution distribution for reference TPR events (b) Momentum resolution distribution for modified TPR events

Figure 5.21: Momentum resolution distributions comparing TPR before and after changes. Figure (a) shows the momentum resolution distribution for final state events in the simulation using only the reference TPR algorithm before any changes were made to CircleFit and FZFit algorithms. On the Y-axis is event counts, this is a histogram. The momentum resolution is calculated as the difference between the reconstructed helix momentum and the Monte Carlo truth momentum. Figure (b) shows the same momentum resolution distribution, with event counts on the Y-axis, but for events reconstructed with only the TPR algorithm after applying the CircleFit and FZFit changes. Again, the different colored lines in each plot indicate different final state particle types.

to the left at the beginning of this investigation.

Comparing the two distributions in Figure 5.21, with the left-hand distribution showing the reference TPR before CircleFit and FZFit modifications and the right-hand side showing the TPR algorithm after the changes were applied, it is clear that the changes to the parameters of interest improve momentum resolution in the TPR algorithm. The momentum resolution distribution is less skewed after the CircleFit and FZFit modifications, evident by the decrease in the mean and visual inspection. Overall, these improvements to the TPR algorithm resulted in increased acceptance for reconstructing cosmic-induced events with TPR alone, better momentum resolution for the TPR algorithm, and increased efficiency for reconstruction within mixed samples containing CEs.

5.5.3 Optimizing the CRV Dead-Time Window

The last study involves the vetoing of cosmic events that have a signal in the CRV and optimizing the length of the CRV veto window, or dead-time window. As we discussed earlier in this chapter, if an event passes any trigger path, the CRV event data is aggregated into the event building process. When a signal is present in the CRV, a dead-time window is applied to data from all detectors at that time to ensure that no cosmic particle will be mistaken for a conversion electron event from the beam. This method of mitigation for cosmic rays works well, but since applying the veto window prevents recording any data that comes in simultaneously with a cosmic ray, valuable data from beam muons which decay from the stopping target may be lost. Therefore, it is important

to tune our dead-time window to the shortest length possible while ensuring that all residuals from cosmic ray particles are rejected. Improvement of the timing of cosmic ray detection in the CRV and the tracker allows for accurately detect cosmic ray events as fast as possible and optimization of the dead-time window applied upon detection of cosmic ray particles.

The main part of this study will discuss rejection of cosmic events based on the number of time clusters corresponding to a given event. In Mu2e, each reconstructed helical track that is left by a particle traveling through the DS generates a spatial track in the tracker which is associated with a temporal time cluster. Thus, the presence of a time cluster suggests that there is a corresponding helix which has been reconstructed in the tracker for that event. Events with multiple time clusters have multiple tracks, and therefore crossed the tracker multiple times [2].

In the simulated dataset used in this investigation, each event is generated from exactly one cosmic particle entering the detector solenoid, by design. Since we know that there is only exactly one cosmic particle per event, events in this investigation which are associated with more than one time cluster are unambiguously events where the same particle crossed the tracker more than once. For example, an event with two time clusters in this dataset is most likely the result of a cosmic particle first entering the DS traveling in an upstream direction and producing an upstream-moving track in the tracker. Then, the particle travels further up the DS into the graded region of our magnetic field at the TS5/DS junction, and is then reflected back downstream towards the calorimeter, leaving a second downstream track through the tracker in the process. This type of event is referred to as a 'bouncing cosmic'. This characteristic of particles with two time clusters identifying the bouncing events with an upstream component allows for a time gate to be employed for the purpose of removing these bouncing cosmics.

During the Sensitivity Update 2020 (SU2020) campaign to determine the Run I sensitivity reach and background mode estimates [6], the Mu2e offline code repository was forked and an SU2020 branch of the offline repository was created specifically for Run I studies. This fork allowed for faster evolution of development for SU2020 to optimize the Run I estimates as well as the ability to use the code for future comparisons with the main branch. The dataset used for this investigation is called the cry31s91b0 cosmic dataset in the SU2020 repository. The name refers to the generator that was used to produce the cosmic ray particles, 'cry', and different versions of the simulation generation.

Recall how the CRV applies a veto in the presence of a signal. If a particle leaves a detectable signal in the CRV, the event must fulfill specific criteria to be reconstructed as a cosmic ray: the signal must be detected in at least three of the four active layers of the CRV dicounters, the three 'hit strips' of active counters in different layers must be adjacent, and the signals in these 'hit strips' must all occur within a 5ns time window. If all three of these conditions occur, a CRV 'stub' is produced, the event is reconstructed as a cosmic ray muon, and the dead-time window is triggered

to ensure that the signal does not mimic that of a conversion electron. The current dead-time window employed by the CRV to reject cosmic events upon the detection of a cosmic in the CRV is 125ns [10].

Cosmic ray particles which have a CRV stub produce events with the following data products:

- **CRV Digis:** the raw ADC outputs of the CRV SiPMs, in ADC counts. These outputs may be plotted as waveforms, but not all digis for an event are typically useful. Collections of digis for an event normally contain noisy waveforms, which are cut out before the next step. The digitization period of the CRV is 12.55ns, so digi waveform resolution is limited by this bin size;
- **CRV Reco Pulses:** the 'good' waveforms from the CRV digis. Reco pulse waveforms have a measurable peak and are ordered by time. There are two times associated with each reco pulse: one labeled simply 'time' which corresponds to the time of the waveform peak, and one labeled as 'LE time' which corresponds to the leading edge time of the waveform, taken at $1/e$ on the rising edge of the waveform;
- **CRV Coincidence Clusters:** the collections of at least 3 CRV counters where cosmic ray particles were detected coincident in space and time. We are now calling these type of objects 'CRV stubs', as in the stub of a cosmic ray track through the CRV. CRV coincidence clusters are associated with a position and a time: the time of the coincidence cluster is taken as the earliest reco pulse time and the position is the central position of the relevant CRV counter;

After covering the basic data products that are part of cosmic ray events in the CRV, we can break down the timing of cosmic events in two sections: the timing of pulses within the CRV bars and the timing of CRV stubs with respect to the timing of the tracks that the cosmic particle makes in the tracker. At the end of this section, we will discuss a particularly relevant distribution which is affected by both the timing within the CRV bars and the timing between the CRV and the tracker.

5.5.3.1 Timing of Pulses within CRV Bars

Prior to this investigation, signals within the CRV bars were simply collected, fit, and stamped with their leading time and peak time. When printing the timing data within an event, any pulses that are detected are ordered only by their peak time, but they carry with them information about the counter and specifically which SiPM on the dicounter manifold from which they were collected. The bar and SiPM labels are valuable information for determining initial position and timing of the particle as it interacted with the CRV. Before this work, the information about which bar and SiPM each pulse was collected from was unused in the context of timing, and physical propagation times were not taken into account when determining the timing of the CRV signals.

5.5.3.2 Timing of CRV Stubs vs. Tracker Helices

CRV stubs and tracker helices have timing structures which have been developed within their own independent spheres, but evaluating the time of tracker objects with respect to CRV objects will aid in building the most robust cosmic rejection system. After a cosmic signal is detected in the CRV, some of the tracker and calorimeter data is lost since the event does not proceed through the full processing chain, but the timing of CRV objects to tracker objects is still available and may be used for rejection.

Event displays were very helpful in starting development of this work to illustrate how cosmic particles traverse the DS. When using an event display to view some of the cosmic events in the cry31s91b0 dataset which was used throughout these improvements, one feature was initially hard to ignore: in many events, the CRV stub signal in the CRV bar appears to have been detected *after* the cosmic particle had already passed through the CRV and made a helix in the tracker. This scenario is unrealistic and unphysical, of course, and is a result of missing key corrections to CRV timing which will be discussed later in this study.

5.5.3.3 dT vs. $Z(\text{CRV})$

One distribution that showcases areas to improve upon the reconstruction algorithm is a plot of the difference in time between a cosmic ray particle's track in the tracker and its CRV stub versus the position in Z of the CRV bar where the cosmic crossed the CRV. This distribution is shown below for the cry31s91b0 SU2020 dataset [11].





Figure 5.22: Time difference between track and CRV stub versus Z. This distribution shows the difference in time between the timing of the track in the tracker and the CRV stub in the CRV versus the position along the beamline in Z of the CRV stub. The events with Z around 5000mm enter near the stopping target and events near 15000mm enter near the calorimeter.

The width of the dT versus $Z(\text{CRV})$ distribution along the Y axis determines the necessary length of the CRV dead-time window which we must apply upon detection of cosemics in the CRV.

Any data from the tracker and calorimeter which may contain the traces from a cosmic particle, regardless of where the cosmic originates from, may mimic a CE signal. So, the timing of a CRV stub relative to a tracker helix is very important. If we know that CRV stubs precede their tracker helix by 50ns or less, then we may only need to apply a dead-time window of 50ns for each cosmic event. In Figure 5.22, it was discovered that the two 'legs' of this distribution signify two distinct groups of events: cosmics which enter the CRV traveling downstream in the same direction as the beamline and cosmics which enter the CRV traveling upstream in the opposite direction of the beamline.


Further analysis into the time cluster data of the events in the cry31s91b0 dataset revealed that there are two classes of events within the set: events with only one time cluster and events with more than one time cluster. In addition, some of the events with more than one time cluster contained an unexpectedly high amount of time clusters. Some events were found to have eight or more time clusters, which is unphysical considering the environment inside the Mu2e DS. Upon deeper investigation, it was discovered that some of the time clusters in the events with high numbers of time clusters were being duplicated. To avoid time cluster duplication, a duplication checking algorithm was developed to count the number of 'effective' time clusters, the clusters which did not have duplicates. Figures 5.23a and 5.23b show time cluster distributions before and after applying effective time clustering to eliminate duplicates.



(a) Number of default time clusters in cry31s91b0 (b) Number of effective time clusters in cry31s91b0

Figure 5.23: Results of effective time clustering on cry31s91b0. Figure (a) shows the number of time clusters on the X axis with number of events on the Y axis in a log scale. This is the default number of time clusters that was associated with each event before discovering the duplicated time clusters. Figure (b) shows the number of effective time clusters on the X axis with the same number of events on the Y axis in a log scale. Effective time clusters condense any duplicated time clusters.

After creating effective time clusters, the dT vs. Z(CRV) distribution may be split according to how many effective time clusters each event contains. Figures 5.23a and 5.23b show the dT vs. Z(CRV) distribution for events with only one effective time cluster (Figure 5.23a) and for events with more than one effective time cluster (Figure 5.23b).



(a) dT vs. $Z(\text{CRV})$ for events with 1 effective time cluster
 (b) dT vs. $Z(\text{CRV})$ for events with ≥ 1 effective time cluster

Figure 5.24: dT vs. $Z(\text{CRV})$ for events with different numbers of effective time clusters. Figure (a) shows the dT vs. $Z(\text{CRV})$ distribution for events with only one effective time cluster. While events with only one time cluster span both legs of the dT vs. $Z(\text{CRV})$ plot, these events usually have a relatively small dT . Figure (b) shows the dT vs. $Z(\text{CRV})$ distribution for events with more than one effective time cluster. Events with more than one effective time cluster are all located in the upper leg of the dT vs. $Z(\text{CRV})$ distribution with high dT times. These are events with distinct upstream and downstream tracks and start at high Z .

After separating the dT vs. $Z(\text{CRV})$ distribution by the number of effective time clusters in each event, it becomes more apparent what type of events are in each leg of the distribution. Events with more than one time cluster are all in the upper leg of the dT vs. $Z(\text{CRV})$ plot because these are events that truly have upstream and downstream track components. These events enter the DS near the calorimeter, evident by the high $Z(\text{CRV})$ position, and most likely have an initial upstream traveling trajectory, generating an upstream track with an early time cluster. When the particle travels to the upstream side of the DS, it is reflected by the graded magnetic field in that region, and travels back downstream, generating a downstream track and a second, late time cluster. The majority of the events with only one effective time cluster are in the central region of the dT vs. $Z(\text{CRV})$ plot, but some of these events span into the upper leg of the distribution. These events with only one effective time cluster are events which only cross the DS once. Since a majority of the events in this category have a relatively low $Z(\text{CRV})$ that corresponds with the upstream end of the DS, these events primarily travel in the downstream direction and only make one track in the DS.

5.5.3.4 Improvements to CRV Stub Timing

There are two main improvements that have been recently added to the su2020/STNTUPLE infrastructure which build on the timing scheme that we detailed above and allow for us to more accurately depict timing of cosmic ray events. These two main improvements are the addition of signal propagation time within CRV bars and accounting for time of flight through space between the CRV and tracker detectors.

In scintillators, the speed of light is limited by the index of refraction of the scintillation

material which light is traveling through. This holds true for the scintillating bars of the CRV; as cosmics enter the CRV counters, the particles produce light within the counters which must travel to the ends of each bar where the SiPMs are located. So, if a cosmic particle crosses through a CRV counter near the center along the length of the counter, there is a nonzero propagation time that the light from the cosmic signal takes to travel to the ends of the counter and reach the SiPMs. The timing from both ends of the counter can be used to solve for the time and location along the counter at which the cosmic particle passed through the counter. The specific system of equations to solve that determines the corrected time and position along the counter where the cosmic was detected is included in Appendix A.

The results of applying this correction is shown in a corrected dT vs. $Z(\text{CRV})$ plot, Figure 5.25. If we compare the distribution in Figure 5.25 with Figure 5.22, it is clear that this correction mainly affects the bottom leg of the dT vs. $Z(\text{CRV})$ distribution.




Figure 5.25: Propagation corrected time difference between track and CRV stub versus Z . This distribution shows the difference in time between the timing of the track in the tracker and the CRV stub in the CRV after the CRV stub timing has been corrected from propagation time in the CRV scintillator versus the position along the beamline in Z of the CRV stub.

The second missing piece of making cosmic events more realistic is taking into account the time it takes for the cosmic to travel in flight from the CRV until it reaches the tracker. This is extremely evident if we turn back to looking at the dT vs. $Z(\text{CRV})$ plot: cosmics which pass through the CRV farther from the tracker naturally have a larger time difference between the track and CRV. This is exactly the effect that causes the two linear legs of the dT vs. $Z(\text{CRV})$ distribution, with the slope of the linear correlation representing the time of flight of the particle.

Since this effect is so readily identifiable in data as the linear correlation between dT and $Z(\text{CRV})$, we can apply a linear fit to the data and correct the dT vs. $Z(\text{CRV})$ distribution using the linear fit. Now, after applying the correction for the time of flight in Z , the distribution of dT vs. $Z(\text{CRV})$ represents more of a band than a 'V' shape. The width in Y of this resolution is now driven

by the location of each cosmic in (X,Y), where the time of flight in X and Y is uncorrected.




Figure 5.26: Corrected time difference between track and CRV stub versus $Z(\text{CRV})$. This distribution shows the difference in time between the timing of the track in the tracker and the CRV stub in the CRV after the CRV stub timing has been corrected from propagation time in the CRV scintillator and empty space between the CRV and the tracker versus the position along the beamline in Z of the CRV stub.

5.5.3.5 Dead-Time Window Optimization Results

The distribution in Figure 5.26 can now be used to determine how long the CRV dead-time window should be. Since all of the events in this dataset are cosmic ray events by design, all of these events should be vetoed by the CRV dead-time window. Thus, the length of the dead-time window must be long enough that the cosmic particle that entered the DS through the CRV is no longer detectable in the tracker or calorimeter. Figure 5.26 shows the time difference between a particle being detected in the CRV and that same particle being detected in the tracker, so the width of this distribution along the Y axis may be used to determine the necessary length of the dead-time window.

Upon inspection of Figure 5.26, a vast majority of the events in the distribution have a dT between -50ns and $+30\text{ns}$. 98.5% of the events in the distribution fall within this 80ns window, and a reduction in the dead-time window from 125ns to 80ns is a 36% reduction in dead time. To check how many events may fall outside of this new window, a small simulation job was completed that counted events that were rejected by various event cuts. The result of this simulation indicated that only 0.95% of events that should be rejected by this cut would fall outside of this window, 101/10617 events. The reduction in dead-time paired with the small fraction of events falling outside of this new cut window makes a compelling case for shortening the CRV dead-time window.

References

- [1] R. K. Kutschke, *J. Phys.: Conf. Ser.* **396**, 022028 (2012).
- [2] G. Pezzullo, (private communication).
- [3] D. Brown, G. Pezzullo, and P. Murat, in *Proceedings of CTD 2020* (2020) p. 1, [PROC-CTD2020-59].
- [4] C. Green, J. Kowalkowski, M. Paterno, M. Fischler, L. Garren, and Q. Lu, *J. Phys.: Conf. Ser.* **396**, 022020 (2012).
- [5] Geant4 Collaboration, “Geant4 Overview,” (2023), [Online] Available at: <https://geant4.web.cern.ch/about/>.
- [6] F. Abdi, *et al.* (Mu2e Collaboration), “Mu2e Run I Sensitivity Projections for the Neutrinoless $\mu^- \rightarrow e^-$ Conversion Search in Aluminum,” (2022), arXiv:2210.11380 [hep-ex].
- [7] A. Edmonds, S. R. Soleti, R. Culbertson, Y. Oksuzian, and D. Brown, “Trkana,” (2019), [Online] Available at: <https://mu2ewiki.fnal.gov/wiki/TrkAna>.
- [8] R. Bernstein, “Electrons per Antiproton at the Stopping Target,” (Mu2e Simulations Meeting 2018-04-12), mu2e-doc-16775-v8.
- [9] A. Edmonds, (private communication).
- [10] L. Bartoszek and *et al.* (Mu2e Collaboration), *Mu2e Technical Design Report*, Tech. Rep. (Fermilab, 2015) arXiv:1501.05241 [ins-det] .
- [11] P. Murat, (private communication).
- [12] Geant4 Collaboration, “Shielding - physics list guide,” (2020), [Online] Available at: https://geant4-userdoc.web.cern.ch/UsersGuides/PhysicsListGuide/BackupVersions/V10.7/html/reference_PL/Shielding.html.
- [13] Geant4 Collaboration, “Qgsp_bert - physics list guide,” (2020), [Online] Available at: https://geant4-userdoc.web.cern.ch/UsersGuides/PhysicsListGuide/BackupVersions/V10.7/html/reference_PL/QGSP_BERT.html.

- [14] Geant4 Collaboration, “Qgsp_bic - physics list guide,” (2020), [Online] Available at: https://geant4-userdoc.web.cern.ch/UsersGuides/PhysicsListGuide/BackupVersions/V10.7/html/reference_PL/QGSP_BIC.html.
- [15] Geant4 Collaboration, “Ftftp_bert - physics list guide,” (2020), [Online] Available at: https://geant4-userdoc.web.cern.ch/UsersGuides/PhysicsListGuide/BackupVersions/V10.7/html/reference_PL/FTFTP_BERT.html.
- [16] G. DeFelice, “SU2020 Antiprotons Generation,” (Mu2e Simulations Meeting 2020-05-14), mu2e-doc-33187-v1.
- [17] J. D. Roy, Ph.D. thesis, Syracuse University (1974).
- [18] G. Pezzullo, “Trigger Working Group Track Triggers Update,” (Mu2e Trigger Meeting 2020-01-13), mu2e-doc-30865-v1.
- [19] Free Software Foundation, Inc., “Gdb: The gnu project debugger,” (2023), [Online] Available at: <https://www.sourceware.org/gdb/>.
- [20] Free Software Foundation, Inc., “Gnu operating system,” (2023), [Online] Available at: <https://www.gnu.org/home.en.html>.
- [21] M. Devilbiss, “Tuning CPR and TPR Reconstruction,” (Mu2e Trigger Meeting 2020-08-03), mu2e-doc-34690-v1.
- [22] M. Devilbiss, “Recent Studies on CPR vs TPR Reconstruction for Cosmics,” (Mu2e Computing and Software Meeting 2020-09-02), mu2e-doc-35050-v1.
- [23] M. Devilbiss, “Improvements in Cosmic TPR Acceptance: Comparing to Reference,” (Mu2e Fall 2020 Collaboration Meeting 2020-10-20), mu2e-doc-35657-v1.
- [24] Fermilab, “The art Event Processing Framework,” (2009), [Online] Available at: <https://art.fnal.gov>.
- [25] R. Bernstein, “Pions per Antiproton Annihilation Comparison,” (Mu2e Simulations Meeting 2021-05-05), mu2e-doc-38067-v2.

Chapter 6 Estimating the CRV Aging Rate

While the Mu2e collaboration prepares to assemble its detector apparatus, pieces of each detector are arriving Fermilab for testing and storage. Nearly all CRV modules have been shipped to Fermilab and are being tested for functionality upon arrival. A few modules are being used for more extensive studies on detector response and long term aging and stability. The CRV module test stand at FNAL came online for the first time in May 2021 and has been collecting cosmic data ever since. This chapter will describe the test bench setup in Wideband and the electronics involved with collecting data using the test bench. The focus of this chapter will be to report the first study using Wideband data to estimate the aging rate of the CRV. First, the problem of scintillator aging will be discussed. Then, the methodology for estimating the overall CRV aging rate will be outlined and the results of this year-long study will be discussed.

A large part of this chapter discusses an investigation into the quality of Wideband data and how the response of the CRV modules depend on variables like the environmental temperature or the bias voltage of electronic components on the FEBs. This work was the first of its kind to discover and begin to characterize the dependencies of the CRV response to these parameters. This work has been important in understanding how to analyze the first data to be collected from full CRV modules. In addition, the discoveries made about the dependence of the CRV on temperature and electronics bias voltages were instrumental in calibrating CRV channels and have already been implemented in official CRV calibration procedures. Based on the lessons learned, the Wideband test bench has evolved to better monitor temperatures in the environment and on electronics components and more accurately calibrate the bias voltages of the FEBs after the conclusion of this study.

6.1 Scintillator Aging

Scintillator aging is a phenomena that has been studied for decades, primarily for physics and astronomy applications. At Fermilab, extensive scintillator aging studies were performed for the Solenoidal Detector Collaboration (SDC) for the Superconducting Super Collider (SSC) in 1993 [1], the MINOS detectors, and the MINERvA detector [2]. Similarly, the T2K detector in Japan uses plastic scintillators and recently published a study on how their scintillator has aged over the past 10 years [3]. During these experiments, it has been observed that scintillator-based

detectors lose efficiency over time typically on the order of 1-3% per year. However, the aging rate is not always consistent over time, and it has been observed that generally scintillators age rapidly shortly after production and tend to age more slowly with time [4]. With this trend in mind, it is most beneficial to study aging trends over a long period of time to try and understand the stable aging rate of the scintillator.

There are several factors that influence long-term stability and performance of scintillator-based detectors, though the specific degradation mechanisms leading to performance reduction are not fully understood. These factors include oxidation, diffusion of small molecules, and crazing, which can occur at the interface with other materials in contact with the scintillator polymer [4]. Moreover, scintillators that are in uncontrolled environments age more rapidly than scintillators in controlled environments. Many degrading effects in plastics are enhanced by extreme temperatures or humidity. Any organic material is susceptible to aging from environmental effects, scintillators are particularly complex due to their multi-component mixtures and varied preparation techniques that could involve proprietary parameters [4].

Aging studies and simulations were started immediately after the designs for the CRV scintillator extrusions were finalized. Initial results from prototype counters that were produced in 2014 and 2015 showed promising aging results around 3% per year. However, the pre-production counters that were produced in 2016 yielded aging rates around 8% per year [5]. The differences in these two aging rate studies early on in the production phases of the CRV counters has been a point of concern in the collaboration and was the driving factor for further studying the aging in Wideband in production counters. It is important for Mu2e to understand the aging of the CRV and evaluate if there are any impacts on the performance CRV from aging in the counters.

6.2 Wideband Test Bench Setup

As Mu2e prepares to assemble the CRV, completed modules are being shipped from University of Virginia and have been arriving at Fermilab. The CRV modules are being stored in a building called Wideband, where they are organized into stacks of four modules separated by foam padding. The Wideband storage area and stacks of CRV modules are shown in Figure 6.1a. Each stack is covered with a tarp to protect the modules from dust, debris, and any potential leaks in the building. The stacks of modules rest on four inches of foam padding between the stack and the floor to protect the modules from any potential flooding. The Wideband building is mostly underground, so covering the modules with tarps and using foam padding under and between the modules is a precautionary measure to protect the modules from environmental damage in case of an accident. The desks and blue electronics rack in the center of Figure 6.1a are working areas for the CRV test bench operators. The trigger electronics and power sources for the test bench are housed in the

electronics rack and the laptop shown to the right in Figure 6.1b is a local laptop that runs the test bench graphical user interface (GUI) and records data from the test bench. The module that was used for these aging studies is the top-most module of the stack shown in Figure 6.1b, module 127. Module 127 is a side module in the CRV-L sector.



(a) Wideband CRV module stacks



(b) Wideband CRV module stack close-up

Figure 6.1: Photographs of the Wideband building at Fermilab. Figure a) shows a wide view of the Wideband CRV module storage area. Under each tarp is a stack of four CRV modules. Figure b) shows a closer view of one stack containing the module of interest for this study, module 127. In Figure b), module 127 is connected with HDMI connectors to the FEBs housed in the aluminum enclosure to the right.

To collect cosmic data using the modules, several layers of readout electronics are used.

While the Wideband test stand is not using the full Mu2e DAQ system, it uses most of the front-end electronics that will be employed in the final experiment’s full DAQ chain. Each CRV dicounter pair is outfitted on each end with CMBs, which hold the SiPMs which detect the light signals that originate from cosmics passing through the CRV bars. When discussing module data, we refer to ‘channels’, with one SiPM serving each channel. Thus, each CMB serves four channels of the module. One module contains 128 channels. CMBs are mounted on top of the fiber guide bars at the end of each dicounter. On the outside of each CMB there is an HDMI connector, which connects the CMB to its FEB. Each FEB is connected to the ROC board via one of the Ethernet (RG45) ports. In this test stand, the ROC has little functionality; it provides the trigger to the FEBs, as well as the power via one of the Ethernet (RG45) connections. Data flows directly from the FEBs to a local laptop via the second Ethernet (RG45) port on each FEB. The second Ethernet (RG45) connection plugs into an Ethernet (RG45) switch with eight ports to route data from all FEBs to the local laptop. In the final system, the second Ethernet (RG45) port is used by the ROC to send data to the full DAQ system. A diagram of the electronics used in the test bench is shown in Figure 6.2.



Figure 6.2: Wideband test bench electronics diagram. A block diagram of the electronics used to record data at the Wideband test bench [5]. Each block is labeled, beginning with a CMB on the left, an FEB in the center, and a ROC on the right. Short descriptions accompany each block.

The last critical component of the Wideband test bench is the trigger system. It consists of three scintillating trigger paddles, a coincidence module, and a NIM logic gate generator. Each trigger paddle is attached to a photomultiplier tube (PMT). The paddle efficiency and optimal operating voltage of the PMT were determined at University of Virginia at the time of production for a subset of the trigger paddles. For aging studies, three trigger paddles were used to form a three-fold coincidence via the coincidence module. The trigger paddles are distributed around module 127, one on top of the module and two beneath the module. The paddles are aligned to overlap as closely as possible. Using three trigger paddles stacked on top of each other rather than a single paddle for triggering improves the ‘purity’ of the cosmic coincidence and restricts the allowed angular distribution of the paths of incoming cosmics. With the module in between the trigger paddles, any incoming cosmics that trigger the paddles should also be detected by the

module. The trigger paddles are shown around module 127 in Figure 6.3. The trigger paddles are



Figure 6.3: Photograph of the Wideband test bench trigger paddles. A photograph of the CRV module stack containing module 127 and the trigger paddles surrounding module 127. The three PMTs at the ends of the three trigger paddles are visible on the left side of the image. The top trigger paddle is visible as the black wrapped bar laying across the top of the stack. The lower 2 paddles also lay across the stack, aligned underneath the top trigger paddle for maximum overlap.

powered by a standalone high voltage power generator to a voltage of approximately 1400V on each PMT. Before operating the trigger paddles, I determined the optimal operating voltage of each PMT, detailed in Appendix 2. Since the optimal operating voltage had been determined for some panels at the time of production, re-evaluating the optimal operating voltage would look for signs of aging on the trigger paddles as well as verify the optimal operating voltage that was determined at production. When all three paddles receive a signal, the coincidence module sends a NIM signal to the ROC, which tells the FEBs to begin recording data. A trigger begins a spill cycle which consists of a nominal 120 second spill-on, during which cosmic data is being recorded, followed by 60 seconds of spill-off, during which the data is being transferred.

When a trigger arrives at the ROC, a waveform is recorded with 127 ADC samples, corresponding to a length in time of 1600ns. The data that is written out for each trigger uses the FEB buffers, allowing us to get data from before the trigger arrives. The length of the pre-signal region may be adjusted using the configuration file that is included with the processing code. An example waveform is shown in Figure 6.4. In this waveform, the large peak is the signal, occurring around 1100ns. This is roughly the time of the trigger. The trigger time may be a bit later, because it is delayed due to photon travel time in the scintillator and electronics response times. The interval of time before the signal is referred to as the pre-signal region of the event. In the pre-signal region,




Figure 6.4: Raw waveform of a cosmic event in the CRV. An example of a waveform collected in a random channel in response to a single trigger. The y-axis is in units of ADC counts and the x-axis shows time in ns. The signal is the large peak around 1100ns. Here, the first 70 ADC samples, from 0 to 880ns is the pre-signal region, and subsequently 880ns to 1600ns is the signal region.

dark counts are visible since this region is usually undisturbed by signals. In Figure 6.4, a single PE dark count is visible around 400ns. The length of the pre-signal region may be modified in a configuration file during data processing, which will be discussed later in this section.

6.2.1 Collecting Data with the Test Bench

To start a run and begin collecting cosmic data with the test stand, a few steps must be performed. First, the trigger paddles should be functioning and their high voltage power supply must be turned on. Next, parameters like the spill-on length, spill-off length, and external trigger use must be passed to the ROC via the local Wideband laptop. A terminal is used to pass these settings to the ROC from the command line. The next steps are to set up the communication to the FEBs, which is done using a special purpose run control (RC) GUI. The GUI has buttons for each FEB that begin colored gray and turn green or red to indicate whether or not communication has been successfully established with each board. After initiating the connection with the FEBs, the SiPM bias voltages are specified by loading configuration files into the GUI using a command line interface within the GUI. The bias voltages are around 50V for each channel, and were determined by collaboration members at Kansas State University.

The RC GUI controls the start and end of a run, and automatically increments the run number to appropriately name output files. While collecting data, the state of each FEB is shown on the main RC GUI screen using a numerical ‘spill status’ to denote whether each FEB is transferring

data, collecting data, or stuck. The number of triggers in the last spill and the cumulative total number of triggers collected by each FEB during the run are also shown on the screen. Generally, all FEBs should be getting the same number of triggers, so these numbers should match at all times. Data transfer may also be monitored from within the RC GUI, as the size of each spill and speed of transfer are also written out on the main screen between spills. Again, these spill sizes should match for all FEBs since all boards should be receiving the same number of triggers. If there is a problem with data transfer, an error message will be displayed on the screen.

The raw files that are made by the Wideband DAQ software are written out locally to the laptop in Wideband. These raw files are large, on the order of 10GB per day of data taking, so they must regularly be transferred to the /mu2e/ tape-backed disks for storage. This test stand uses the Mu2e file transfer service (FTS) that is currently under development by Ray Culbertson to transfer files to disk. Each file from Wideband is labeled with a unique 6-digit ‘run number’, which is assigned at the start of each data collection period. Since the raw Wideband files are so large, the files are often split into chunks. The resulting chunks contain 1000 spills each. Sometimes, the final chunk is small and only includes few spills. After files are split, they also get a 3-digit ‘subrun number’, starting at 000 for each run number and incremented based on the number of chunks the file was split into. Files from Wideband also have one more identifier, called the configuration number. Each configuration number refers to a specific set of FEBs connected to a specific module. Using this identifier, datasets from the same setup may be appropriately selected and analyzed together. For aging estimations, it is crucial to keep all elements of the test stand as stable as possible to track only the effect that passing time has on the module of interest. Once the output files are properly named, split, and transferred to the /mu2e/ disks, they can be used with the Mu2e codebase for further analysis/

6.2.2 Wideband Data Processing Code

The raw files that are moved to the /mu2e/ disks are a mixture of text and binary files, so they must be processed to organize the data into meaningful, human-readable files and plots. The processing code for the Wideband test stand was written and is maintained by Ralf Ehrlich, a postdoc from University of Virginia. Processing raw files occurs in three steps: parsing, calibration, and reconstruction. Each of these steps is described in the list below.

- Parsing: takes the raw files and converts them into more usable ROOT files.
- Calibration: reads the parsed ROOT file and creates a .txt and a .pdf file for each input parsed file. During this step, dark pulses in the pre-signal region of events are located by searching for peaks in the ADC counts vs time domain that correspond to integer numbers of PEs. Typically, one- and two- PE peaks are reconstructed and fit to determine the pedestals and

calibration constants for each channel, which are then stored in the output .txt file. The calibration constant for a given channel is determined by measuring the pulse area as given by the integral under the one-PE peak that was found in the pre-signal region of an event. The output .pdf file contains plots of the pedestal and calibration values of each channel.

- **Reconstruction:** analyzes pulses that occur in the signal region of each event. This step reads the parsed ROOT file and creates a .pdf file and a new ROOT file containing pulse information for each event in each channel, such as PE distributions and timing.

Each step of processing is performed using a .cpp script which makes an executable that can be used on a raw file. In addition to the scripts, there is also a configuration text file, config.txt, that must be included with this workflow. The configuration text file sets the value of some relevant constants for analysis and output destinations for each stage of processing. The time periods within an event that correspond to the beginning and end of the pre-signal and signal regions of an event, as was mentioned above, are a few of the constants included in config.txt. The location of the signal region also depends on how the time delay of the FEBs is set relative to the arrival of the trigger. All of these values may be adjusted as needed by editing the values of the parameters in the config.txt file.

6.3 Wideband Aging Measurement Strategy

With the test stand active, the general strategy for measuring CRV aging is to track the amount of light, quantified by the number of PEs, that each channel in the module observes over time. The peak number of PEs collected in a given channel during a run is referred to as the ‘light yield’. Due to aging effects, a decrease in light yield over time is expected. To perform these studies, a separate analysis workflow was developed and maintained. The starting point for these aging studies uses the reconstruction ROOT files generated by the processing code described in the previous section. Within the reconstruction files, there are several PE distributions that may be used to begin extracting the light yield from every channel. The PE distributions in the reconstruction files generally come in two flavors: corrected and uncorrected. Some of the electronics used in the test bench have temperature-dependent response, so the corrected data includes adjustments performed by Ehrlich for different electronics effects. For aging studies reported here, the uncorrected distributions are used as a starting point and temperature corrections are calculated and applied separately in order to determine the light yield from each channel. The purpose of separately applying temperature corrections at this stage was to iterate and improve on the corrections that Ehrlich otherwise applies during the data reconstruction step of processing.

The corrections applied to Wideband data in both Ehrlich’s data processing code and my independent analysis code are meant to correct for the SiPM quantum efficiency with respect to environmental temperature. To measure the CRV SiPM response versus temperature, a test beam

run was performed in April 2022. For the test beam run, a short, 1-meter long, CRV dicounter was outfitted with a CMB. To manipulate the environmental temperature, a heating strip was wrapped around the FGB which holds the CMB to heat the SiPMs that are in contact with the dicounter. The dicounter was then placed in a 120 GeV proton beamline. The heating tape was turned on until the temperature sensor on the CMB reached 35 degrees Celsius, and subsequently was turned off while cosmic ray data was recorded as the SiPMs cooled over time. Figure 6.5 shows the resulting plots that were collected of the light yield of cosmic events versus temperature reported by the temperature sensor on the CMB. Recall that each dicounter has four SiPMs and thus four channels from which to collect data. During this test beam run, one channel did not successfully record cosmic events. As a result, only three channels are shown in Figure 6.5.

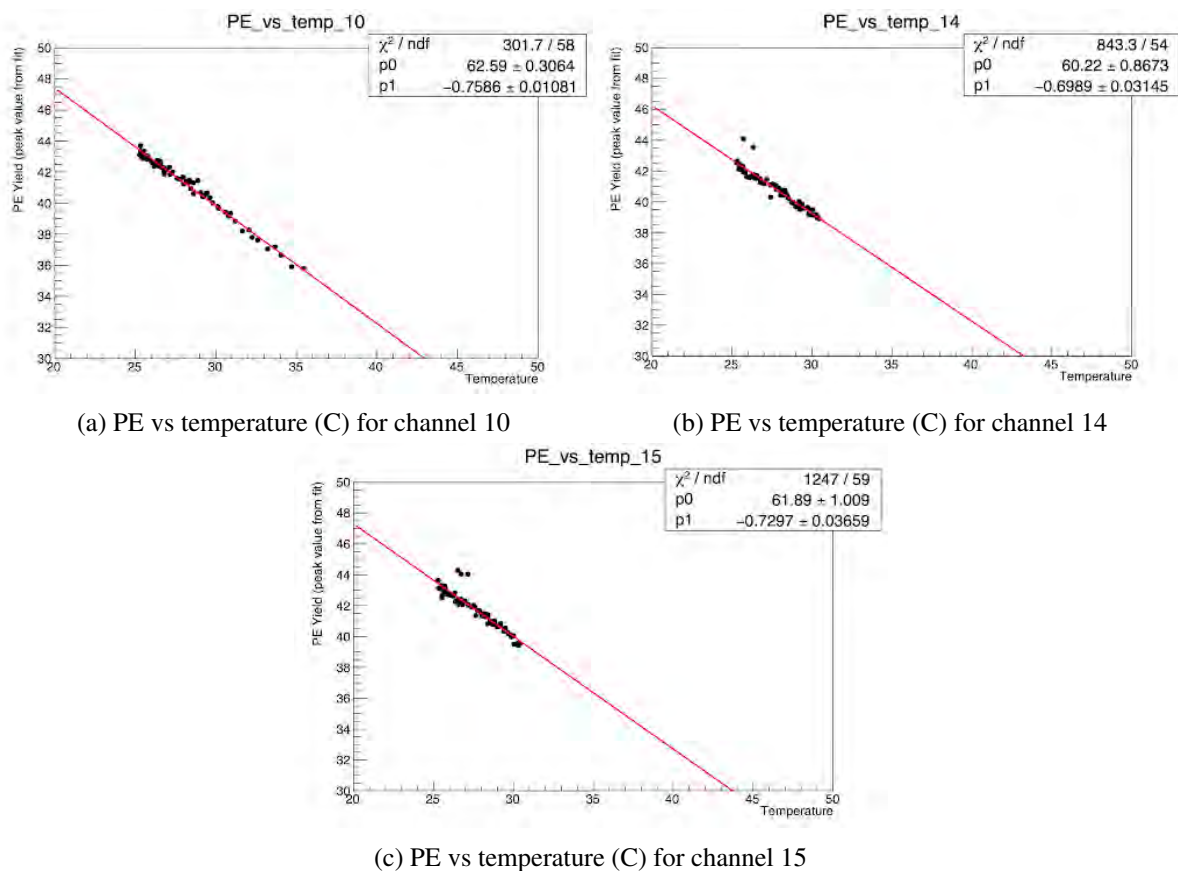


Figure 6.5: Test beam data for PE versus temperature SiPM corrections. Light yield, in PEs, versus temperature, in degrees C, from the data collected during the test beam run in April 2022 with the heating tape applied to the dicounter in the test beam. The three different figures show the three functional channels in the heated dicounter: figure a) shows channel 10, figure b) shows channel 14, and figure c) shows channel 15.

Figure 6.5 shows the results of light yield vs temperature during the test beam run. There is a clear inverse relationship between temperature and SiPM quantum efficiency, as expected. The

three channels that were able to be read out all yielded a similar fit to the PE vs temperature data, indicated by the lines of best fit in Figure 6.5. Since one CMB serves all of the channels in a dicounter, three separate runs were taken to record independent temperature data for each channel. To extract the temperature correction, the best slope was calculated by minimizing the chi squared error of the distribution of the three plots simultaneously. This can be compared with how the temperature correction is applied within Ehrlich's processing code to generate the 'corrected' set of plots. Ehrlich calculated his temperature correction by cutting a few outlier points and then taking the average of the three negative slopes. The two methods result in percent-level differences in the negative slope parameter used for the temperature corrections.

The slope thus calculated is used in the system of equations below to determine the temperature-corrected PE values for the data:

$$PE_0 = \text{slope} \times T_0 + PE_{\text{int}} \qquad PE_{\text{corr}} = \text{slope} \times T_{\text{ref}} + PE_{\text{int}} \qquad (6.1)$$

$$PE_0 - PE_{\text{corr}} = \text{slope} \times (T_0 - T_{\text{ref}}) \qquad (6.2)$$

$$PE_{\text{corr}} = PE_0 - \text{slope} \times (T_0 - T_{\text{ref}}) \qquad (6.3)$$

where T_{ref} is a reference temperature, which was set at 25 degrees Celsius. The temperature T_0 is read out from the CMBs once per spill, as well as the uncorrected PE_0 value. As a result, the uncorrected PE_0 value is linearly shifted to the corrected PE_{corr} value. We believe that such linear correction is valid for temperatures close to the reference temperature, due to the fact that the PDE and gain of the SiPMs is also approximately linear in this region as shown in Figure 6.6. The overvoltage applied to the SiPMs according to the FEB configuration files is in the range of 2-4 volts, so the left hand side of the plot is the region that we are concerned with. The left end of the PDE curve is approximately linear.

Before applying the spill-by-spill temperature corrections, quality cuts were applied to the data taken for module 127. During these quality cuts, runs are marked as 'bad' or 'good'. If a run is too short and the PE distribution is too sparse to perform a stable fit or if a run is missing any of the relevant information needed for the aging analysis, is it labeled as bad and excluded from analysis. If the CMB temperature readout was stuck, then the temperature information is missing for corrections, then the run is marked bad. If the bias voltages on the FEBs were set incorrectly or not at all, the run is also marked as bad. This bad run determination is entered in a Google sheet of all the runs taken since the start of the Wideband test bench that lists for each run the start date, the number of spills, the number of events and other relevant test stand status information. Bad runs are then excluded from analysis. Examples of uncorrected and corrected PE distributions in a randomly chosen channel from a randomly chosen good run are shown in Figure 6.7.




Figure 6.6: General characteristics of Hamamatsu S13360-2050VE SiPM. This figure shows general characteristics as measured by the manufacturer of the SiPMs chosen for the CRV. The dashed line along with the right-hand Y axis describe the PDE of the SiPMs as a function of the applied overvoltage. The solid line represents the gain, and the dash-dot line represents the crosstalk probability. Taken from Hamamatsu S13360-2050VE data sheet [6].

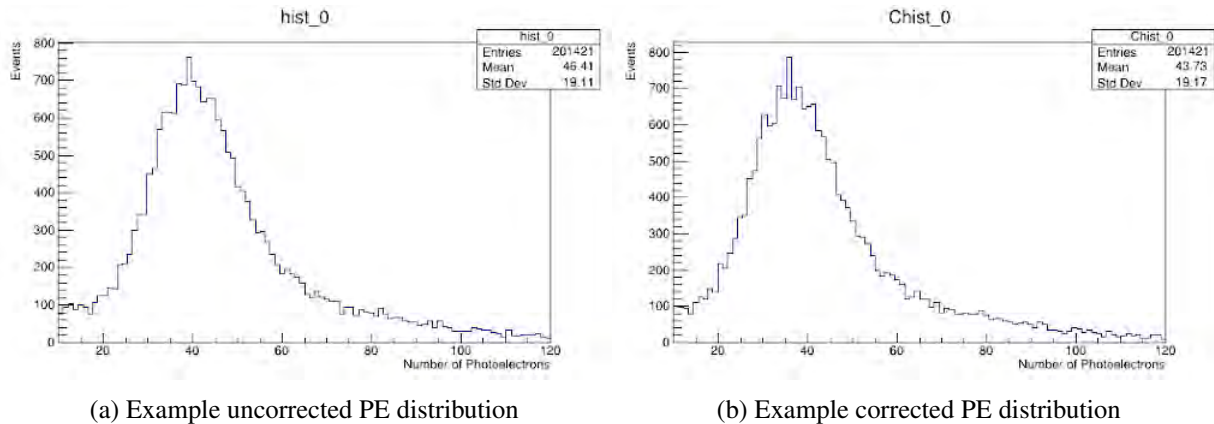


Figure 6.7: Effect of temperature correction on a PE distribution. A side-by-side comparison of an uncorrected PE distribution with a corrected PE distribution to illustrate the effect of these temperature corrections. Figure a) shows an uncorrected PE distribution. In figure b), the same PE distribution is shown, now corrected after applying temperature corrections.

In Figure 6.7, the x-axis does not begin at 0 PE, instead it begins at 10 PE. Events with less than 10 PE are cut from the distribution because they likely originate from dark counts, not cosmic muons. After applying the temperature correction, the distribution is slightly shifted. Temperature correction proceeds on an event-by-event basis while the temperature sensors on the CMBs record temperature once per spill, every three minutes. Within every 3-minute spill, there are about 150

single events recorded. Each event is corrected according to its corresponding temperature.

6.3.1 Temperature and Bias Voltage Stability

To perform the temperature corrections that we have been discussing, it is necessary to have accurately reported temperature from the CMBs on the module. We can plot the temperature of the CMBs, as well as the FEBs, to ensure that temperature recordings are stable and sensible with respect to environmental conditions at Wideband. Figures 6.8a and 6.8c shows the temperature in FEBs 0 and 1 versus hours since the beginning of the subrun. Likewise, Figures 6.8b and 6.8d shows the temperature recorded in CMB0 and a random CMB versus hours since the beginning of the subrun for good runs. The x-axis was made by converting spill numbers to length of time in hours.

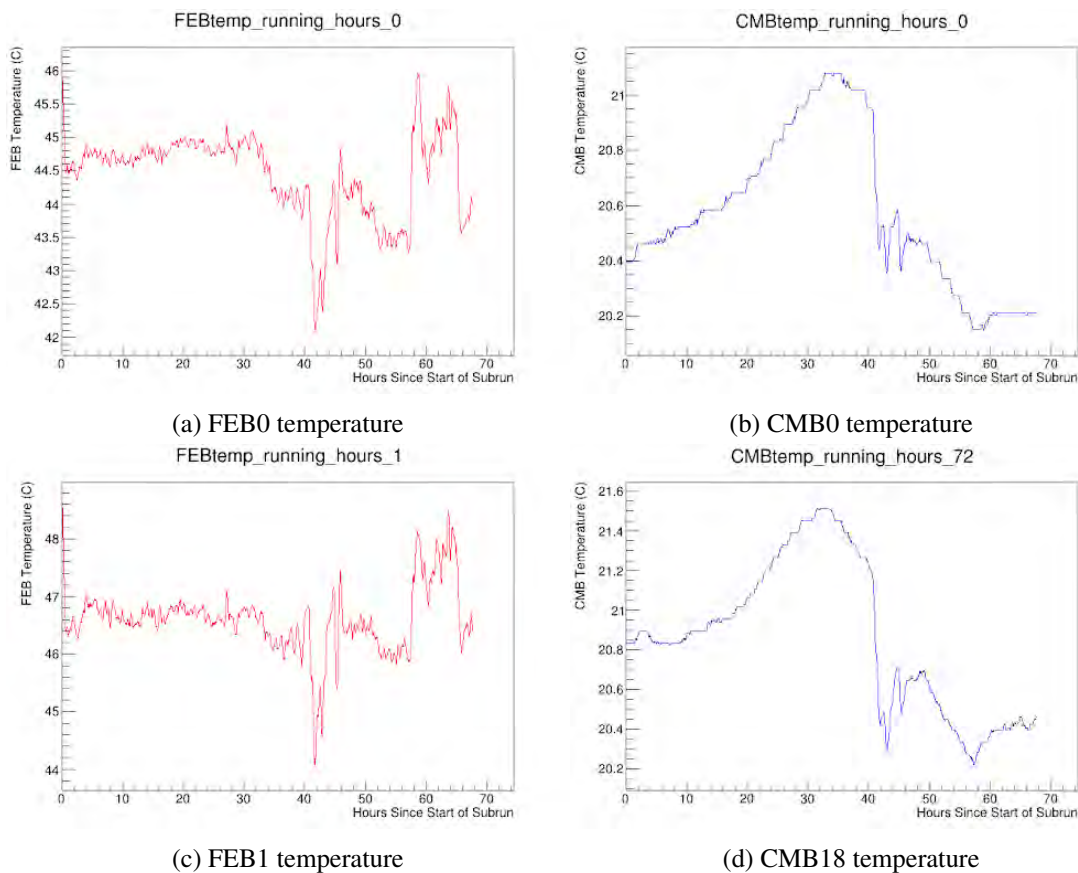


Figure 6.8: Example CMB and FEB temperature plots. Temperature distributions collected for both FEBs and two CMBs on module 127. Figures a) and c) show temperature distributions for FEBs 0 and 1, respectively, in degrees C as recorded by the sensor on the FEBs versus hours since the start of the subrun. Figures b) and d) show temperature distributions for CMBs 0 and 18, respectively, in degrees C as recorded by the sensor on two different CMBs with the same axis. Running averages were taken over a period of 3 spills, roughly 10 minutes, to slightly smooth the plots.

From these plots, it is clear that temperature that is recorded from the CMBs is fairly consistent, within one degree over a time period that spans several days. The temperature that is recorded from the FEBs has more variation since the FEBs are not in direct contact with the module's thermal mass. Still, the FEB temperature is relatively stable and only varies a few degrees over several days. By inspection of the temperature plots like the plot shown above in Figure 6.8, slight diurnal temperature variations are observed. Wideband is a large and rather old building, so diurnal variations were expected. Large temperature fluctuations, like the one seen in Figure 6.8 around 40 hours are validated by comparing temperatures reported by different CMBs, like in Figures 6.8b and d, or by the second FEB over the same period of time, like in Figures 6.8a and c. Since both sets of FEB and CMB temperatures shown in Figure 6.8 track each other relatively well, we can hypothesize that the rapid change in temperature comes from a real effect, like the Wideband garage door opening during the measurement.

Temperature variations can affect the bias voltage of the SiPMs in the module. Therefore, it is important to monitor the reported bias voltages for each channel and check for deviations from the value it was set to at the beginning of the run. Each FEB controls the bias voltages of 64 channels using general control buses that are set by the FEB configuration files at the beginning of the run. There are eight general control buses on each FEB, therefore the value assigned to one bus is used to determine the voltage for eight channels. The configuration files set an overvoltage value for each channel, which is used with the bus voltage to set each individual SiPM to its optimal operating voltage as determined by the team at Kansas State University.

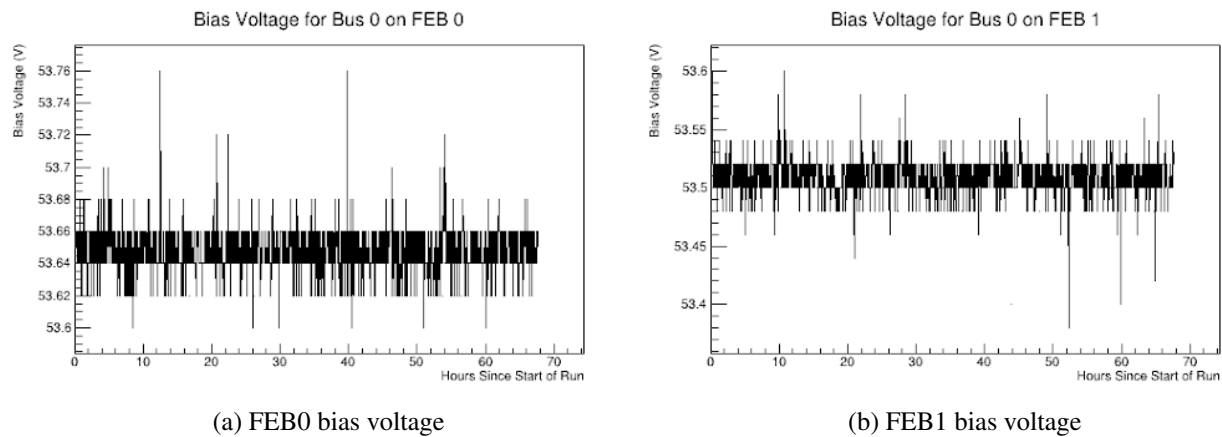


Figure 6.9: Example FEB bias voltage plots. Bias voltage distributions collected for both FEBs on module 127 for an example run. Figure a) shows the bias voltage distribution of bus 0 for FEB0, in volts versus hours since the start of the subrun. Figure b) shows the bias voltage distribution of the same bus, bus 0, for FEB1, in volts with the same axis.

The plots in Figure 6.9 show that the bias voltage on bus 0 is stable on both FEBs. The small variations that we see here are not large enough to be concerning. The bias voltages on bus 0

are slightly different for each FEB, this is expected because each FEB has its own configuration file. Similar small variations are seen in all buses upon inspection of analogous plots for the different buses on each FEB.

6.4 Results

Now that we understand how the test bench functions and have confidence in the electronics temperature and bias voltage stability on the FEBs within their environment, we can analyze the cosmic ray data that is collected and look for signs of aging. To look for aging, the light yield, in PEs, is plotted over time and any measured decrease in light yield over time is used to estimate the rate of aging in the detector. This section will present and analyze plots of light yield over time for data collected using module 127 in Wideband. Since the aging plots are presented in PEs over time, the aging rates will be presented in the same units: loss of PEs over time, as a percent per year rate. The light yield of each channel at time = 0 is taken as the channel's maximum light yield, and thus the amount of PEs lost over time may be turned into a percentage of the total channel's yield lost per year. For this investigation, plots that are referred to as "aging plots" are the plots of light yield, in PEs, over time with a linear fit.

6.4.1 Light Yield Distributions

The first step in extracting the light yield in each channel over time is to apply a reasonable fit to the light yield distributions that are collected in each channel. Then, the peak of the fit may be used as the channel's light yield on the date when the data was recorded. From Figure 6.7, it is clear that the temperature corrected PE distributions are not Gaussian, and have a rather long tail extending out to high PE values. Fit optimization studies found that a Gaussian plus a second-degree polynomial returns fits with lower chi-squared values than a Gaussian plus a Landau or a Gaussian plus a Landau plus a second degree polynomial. An example PE distribution with a good fit is shown in Figure 6.10. In Figure 6.10, the Gaussian plus second-degree polynomial fit is shown in red. The blue fit in Figure 6.10 is a pure Gaussian fit, which supplies initial values for the Gaussian plus second-degree polynomial fit to improve fit accuracy.

There are some instances where the fit to the PE distribution fails. The most common failure mode is when a channel has a low number of events. These are referred to as 'low occupancy' PE distributions. In these cases, the PE distribution does not have a clear, well defined peak feature from the fit. An example of a low occupancy PE plot with a bad fit is shown in Figure 6.11. Similar to Figure 6.10, Figure 6.11 shows two fits, where the red is the Gaussian plus second-degree polynomial fit and the blue is a pure Gaussian fit to seed initial values for the Gaussian plus second-degree polynomial fit. Many low occupancy PE distributions are either from channels that are on




Figure 6.10: PE distribution example with fit. A temperature-corrected PE distribution with its final Gaussian + polynomial fit shown in red. The blue line shows a restricted-range pure Gaussian fit, which sets initial values for the final fit. The number of PEs in a single event is shown on the x-axis and the number of events for each PE value is shown on the y-axis.




Figure 6.11: Low occupancy PE distribution example. A temperature-corrected PE distribution with low occupancy. The red line is the final Gaussian + polynomial fit and the blue line shows a pure Gaussian which sets the initial values of the final fit. The number of PEs in a single event is shown on the x-axis and the number of events for each PE value is shown on the y-axis. In the statistics box, the integral is below 2000, which indicates that this PE distribution is too sparsely populated to have a proper fit. Channel 62 is a channel on the edge that is always sparsely populated.

the very edge of the counter, and thus may not have full coverage from all three trigger paddles,

or from the last subrun file in a run which was split, as sometimes these last files contain a small number of spills and thus few events in total. Requiring that the integral value of the PE plot be above 2000 removes most of these fit failures.

To prevent any remaining channels with bad fits from skewing the aging plots, a cut on the total chi-squared value of the fit is applied. Figure 6.12 shows the histogram of the chi-squared values from each PE distribution in the entire set of data after applying the low occupancy cut. We

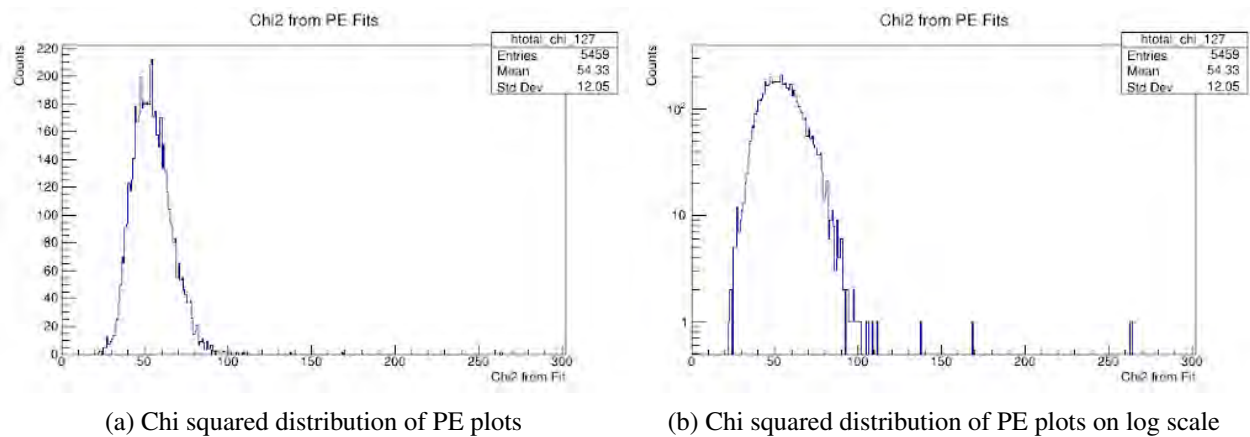


Figure 6.12: Chi squared distributions of fits of PE plots. Figure a) shows the distribution of total chi squared values from the fits of every PE distribution that is included in the dataset. Figure b) shows the same total chi squared distribution with a log scale on the y-axis. The log plot emphasizes the few PE distributions that contribute high chi squared values.

choose to remove light yield data points obtained from fits with a chi square greater than ~ 90 , i.e. more than three standard deviations away from the bulk of the chi square distribution. Less than 5% of histograms are removed by this cut.

After these cuts have been applied, we have collections of corrected PE distributions with accurate fits in each channel at different times. To make aging plots for each channel, the Gaussian mean value of the final Gaussian plus polynomial fit from each good PE distribution is extracted as the light yield value for each channel in a given run. The Gaussian mean is a good metric to use for the light yield in a given channel because the Gaussian part of the fit describes the bulk of the distribution. The starting date of each run is known, so the light yield values in a given channel may be plotted over time using a scatter plot. The resulting plots of PE values versus time may then be fit to determine the aging rate of each channel. An example aging plot is shown in Figure 6.13. The errors on the plot were extracted from ROOT as the errors on the slope parameter of the Gaussian plus polynomial fit.

Oftentimes scintillator aging rates are quantified as the percentage of the total light yield lost per year. For this reason, the time axis (x-axis) of these plots is chosen carefully. Instead of simply using the month, day, and year of each run on the x-axis, the date is converted to a 'fraction




Figure 6.13: Example Wideband module 127 aging plot. This aging plot is from channel 42. The x-axis represents time, given by the time fraction of the number of years elapsed with respect to 5/7/2021. The y-axis is the light yield, in PEs. The statistics box in the upper right corner shows the intercept (p_0) and the slope (p_1) of a linear fit of the data.

of a year' or 'time fraction' with respect to the first day that the Wideband test bench came online. So time 0 represents data taken when the test bench first came online on 5/7/2021. Data taken on 11/5/2021, i.e. roughly 6 months after the Wideband start date, is represented with a time fraction of 0.5. By using this time fraction on the x-axis, the aging rate per year can be easily calculated. The slope of the linear fit on each aging plot gives the number of PEs lost per year. By dividing the number of PEs that is lost over the year by the initial light yield of each channel, the percentage of the total light lost per year of each channel may be calculated. As was mentioned above, this percentage of light lost per year is the aging rate in a given channel. For example, from Figure 6.13, we observe that the light yield of channel 42 decreased by ~ 3 PE within the year. Compared to the initial light yield of 40 PE, given by the first point in the plot, the aging rate of this channel is roughly 7.5% per year.

Aging plots like the one in Figure 6.13 are created for each channel in module 127 and aging rates may be extracted for each channel from the corresponding plot of PEs over time. The two plots in Figure 6.14 show the distribution of the aging rates for all the channels in module 127 and the measured aging rate for each channel. From Figure 6.14b, the average aging rate for module 127 is 10.42% per year. In Figure 6.14a, no aging rate is reported for channels 31, 32, and 60-63 because these channels are on the very edge of the module, and do not have full coverage from the trigger paddles. As a result, nearly all of the PE distributions from these channels fall into the low occupancy category and are removed from analysis. There are not enough data points on the plot of




Figure 6.14: Aging rates for module 127. Distributions of the aging rates estimated in all channels of module 127. Figure a) shows the 2-dimensional distribution of aging rates by channel number, where aging rate reported as in percentage of light yield in each channel lost per year. Figure b) shows the 1-dimensional distribution of the same values for all 128 channels of module 127. The plot on the right is the projection of the y-axis of the plot on the left.

PEs over time to make a reasonable aging rate estimate in these channels.

Figure 6.14a seems to indicate different average aging rates for the first half vs the second half of the channels. The two groups of channels in the two halves of Figure 6.14a are served by different FEBs. Because the scintillating bars for a given module are all produced at the same time, there is no reason to suspect that different layers of the same module should have drastically different aging rates. Plotting the aging rates of the channels served by FEB0 and by FEB1 separately in 1-dimensional plots gives a good indication of how much the aging rates from two halves of the module differ from each other. In this configuration, FEB0 serves the bottom two layers of module 127, while FEB1 serves the top two layers. Aging rate distributions for the channels served by each FEB are shown in Figure 6.15.

Taking the mean of each of the distributions shown in Figure 6.15 as the average aging rate associated with the two halves of the module suggests that the bottom half of the module is aging slower than the top half. The channels on the bottom half of the module, served by FEB0, appear to have an average aging rate of $\sim 8.6\%$, while the channels on the top half of the module, served by FEB1, appear to have an average aging rate of $\sim 12\%$. The drastic difference between the average aging rates in the two halves of the module prompted a further investigation on the quality of the data taken over such a long period, which is the main focus of the next subsection and the rest of this chapter.

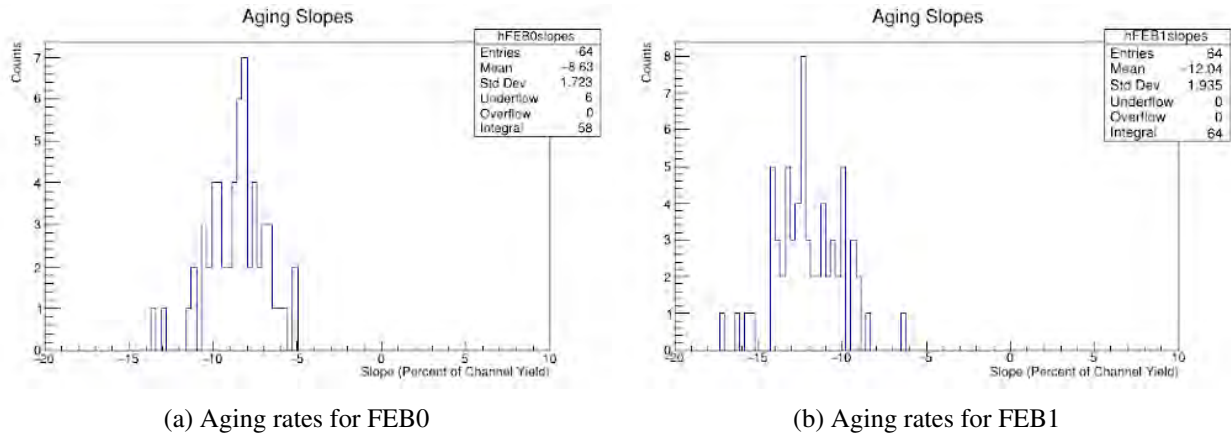


Figure 6.15: Aging rates for FEB0 and FEB1 of module 127. Distributions of the aging rates estimated in all channels of module 127. Figure a) shows the 1-dimensional distribution of aging rates in percentage of light yield lost per year from FEB0, which is connected to the bottom two layers of module 127. Figure b) shows the analogous 1-dimensional distribution of aging rates in percentage of light yield lost per year from FEB1, which is connected to the top two layers of module 127. The bottom half of the module includes the edge channels which are not sufficiently covered by the trigger paddles, these channels are the overflows seen in the plot on the left.

6.4.2 Investigating Data Quality

Upon a closer review, the data that had been collected in Wideband from May 2021 to July 2022 showed a few features that hinted at possible instabilities in the test bench. This subsection will discuss two peculiarities discovered in the data and what was revealed about the test bench at those times.

The first hint of suspicion arose from an independent analysis performed by Yuri Oksuzian. Oksuzian's analysis used the corrected PE distributions directly produced by Ehrlich's processing code which I did not use. The method that Ehrlich uses to produce corrected PE distributions and extract light yields from each channel is similar to mine. However, we apply slightly different temperature correction equations, which are consistent with each other. Ehrlich also uses a different fit function to fit the PE distributions; his final fit is a Gaussian distribution convoluted with a Landau distribution. When Oksuzian performed this analysis, the large difference in aging rates between the channels served by FEB0 and FEB1 was indeed observed. In addition to producing aging results, Oksuzian also plotted the calibration constants associated with the calibration corrections that are performed as part of Ehrlich's analysis workflow, which involve finding a single PE peak and calculating the integral under that peak. This calibration constant is also referred to as the single photoelectron (SPE) value. In Ehrlich's workflow, the temperature correction is always applied on top of the calibration corrections calculated in the calibration step of the data processing procedure. Oksuzian plots the SPE value over time, as seen in Figure 6.16. In Figure 6.16, the black dots are the SPE values over time. The red curve shows the environmental temperature. The blue dots show

the PE values over time, complete with the green linear trendline to estimate aging.




Figure 6.16: Example module 127 aging plot. Example of an aging plot from channel 40 from Oksuzian using Ehrlich’s PE results [7]. The x-axis on this plot is slightly shifted compared to that of Figure 6.13, but uses the same units of time fraction. Oksuzian’s time fraction is calculated as the number of years with respect to 7/8/2021, which is the date where the first ‘good’ run was started. The blue points represent the light yield values over time, normalized such that the initial light yield is 1. The green trendline is the linear fit applied to these PE values. The red curve represents the temperature, as recorded by a Bluetooth temperature sensor mounted to a staircase of Wideband near the module storage area. Finally, the black dots represent the SPE values for this particular channel over time.

The black dots in Figure 6.16 track the SPE values of a particular channel over time. Because PEs are quantized units, the single PE distributions should always have the same shape over time, once corrected for environmental factors like temperature variations. This implies that the SPE values extracted from the data in Wideband should be constant over time. Looking at the black dots in Figure 6.16 above, there appears to be a region in which the SPE value calculated in this channel suddenly increases. The transition occurs around the x-axis value of 0.7 in Figure 6.16, which corresponds to the date of 3/17/2022. This sudden increase of SPE value was observed in all channels in module 127. This begs the question: what happened on 3/17/2022?

Upon review of messages exchanged on Slack in the Wideband channel, it was discovered that on 3/17/2022, there were fans installed near the FEBs to aid with cooling the boards. The FEB boards are encased in aluminum enclosures that are only open on two ends; the fans were set up to blow air through the enclosures. At this point, heat sinks had not been installed on the FEBs. As a result of installing the fans, the temperature of the FEBs was reduced by 20 degrees Celsius [8]. A plot of average FEB temperature over time is shown in Figure 6.17. Much like in the SiPMs of the CMBs, changes in temperature could also affect the gain of electronics on the FEBs. According to

Ehrlich, a 30 degree Celsius change in environmental temperature corresponds to a 0.1V change in bias voltage of the FEBs, based on measurements taken at University of Virginia. There are temperature-dependent resistors on the FEBs as part of the circuit that regulates bias voltage. With this in mind, it is plausible that the installation of the fans changed the bias voltage of the FEBs, leading to an increase in the SPE value.

The second hint of suspicion about data quality arose while monitoring any other temperature changes that may have occurred throughout the lifetime of the Wideband test bench. Before the sudden increase in SPE values was discovered, the only handle on temperature that could be plotted from the reconstruction files was temperature as recorded by the CMBs. Recall that each CMB includes a temperature sensor that records data once per spill and that CMBs are mounted to the end of each dicounter in the module. Because CMBs are mounted directly on the module, the plastic scintillator that comprises the bulk of the module acts as a thermal mass. The modules are sitting in the Wideband lab, so generally the temperature of the modules fluctuate with the temperature of the environment. Since the increase in SPE values is related to temperature changes of the FEBs, an additional variable was added to the processing code that allows for the FEB temperature to also be recorded and plotted. Plots for the average temperatures for each run recorded by a random CMB and FEB0 are shown in Figure 6.17.

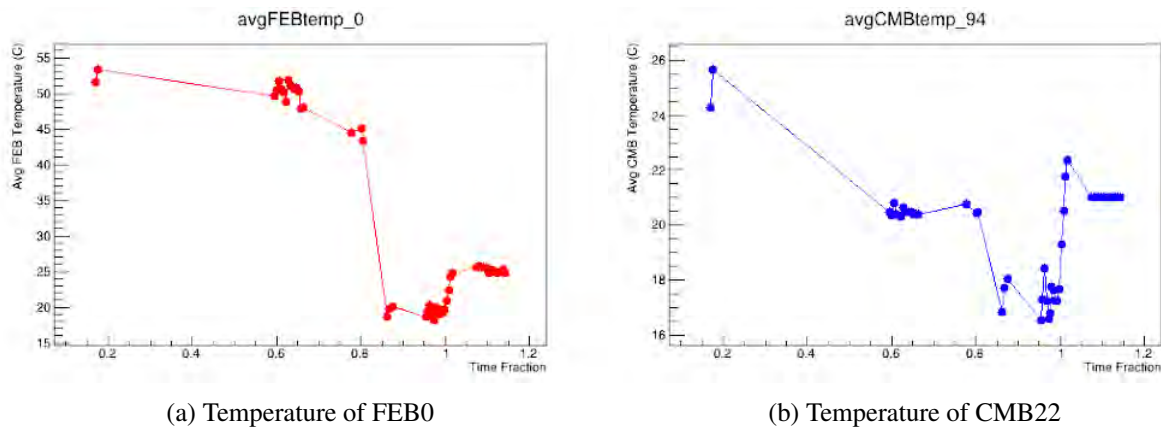


Figure 6.17: Average temperatures for an FEB and a CMB of module 127. Figure a) tracks the average FEB temperature in degrees Celsius of a given channel over time. Figure b) tracks the average CMB temperature in degrees Celsius of a given channel over time. The x-axis on both of these plots is the time fraction in number of years with respect to 5/7/2021, same as our aging plots.

Figure 6.17 shows large variations recorded by both the CMBs and the FEBs. In both plots, there is a sharp decrease in the average temperature observed just after the time fraction of 0.8. This date matches the date of the sudden SPE increase observed in Oksuzian's plots, 3/17/2022. So, the drop of roughly 20 degrees Celsius in the FEB temperature corresponds to the addition of fans near the FEB enclosures. At the same time, the air conditioning and heating in Wideband was

turned off with the idea of letting the environment to cool as much as possible and measure the effect of temperature on the light yield of the module. The expected drop in ambient temperature is confirmed by the measured average CMB temperature and the hall temperature as monitored by a Bluetooth sensor installed in the staircase. The temperatures shown in Figure 6.17 is representative of the temperatures that were recorded by the other CMBs and FEBs in use, as they were similar upon visual inspection.

The second sudden jump in temperature around $x = 1.0$ in Figure 6.17 is similarly suspicious. The amount of sudden temperature increase around this time is concerning due to the effects that temperature changes can have on the gain and other properties of sensitive electronics on both the CMBs and the FEBs. A closer review of the Wideband Slack channel uncovered that on 5/11/2022 the building air conditioning was turned back on. Fans were still cooling the FEBs at this time. In addition, the plots of the average CMB temperature for the runs taken at this time show that the temperature was not properly updating.

With these two features of the data in mind, we split the Wideband data collected until this point into three periods: before $x = 0.8$, between $x = 0.8$ and $x = 1.0$ and after $x = 1.0$. The first period refers to data collected before 3/17/2022. The second period refers to data collected between 3/17/2022 and 5/11/2022. The third period refers to data collected between 5/11/2022 and 6/24/2022. The aging analysis was performed on each subset of data to look for the effect of these changes on aging rates. First, the aging rates obtained from analyzing periods 2 and 3 are presented, as these are periods of which data quality is under suspicion.




Figure 6.18: Period 2 aging rates by channel for module 127. Distributions of the aging rates estimated in all channels of module 127, in percentage of light yield lost per year by channel number for period 2, spanning between March and May.

Figure 6.18 shows the 2-dimensional plot of aging rates by channel number using only period 2 data for analysis. The upper and lower layers of the module do not have the same obvious difference in average aging rate that was observed before splitting the data into periods. However, the aging rates calculated from this period of data appear to have large channel-to-channel variations with many channels showing positive aging rates.




Figure 6.19: Period 3 aging rates by channel for module 127. Distributions of the aging rates estimated in all channels of module 127, in percentage of light yield lost per year by channel number for period 3, spanning between May and the end of June.

Figure 6.19 shows the 2-dimensional plot of aging rates by channel number using only period 3 data. The two halves of the graph have a large difference in average aging rate and a large range for the aging rates, likely due to the fact that this data only spans a short period of time and are affected by large statistical fluctuations. After excluding data from periods 2 and 3, the aging rates for period 1 only are discussed in the next section as the final aging analysis for module 127. We decided to present these as the final result of the aging analysis because they are taken with the most consistent conditions.

6.4.3 Final Aging Analysis

In this section, aging results using only high quality data from period 1 are presented. Figure 6.20 shows the 2-dimensional plot of aging rates by channel number using only period 3 data.

In Figure 6.20, the aging rates from period 1 are shown by channel number. The lower and upper channels of module 127 have a similar average aging rate, and small channel-to-channel fluctuations. The few channels that result in a positive aging rate are channels located near the




Figure 6.20: Period 1 aging rates by channel for module 127. Distributions of the aging rates estimated in all channels of module 127, in percentage of light yield lost per year by channel number for period 1, spanning between July 2021 and the end of February 2022. This distribution includes all ‘good’ runs collected at Wideband before 3/17/2022.

edge of the module which have more sparsely populated aging plots. The 1-dimensional histogram




Figure 6.21: Period 1 aging rates for module 127. 1-dimensional distribution of the aging rates estimated in all channels of module 127, in percentage of light yield lost per year for period 1, spanning between July 2021 and the end of February 2022. This distribution includes all ‘good’ runs collected at Wideband before 3/17/2022.

of aging rates from period 1 is shown in Figure 6.21. The underflow channels come from edge channels where the aging plots are too sparsely populated to be fit. The final estimation of the CRV

aging rate is extracted from Figure 6.21 using the mean and standard deviation as $8.06\% \pm 0.33\%$ per year. The error is given by the standard error on the mean.

6.5 Conclusions

An aging estimation for the CRV was performed using data collected from the Wideband test bench between May 7, 2021 and June 30, 2022. Initially, the study was performed using all of the data between May 2021 and June 2022. However, the results indicated that the upper and lower halves of the module were aging at different rates, which prompted a closer review of data quality. The data was split into three periods after discovering environmental changes that affected test stand electronics: Period 1 from 5/7/2021 to 3/16/2022, Period 2 from 3/17/2022 to 5/10/2022, and Period 3 from 5/11/2022 to 6/30/2022.

Period 1 contains the only high quality data for module 127. Aging analysis using only Period 1 data yields an estimated aging rate of $8.06\% \pm 0.33\%$ per year. This aging rate is higher than expected, when comparing to the aging rate of other experiments at FNAL that use similar scintillator-based detectors. As a result, the investigation discussed next in Chapter 7 was launched to determine whether or not one individual component of the CRV bars could be the main contributor to the high estimated aging rate during this time.

The Wideband test bench has been continuously under development to improve the quality of data collection and use other CRV modules to estimate the aging rate of different sides of the CRV. Since this investigation, many changes have been made to the test bench such that the same configuration used with the module 127 studies above is no longer in use. These changes include adding lead bricks above and below the module that is being measured to improve trigger purity by reducing the possible incident angles that cosmic rays may have when traveling through the CRV modules. When analyzing data for signs of aging, the configurations of the datasets should have the same environmental and electronics setup to control for any variables that may impact light yield. New aging estimation campaigns are underway as the test bench improves and data collection configurations change [5].

References

- [1] V. Senchishin, V. Koba, O. Korneeva, V. Seminozhenko, V. Kovtun, I. Zalubovsky, I. Chirikov-Zorin, J. Budagov, F. Markley, and G. Bellettini, *New Radiation Stable and*

Long-Lived Plastic Scintillator for the SSC, Tech. Rep. (Fermilab, 1993)
FERMILAB-TM-1866.

- [2] L. Aliaga, *et. al.* (The MINERvA Collaboration), Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip. **743**, 130 (2014).
- [3] K. Abe, *et. al.* (The T2K Collaboration), “Scintillator Ageing of the T2K Near Detectors from 2010 to 2021,” (2022), arXiv:2207.12982 [ins-det].
- [4] P. Feng, “Polymer Aging of Plastic Scintillators,” (2015-07-01),
https://indico.fnal.gov/event/23044/contributions/193049/attachments/132238/162386/scintillator_aging_feng_2015.pdf.
- [5] Y. Wu, “Mu2e CRV Module Aging Rate Studies,” (Mu2e Biweekly Meeting 2023-09-14), mu2e-doc-46508-v1.
- [6] Hamamatsu, “MPPC S13360-2050VE,” (2016), datasheet, Available at:
https://www.hamamatsu.com/us/en/product/optical-sensors/mppc/mppc_mppc-array/S13360-2050VE.html.
- [7] Y. Oksuzian, (private communication).
- [8] R. Ehrlich, (private communication).

Chapter 7 Measuring Aging in Titanium Dioxide Cladding Material

After obtaining the first CRV aging estimate from the test stand at Wideband as discussed in Chapter 6, discussions began about what could be the root cause of the high aging rate that is observed for the Mu2e CRV. Aside from electronics, there are three main components to a CRV counter: the bulk scintillator, inner wavelength shifting fibers, and an outer coating of reflective cladding. If one of these three components is aging rapidly, perhaps it could explain the high aging rate estimation from Wideband and a remedy could be prescribed before the CRV is assembled. As described in Chapter 4, the counters that are used in the CRV are fabricated at FNAL by the FNAL-NICADD collaboration facility. For this study, I worked with a small group of Fermilab scientists who operate the FNAL-NICADD facility and were part of fabricating the scintillating bars for the Mu2e CRV counters. Dr. Alan Bross served as the leader of this investigation and Dr. Anna Pla and technician Brian Leung provided substantial support throughout this phase of study.

As described in Chapter 4, the scintillator bars for CRV counters are extruded using specific die shapes, the inner channels for the fibers are co-extruded with dry nitrogen injections, and the cladding layer is also co-extruded as a 0.25 mm thick outer layer on the surface of the scintillator bars. The cladding is comprised of 30% titanium dioxide [1]. The wavelength shifting fibers are inserted into the channels after the extrusion process is complete. Each of these three components - the scintillator, the cladding, and the fibers - are plastic-based, and have the potential to degrade over time. The TiO_2 reflective cladding is designed to increase the light yield within the CRV counters by preventing scintillation light from exiting the surface of the bars. Light yield is a common metric for measuring the performance of scintillating detectors. Using cladding with high reflectivity, like titanium dioxide, increases the mean number of bounces that scintillation light takes from the surface of the CRV counters since the light cannot exit through the reflective coating. If there is a decrease in the reflective properties of TiO_2 over time as a result of aging, the mean number of bounces that scintillation light takes within the CRV bars may decrease, resulting in a decrease in light yield.

Part of this investigation was also driven by quality control monitoring at Fermilab. The raw materials for the scintillator and the TiO_2 cladding layer are purchased from third parties with various quality control standards, so the team at FNAL monitors samples from each purchase batch to watch for changes over time from the manufacturers. Specifically, polystyrene is the plastic

base used in the FNAL-NICADD scintillator, and visual differences in color tint and cloudiness have been observed in extruded samples of polystyrene throughout the recent past. Properties like reflectance and transmittance of light and the light yield in polystyrene can have large effects in the detector response for detectors like the CRV. Much effort has been put into finding reliable, pure, and stable sources for polystyrene and titanium dioxide from both domestic and international markets. For each batch of polystyrene that is received at Fermilab, a small sample is extruded and analyzed for optical properties. For each batch of titanium dioxide that is received at Fermilab, the manufacturer typically sends an extruded sample, called a coupon, which is analyzed for optical properties and stored in a binder with other titanium dioxide for record-keeping posterity.

This chapter will detail the aging study that was performed to evaluate the aging rate of the reflective titanium dioxide cladding on the surface of the CRV bars. This study was performed using two sets of samples, the titanium dioxide coupons directly from the manufacturer and short pieces of FNAL-NICADD extruded scintillator bars. For the extruded samples, the scintillator bulk was machined away to isolate the titanium dioxide cladding surface for analysis.

7.1 Lab 6 Aging Measurement Strategy

For this investigation, the strategy for measuring the aging of titanium dioxide is to track the reflectivity of each sample over time and monitor each sample for changes and trends. The reflectance data for each sample was collected using a spectrophotometer, about once per week during 2022 and more sparsely afterwards. By using the reflectivity data for each sample, scatter plots illustrating the reflectivity of each sample over time are created and studied by plotting reflectance versus time.

A HunterLabs UltraScan VIS Spectrophotometer was used to measure the reflectivity of each sample [2]. This instrument measures reflectivity using an optical sphere with an aperture and a dual beam xenon flash light source. The optical integrating sphere is six inches in diameter and is coated with Spectrafect [3], which diffuses the light from the lamp. Each sample is placed flush to the aperture in the sphere using a clamp attached to the instrument. The instrument is then triggered from within the EasyMatchQC software and illuminated with light, where the light either reflects from the sample surface or transmits through it. A lens positioned nearly perpendicular to the sample surface then collects the reflected or transmitted light and directs it to a diffraction grating where the light can be separated into components. These components are then measured by dual diode arrays and converted into usable data [2]. For this study, only the reflectance mode was used to collect data.

Since the wavelength shifting fibers inside of the CRV counters absorb light between 375 and 475 nm, this range of wavelengths is the most important range to monitor for changes. The

HunterLabs UltraScan VIS Spectrophotometer nominally operates between 360 and 780 nm [2], which is sufficient for our wavelength range of interest. When the aging plots are created to monitor reflectance over time, each plot represents data collected for a designated wavelength of light. Aging plots are then created for wavelengths 380 nm to 450 nm, in increments of 10 nm. The HunterLabs UltraScan VIS Spectrophotometer comes with its own data acquisition software, EasyMatchQC, which allows the user to take and record reflectance data [2]. The data is recorded in a table on the EasyMatchQC GUI, so after recording the data with EasyMatchQC, the data was transferred into Microsoft Excel spreadsheets for easier organization and use with analysis codes.

7.1.1 Analysis Codes

All of the analysis code used for this investigation was written from scratch for the express purpose of organizing this reflectance data into arrays by wavelength and then plotting the data over time in two stages. The first stage of the analysis codes convert Excel spreadsheet data into ROOT histogram data. Each file contains labeled reflectance spectra, in histograms in ROOT that now look like the spectra that could be seen using the EasyMatchQC software.

The second stage of the analysis codes generate the aging graphs for further analysis. During this stage, the time axis is created, which is calculated in fractions of years since the first measurement. This time axis is directly used as the x-axis of the aging plots, with reflectance on the y-axis. A linear fit is applied to each plot where the slope represents the change in reflectance over the period of one year, since the x-axis is recorded in fractions of years.

7.1.2 Calibration and Standardization of the Spectrophotometer


The user manual for the HunterLabs UltraScan VIS Spectrophotometer details how to standardize and calibrate the device [2]. It is suggested that the spectrophotometer be standardized every eight hours of operation, between hardware changes, or in response to environmental changes, like a temperature change. Standardization sets the maximum and the minimum of the photometric scale used to record reflectance data. To standardize the spectrophotometer, both a white tile and a light trap are used. The light trap sets the bottom of the scale and the white tile sets the top of the scale for the reflectance mode of the instrument [2]. Standardization tools like the white tile and light trap are included with the instrument. On-screen messages in the EasyMatchQC GUI prompt the user through the standardization procedure. Before placing the light trap, white tile, or any other sample in the aperture of the spectrophotometer, a brush or disposable cloth pad was used to clean off any dust that could interfere with the measurements. The spectrophotometer was also turned on about an hour before taking any data or performing any standardization tasks, as the instrument needs time to warm up the lamp.

In addition to standardization with the white tile, a green tile can be used to monitor the long-term performance of the spectrophotometer. This is referred to in the manual as the "Repeatability Test" [2]. It is suggested in the manual that the green tile test should be performed once per week, however, the technician in Lab 6 who introduced me to the instrument suggested that this test be performed once per month. I performed the green tile test about once per month during my time at Fermilab in 2022. To perform the green tile test, a green tile is placed in the optical sphere aperture. The spectrophotometer then illuminates the green tile and gives an output of XYZ tristimulus color values. The correct values of each XYZ value is listed directly on the green tile itself. EasyMatchQC then compares the measured XYZ color values of the green tile to their nominal magnitude. If the green tile reading varies by more than ± 0.15 XYZ units from the values given on the tile, then the instrument fails the test [2]. Each of my green tile tests passed the comparison to nominal values. Response from the green tile is temperature dependent, and operates well within the typical range of room temperatures. This test is one way of monitoring for any change in data as a result of environmental factors in the room where the instrument is kept.

One more test that was performed to monitor the standardization and calibration of the spectrophotometer was to measure well-known samples after standardizing the instrument and visually check the spectra for large unexpected deviations. Two standards are used, a sample from the NOvA experiment, named N-27-09-NC, and a sample of GoreDRP, a highly reflective and flexible material. These spectra are compared over time using custom-built analysis code that looks for changes in the spectra. After writing the analysis code to compare the standard spectra, I began taking three measurements of each standard before and after collecting sample data. This allowed me to calculate the average spectrum from each set of three standard measurements, then compare the averages before and after taking data to look for any changes as a response to environmental factors.

The NOvA standard was created in October of 2009 and has been used regularly with the Hunter Labs UltraScan VIS Spectrophotometer ever since. In fact, the instrument was purchased by the NOvA experiment to monitor various properties of their detectors [4]. This standard is a small, circular piece of the reflective cladding from the NOvA experiment that was cut and machined after the experiment was decommissioned. Photographs of the front side and of the storage bag indicating the production date are shown in Figure 7.1. An example reflectance spectrum of the NOvA N-27-09-NC sample is shown in Figure 7.2. The steeply curved part of the spectrum from about 380 to 450 nm is the section that was used for visually checking the calibration of the spectrophotometer instrument by looking for unexpected curves in that region.

GoreDRP is a highly reflective, proprietary material that can be used to create reflective surfaces. This material is not extruded in-house like the titanium dioxide cladding used in the NOvA and Mu2e experiments. Instead, samples of GoreDRP were ordered to Fermilab for analysis.



(a) Front side of NOvA N-27-09-NC standard

(b) Storage bag for NOvA N-27-09-NC standard

Figure 7.1: Photographs of the NOvA N-27-09-NC standard. Figure a) shows the front side of the standard, which was placed against the spectrophotometer aperture for measurement. Figure b) shows the storage bag for the NOvA standard. The bag is labeled with the date that the sample was produced, the name, and the location from which it was taken from the NOvA detector.




Figure 7.2: Example NOvA N-27-09-NC reflectance spectrum. The reflectance of the standard surface is shown on the y-axis in percentage versus wavelength of light on the x-axis. Note that the y-axis begins near 20% reflectance and extends to nearly 100% reflectance. The title of the plot indicates that the measurement was taken before measuring samples on 11/22/22.

GoreDRP is a soft and flexible material, so the sample standard that was prepared for use with the spectrophotometer was created by taking a square piece of GoreDRP and fixing it to a rigid plastic disk with adhesive. In this way, we can create a smooth, flat surface to use with the spectrophotometer and measure the reflectance of the surface of GoreDRP. Photographs of the front and back sides of the GoreDRP standard are shown in Figure 7.3. An example reflectance spectrum




Figure 7.3: Photographs of the GoreDRP standard. Figure a) shows the front side of the GoreDRP standard, which was placed against the spectrophotometer aperture for measurement. Figure b) shows the back of the GoreDRP standard, which is labeled with the name of the material.

from the GoreDRP standard is shown in Figure 7.4. This standard was preferred for monitoring the stability of the instrument since it is so highly reflective across all wavelengths. The spectrum is relatively smooth and near 100% reflectivity across the entire wavelength range, so any deviations would be visually easy to find if they were large enough.

7.1.3 *Stability and Aging of NOvA and GoreDRP Standards*

Before looking for results in the two sets of new samples, the Scintillation R & D team decided that the best course of action would be to analyze aging plots for the two sets of standards, NOvA N-27-09-NC and GoreDRP. The NOvA standard was used throughout the entire investigation, starting in February 2022 when this study began. The GoreDRP standard was used beginning in August 2022 after a suggestion during a Scintillation R & D meeting. In addition to monitoring the reflectance of the standards over time for signs of aging, repeated measurements of the standards were used to check the stability of the spectrophotometer instrument. Beginning in September 2022, both standards were measured three times before measuring the samples and three times after





Figure 7.4: Example GoreDRP reflectance spectrum. The reflectance of the standard surface is shown on the y-axis in percentage versus wavelength of light on the x-axis. Note that the y-axis begins near 96% reflectance and encompasses only between 96% and 100% reflectance, indicating the high reflectivity of the standard across all wavelengths. The title of the plot indicates that the measurement was taken before measuring samples on 11/22/22.


measuring the samples. Prior to September, the stability of the instrument was not in question, but during a meeting around this time, the age of the instrument was discussed as a point of concern. By calculating the standard deviation for the three repeated measurements for each standard both before and after measuring the samples, the instrument can be monitored for stability in addition to the standardization tests performed with the Hunter Labs software. For the aging plots of both sets of standards, the standard deviation of the repeated measurements is also used as the error on the reflectance measurements.

Let's discuss how time is plotted on the the x-axis before we present any plots for this investigation. Much like the aging study using the CRV modules in Wideband, these plots are made with ROOT, so it is beneficial to define time as a decimal number, rather than a date. Again, here time is expressed in fractions of a year, where one full year equals 1 and no time elapsed equals zero. Using this metric for time on the x-axis allows for a direct extrapolation of the slope from a linear fit to the aging rate from plots of reflectance over time. In Figure 7.5, the standard deviation in percent of the repeated measurements of the NOvA and GoreDRP standards are plotted over time. The open circles correspond to the measurements taken before measuring the samples and the filled circles correspond to the measurements taken after measuring the samples. On the x-axis, the beginning date is 10/11/22 and the ending date is 6/23/23. Plots like the ones shown below are made for one wavelength at a time for wavelengths between 380 nm and 450 nm in increments of

10 nm. This wavelength range was chosen since this is the regime where the wavelength shifting fibers absorb light from the scintillator. The plots below represent data collected at 390 nm for both standards. Only one wavelength is shown here as an example.



(a) Standard deviation of repeated NOvA standard measurements over time



(b) Standard deviation of repeated GoreDRP standard measurements over time

Figure 7.5: Standard deviation of repeated standard measurements over time at 390 nm. Figure a) shows data from the NOvA standard and Figure b) shows data from the GoreDRP standard. In both plots, the open circles represent repeated measurements taken before measuring any samples with the spectrophotometer and the filled circles represent measurements taken after measuring samples with the spectrophotometer. The red lines indicate the mean of the standard deviations before sample measurements and the blue lines indicate the mean of the standard deviations after sample measurements.

From Figure 7.5, information about the stability of the machine over time may be inferred.

For both sets of standards, the standard deviation of repeated measurements is low across all measurement dates. There is also no quantitative difference between the standard deviations of repeated standard measurements before and after measuring samples in both plots. Both of these observations are evidence that the spectrophotometer is stable. The standard deviation of repeated measurements is consistently low over time and there is no trend with respect to timing of repeated measurements within a data collection session.

Now, with confidence in the stability of the spectrophotometer, we can now produce plots of reflectance over time and look for aging trends. With the data collected on the NOvA and GoreDRP standards, aging plots for the standards can be produced. For the standards, no aging is expected. Similarly to the standard deviation plots in Figure 7.5, the aging plots are produced for one wavelength at a time. Figure 7.6 shows reflectance over time for the NOvA standard at 390 nm.




Figure 7.6: NOvA standard reflectance over time at 390 nm. The reflectance of the standard surface is shown on the y-axis in percentage versus time in fraction of years on the x-axis. The x-axis extends from 3/17/22 until 6/23/23. The trendline is a linear fit where the slope of the fit indicates the aging rate at this particular wavelength.

Figure 7.6 is an aging plot for the NOvA standard. This plot indicates that the reflectance of the surface of the NOvA standard increases over time rather than decreases. Though the y-axis of the plot only spans two percent in reflectance, the positive trend is generally consistent over the entire length of observation. Each point on this plot is an average of the repeated reflectance measurements collected on a given date. The errors on each point in this plot are the standard deviation of the repeated measurements for that date. For dates with only single measurements, the standard deviation from the earliest repeated measurement date is used for the error. The slope of the trendline indicates that the NOvA standard gains +0.75% per year. This was a surprising result

since the NOvA standard is so old, it was expected that aging had already occurred. Based on the error on the slope of the trendline, the aging rate for the NOvA standard is not consistent with zero. Analogous plots may be made for each wavelength in our wavelength range of interest, from 380 nm to 450 nm in increments of 10 nm, from which an aging rate may be found for each wavelength. A summary plot for the NOvA standard showing the aging rate for each wavelength extracted from the slope of the aging plot at each wavelength is shown in Figure 7.7.




Figure 7.7: NOvA standard aging rate summary plot. Aging rates for the NOvA standard at different wavelengths of interest. Aging rates in loss of reflectance per year, extracted from the slope of the trendline of each aging plot, are shown on the y-axis. Wavelength, in nm, is shown on the x-axis.

Figure 7.7 shows that the aging rates are positive across all wavelengths for the NOvA standard. The errors shown for each point on the summary plot come from the fit parameter errors. This result was unexpected and called into question how well the NOvA standard is truly understood or has changed over time. Beyond this investigation, reflectance data from the NOvA standard has not been used in any other recent studies and has not been consistently monitored over time. For this reason, it was decided that the NOvA standard data should not be considered for aging studies, though it was instrumental in the calibration and standardization of the spectrophotometer.

Though the aging results from the NOvA standard produced surprising results, the same analysis can be performed for the GoreDRP standard. A plot of reflectance over time for the surface of the GoreDRP standard is shown in Figure 7.8.

Figure 7.8 is an aging plot for the GoreDRP standard at 390 nm, analogous to Figure 7.6 for the NOvA standard. Similarly to the NOvA standard aging plot, each point on this plot is an average of the repeated reflectance measurements collected on a given date. The errors in Figure 7.8 are the standard deviation of the repeated measurements for that date. For dates with only single




Figure 7.8: GoreDRP standard reflectance over time at 390 nm. The reflectance of the standard surface is shown on the y-axis in percentage versus time in fraction of years on the x-axis. The x-axis extends from 8/31/22 until 6/23/23. The trendline is a linear fit where the slope of the fit indicates the aging rate at this particular wavelength.

measurements, the standard deviation from the earliest repeated measurement date is used for the error. Figure 7.8 indicates no observable aging for the GoreDRP standard at this wavelength. There is no consistent trend in reflectance during the measurement period, with seemingly random distribution in the measurements within a 0.5% range of surface reflectance. In addition, the slope of the linear trendline and its associated uncertainty are consistent with a slope of zero, or no aging. A summary plot for the GoreDRP standard showing the aging rate for each wavelength extracted from the slope of the aging plot at each wavelength is shown in Figure 7.9.

Though not all of the aging rates across all wavelengths for GoreDRP are consistent with zero, as shown in Figure 7.9, most wavelengths have an aging rate that is consistent with zero. The rates that are not consistent with zero are very small and are not in the same direction as to appear to have a trend. Projecting the aging rates into a one-dimensional histogram gives a mean aging rate of $0.03\% \pm 0.17\%$ surface reflectance lost per year, consistent with zero. The error on the mean aging rate is taken as the standard deviation of the histogram. This result of no observed aging in the GoreDRP standard reinforced confidence in the accuracy and long term stability of the spectrophotometer as it is used to collect data on the cladding coupon and extrusion samples.

7.2 Description of Samples

With the analysis strategy laid out above, this section will describe the two sets of samples that were used in this investigation: the manufacturer-produced TiO_2 coupons and the FNAL-





Figure 7.9: GoreDRP standard aging rate summary plot. Aging rates for the GoreDRP standard at different wavelengths of interest. Aging rates in loss of reflectance per year, extracted from the slope of the trendline of each aging plot, are shown on the y-axis. Wavelength, in nm, is shown on the x-axis.

produced and prepared extrusion samples. The two standards will also be described: the NOVA N-27-09-NC standard and the Gore DRP standard.

7.2.1 *Manufacturer-Produced Titanium Dioxide Cladding Coupons*


As mentioned in the introduction to this chapter, each shipment of titanium dioxide pellets that arrive at Fermilab typically come with a set of manufacturer-produced sample extrusions of the cladding. These small sample cards are called "coupons". The technicians at Fermilab have been historically storing these coupons to preserve them for future analysis. The coupons are stored in a large binder and each set of coupons is labeled with the name of the manufacturer, a serial number, a lot number, the type of plastic used to mix with the TiO₂ for extrusion, and the date that the coupon was made. Photographs of a coupon are shown in Figure 7.10.

The front side of each coupon is flat and the back side of each coupon is stepped such that each coupon has two, sometimes more than two, thicknesses. Most coupons have two thicknesses where the thin side of the coupon is roughly 1 mm in width and the thick side of the coupon is roughly 2 mm in width, as shown in Figure 7.10c. Recall that the TiO₂ cladding layer that is extruded onto the surface of the CRV bars is specified to be 0.25 mm, much thinner than the coupons from the manufacturer. The front, flat side of the coupons is a polished surface, so this is the side that was measured for the reflectance data used in the forthcoming aging investigation. The back of the coupons was sometimes textured, which would interfere with measurements, and therefore was not used to collect reflectance data. An example reflectance spectrum from an older titanium



(a) Front side of a TiO₂ coupon

(b) Back side of a TiO₂ coupon



(c) Side view of a TiO₂ coupon

Figure 7.10: Photographs of an example TiO₂ coupon. These coupons are from the manufacturer Clariant. The manufacturer name, serial number, part number, lot identifier, the type of resin or plastic used to mix the sample, and the date of manufacture are listed on the label. Figure a) shows the front of the coupon, figure b) shows the back of the coupon, and figure c) shows the side of the coupon.

dioxide cladding coupon is shown in Figure 7.11. The set of coupons that this example spectrum was collected from is called KE12291 with a manufacturing date of 6/21/16.

When I arrived at Fermilab and began working with the Scintillator R & D team, I began going through the binder of coupons at Fermilab and measuring the reflectance of each coupon in the binder. The oldest set of coupons in the binder are dated to the year 1999. All other sets


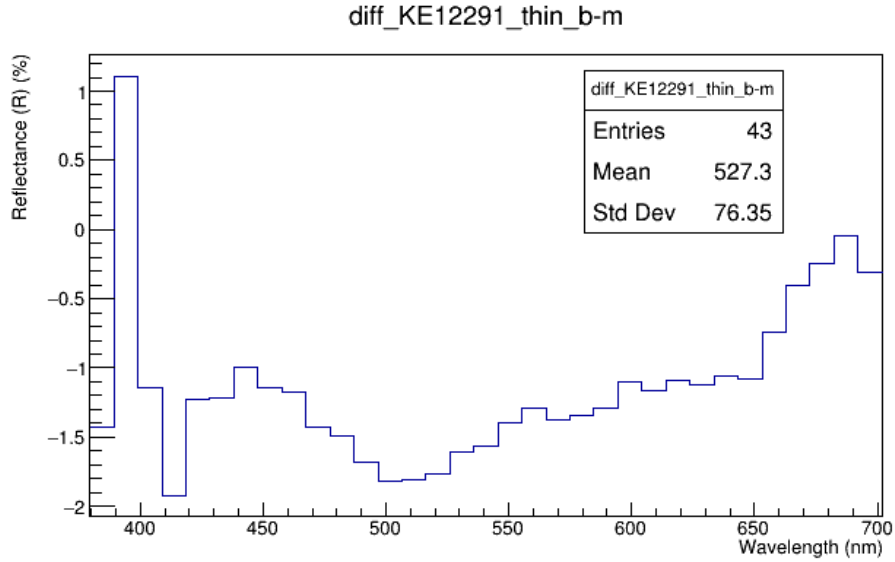


Figure 7.11: Example TiO₂ KE12291 reflectance spectrum. The set of KE12291 coupons was manufactured in June 2016. The reflectance of the sample surface is shown on the y-axis in percentage versus wavelength of light on the x-axis. Note that the y-axis begins near 10% reflectance and extends to over 90% reflectance.


of coupons in the binder are dated between 2014 and 2022. Upon visual inspection, some of the coupons appeared yellow, an effect that tends to occur over time in different types of plastics [5]. Yellowing of plastic over time is a concern with regard to aging because a plastic appearing to look yellow is reflecting yellow light in the 570 to 590 nm range in wavelength. The wavelength shifting fibers within the CRV counters absorb light from the scintillator bulk anywhere between 350 and 500 nm and reemits light anywhere between 450 to 650 nm [6]. If the TiO₂ cladding is yellow on the inner surface of the CRV counters, it could change the way that light is reflected off of the surface because the fibers reemit light in the yellow range.

The reason for remeasuring the reflectance of each coupon available was to compare with older reflectance data that was taken over the past few years. Between 2016 and 2020, many samples were measured when work was being performed to decide what kind of cladding to use for Mu2e. If my measurements in 2022 showed significant deviations from measurements in the years prior, it would indicate aging of the cladding. Two plots are shown in Figure 7.12, both plots show the difference between the older and newer reflectance spectra. Figure 7.12a shows the difference between the average spectrum of KE12291 coupons between 2016 reflectance data and 2022 reflectance data.

The differences in the spectra of both samples shown in Figure 7.12 are relatively small, with the largest difference only resulting in a 2% change for one particular wavelength bin in the spectrum. In addition to the small magnitude of the differences in reflectance, the time that has elapsed between the data points that are being compared is large. For the USPN035597 set



(a) Difference in reflectance spectra of KE12291 between 9/23/16 and 2/8/22



(b) Difference in reflectance spectra of USPN035597 between 11/4/20 and 2/7/22

Figure 7.12: Difference in reflectance spectra over time for two samples. Figure a) shows the average difference in the reflectance spectrum of coupon set KE12291 between 9/23/16 and 2/8/22. Figure b) shows the average difference in the reflectance spectrum of coupon set USPN035597 between 11/4/20 and 2/7/22. Note the y-axes on these plots: both axes roughly encompass $\pm 2\%$ of reflectance across the entire wavelength range.

of coupons, there is almost 1.5 years between the two spectra that are being compared. For the KE12291 set of coupons, there is nearly 5.5 years between the two spectra that are being compared. Observing these small changes over a large amount of time in not only these two samples, but other older sets of coupons, led to insightful discussions about aging over time. It was decided that the

older sets of coupons were not the best choice of samples for long-term aging rate calculations. If no signs of aging could be seen by comparing an old spectrum to a new spectrum, as in Figure 7.12, then measuring the older samples many times over the course of 2022 is also not likely to result in signs of aging.

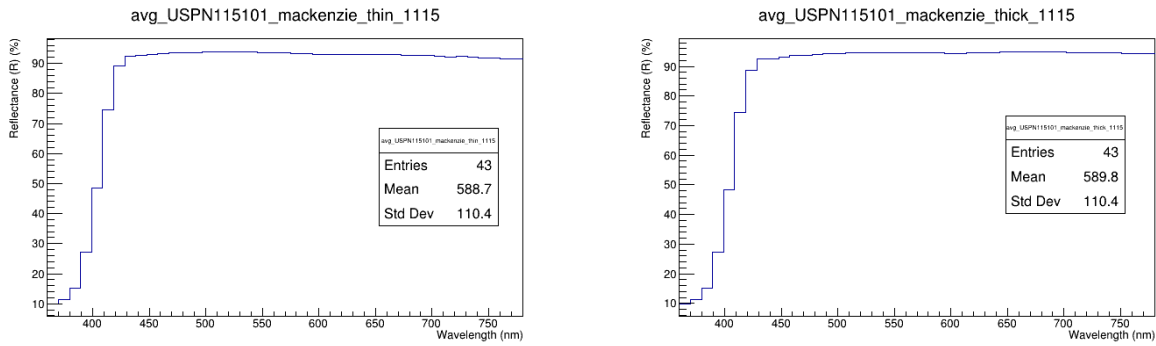
Early on in this investigation, there was a discussion in a Scintillation R & D meeting about how valuable it would be to obtain a new set of coupons and monitor them for aging over the course of the year. Typically, when scintillators age, it has been observed that aging occurs rapidly just after production, and then the age levels out to a small aging rate over time. By obtaining and monitoring a new set of coupons, they could be monitored for this initial rapid aging or any other trends. A new batch of titanium dioxide was ordered in early 2022 and conveniently arrived in early March with a new set of coupons, identified by lot number USPN115101. This is the same as the titanium dioxide used in Mu2e CRV counter production. The set of USPN115101 coupons are the focus on the coupon side of this aging investigation and were measured once per week from early March to early December 2022 to collect reflectance data for this aging investigation. Example reflectance spectra from each side of a USPN115101 coupon is shown in Figure 7.13.

Comparing the two spectra shown in Figures 7.13a and b, there is hardly any difference between the spectrum measured on the thin side versus on the thick side of the manufacturer-produced cladding coupons. The difference is shown in Figure 7.13c. The thin side of the USPN115101 coupons is closer to the thickness of the cladding surface that will be on the Mu2e CRV counters. The largest difference between the thin and the thick sides of the USPN115101 coupons is only a 2% effect and also occurs at higher wavelengths which are not absorbed by the wavelength-shifting fibers within the CRV counters. Based on these comparisons, using thicker cladding around the CRV counters does not improve cladding surface reflectance.


7.2.2 Extruded Titanium Dioxide Cladding Surface Samples

In addition to measuring and monitoring the manufacturer-produced TiO₂ cladding coupons, the Scintillation R & D group also decided later in the year that it would be interesting to measure the reflectivity of cladding that had been extruded at Fermilab. To do this, Dr. Alan Bross took extrusions from Mu2e and machined the scintillator bulk away using a diamond saw. This isolates the reflective titanium dioxide cladding layer from the scintillator bulk to try to prevent the scintillator from interfering with the reflectivity measurements. These samples were created in early September 2022 and one of the extrusion sample pieces is shown in Figure 7.14.

The front side of each extrusion sample is the outer surface of the CRV counter and the back side is the inner surface of the counter, which is the interface between the cladding and the scintillator. The extrusion samples were machined to be as thin as possible without damaging the TiO₂ cladding layer, though there is still some plastic scintillator still attached to the back side of the



(a) Example spectrum from the thin side of USP115101 (b) Example spectrum from the thick side of USP115101
diff_USPN115101_m_thick-thin




(c) Difference between the thick and thin sides of USP115101


Figure 7.13: Example spectra from both sides of a USP115101 coupon. These samples are produced by Aviant. Figure a) shows the spectrum of the thin side of the coupon. Figure b) shows the spectrum from the thick side of the coupon. Figure c) shows the difference between figures b) and a). Note that the ranges on the y-axes of these plots is nearly identical: both plots begin around 10% reflectance and continue to above 90% reflectance at high wavelengths.

extrusion samples that cannot be machined away. The measured thickness of the extrusion samples is around 1.2 mm. There were ten extrusion samples produced for this investigation, though for a given measurement, five of the extrusion samples were chosen at random for reflectance data measurements. Averages were taken between different samples of both titanium dioxide coupons, cladding extrusion samples, and both standards to perform the final aging calculations for this investigation.


In fact, having too much scintillator still attached to the surface of the cladding produced strange reflectance spectra when measured. It was observed that, when measuring a sample with scintillator still on the surface, the entire reflectance spectrum shifts downward and appears to be



(a) Cladding side of an extrusion sample



(b) Scintillator side of an extrusion sample



(c) Side view of an extrusion sample

Figure 7.14: Photographs of an example extrusion sample. These samples were made by Alan Bross. Though the front and back of the sample look identical, one side has a thin layer of scintillator still on the surface. This is visible near the top of figure b) where there is a small strip where the cladding is present at the top surface of the sample and scintillator is not. Figure a) shows the cladding side of the coupon, figure b) shows the scintillator side of the coupon, and figure c) shows the side of the coupon.

less reflective across all wavelengths. This was verified using the GoreDRP standard and holding pieces of scintillator of different thicknesses flush with the surface of GoreDRP. As the scintillator piece in front of the GoreDRP got thicker, the reflectance spectrum decreased in magnitude across

all wavelengths. When the extrusion samples were first machined, it was also observed that the reflectance spectrum appeared to surpass 100% reflectance near 400 nm. This effect was caused by a very thin layer of scintillator on the surface of the sample, where the scintillator contains dopants that fluoresce. When the scintillator is activated, the dopants produce light in this wavelength range as a result of POPOP and PPO fluorescence. After this observation, the extrusion samples were carefully prepared to remove as much scintillator from the inner surface as possible. Example reflectance spectra from the samples are shown in Figure 7.15.

Comparing the spectrum from the exterior cladding surface of the extrusion sample in Figure 7.15 with the spectrum interior scintillator surface illustrates the effect that excess polymer has on the reflectance of the surface. The spectrum from the scintillator surface levels off at high wavelengths to around 60% reflectance, whereas the cladding surface levels off near 90% reflectance at the same wavelengths. For this reason and due to the effect of excess surface scintillator on the reflectance spectra, only the outer cladding side data is considered for aging rate calculations.


7.3 Results

With the descriptions of our samples and standards, as well as discussions about the stability of the instrument and analysis methodology in mind, we can now analyze the reflectance spectra and create plots of reflectance over time for our two sets of samples. This section will present and analyze the plots of reflectance over time for each set of samples that were measured to extract aging rates corresponding to each material. Since the aging plots are presented in reflectance, as a percent, over time, the aging rates will be presented in the same units: loss of reflectance over time, as a percent per year rate. For this investigation, plots that are referred to as "aging plots" are the plots of reflectance over time with a linear fit.


7.3.1 Aging of Titanium Dioxide Cladding Coupons

Now, with confidence in the spectrophotometer stability, the reflectance over time of the titanium dioxide cladding coupons from third-party manufacturers may be analyzed for signs of aging. As mentioned in the introduction to this chapter, one set of coupons were primarily analyzed for this investigation, USPN115101, as they were newly produced and received in March 2022. This set of coupons was chosen because they could be monitored starting shortly after production, where the most accelerated aging is theorized to occur. Thus, the newly produced USPN115101 coupons became the focus of this investigation in March. An average plot of reflectance over time for the thin side of the USPN115101 cladding coupon set at 390 nm is shown in Figure 7.16.

Like the aging plots presented for the standards, each point in Figure 7.16 is an average of the reflectance spectra collected for a given date. For the USPN115101 coupon set, the spectra



(a) Spectrum from the cladding side of an extrusion sample



(b) Spectrum from the scintillator side of an extrusion sample

Figure 7.15: Example spectra from both sides of an extrusion sample. These samples were produced in-house at FNAL. Figure a) shows the spectrum of the exterior cladding surface. Figure b) shows the spectrum from the interior scintillator surface which was machined away to isolate the cladding layer. Note the differences between the y-axes on these plots: they both begin below 10% reflectance, the cladding side extends to nearly 100% reflectance while the scintillator side only extends to 80% reflectance.

collected on each of the four coupons is averaged together rather than measuring one coupon repeatedly. The error associated with each point is the standard deviation of the measurements on the four coupons. The trendline on the data indicates a negative slope, which is what is expected as a result of aging. According to Figure 7.16, the thin side of the titanium dioxide cladding coupons




Figure 7.16: USPN115101 reflectance over time at 390 nm. The reflectance of the sample surface is shown on the y-axis in percentage versus time in fraction of years on the x-axis. The x-axis extends from 3/17/22 until 6/23/23. The trendline is a linear fit where the slope of the fit indicates the aging rate at this particular wavelength.

from the manufacturer age by losing $0.46\% \pm 0.06\%$ of surface reflectance per year at 390 nm. Analogous plots may be made for each wavelength in our wavelength range of interest, from 380 nm to 450 nm in increments of 10 nm, and also for the thick side of the coupons. A summary plot for the USPN115101 coupon set showing the aging rate for each wavelength extracted from the slope of the aging plot at each wavelength is shown in Figure 7.17.

Figure 7.17 shows that the aging rates that were collected at each wavelength for both the thin and thick sides of the coupons are small and negative. There is some structure in aging rate with respect to wavelength of light, it appears that aging occurs most rapidly at small wavelengths, near 380 nm, then as wavelength increases, aging rate decreases until around 420-430 nm, where the aging rate again increases again. At high wavelengths, the aging rate appears to slow. This structure is seen in both the thin side and the thick side reflectance data collected on the USPN115101 coupons. The cause behind the structure is not known. Overall, the results suggest that the titanium dioxide cladding material is undergoing small aging effects. The thin side of the USPN115101 coupons age at a mean rate of $0.33\% \pm 0.12\%$ of surface reflectance per year and the thick side of the USPN115101 coupons age at a mean rate of $0.15\% \pm 0.086\%$ of surface reflectance per year. These aging rates are given by the mean rate across all wavelengths with the standard deviation of the distribution of aging rates as the error. Though these aging rates are seemingly very small, on the order of half a percent change in surface reflectance or less, they could be significant. Any change in the properties of the cladding surface could impact the light yield from single cosmic ray




Figure 7.17: USPN115101 coupon aging rate summary plot. Aging rates for the USPN115101 titanium dioxide cladding coupon samples at different wavelengths of interest. Aging rates in loss of reflectance per year, extracted from the slope of the trendline of each aging plot, are shown on the y-axis. Wavelength, in nm, is shown on the x-axis.

events within the CRV counters. Light within the CRV counters must propagate into the wavelength-shifting fibers and down to the sensors at the ends of the counters to be detected, and thus the light bounces off the inner surface of the counters multiple times as it propagates through the bars. If the mean number of bounces that light takes off the inner surface is reduced, it could reduce the light yield in the counters because the light within the counters may need more than one chance to cross the interface between the scintillator and the fiber. By reducing the number of bounces, the number of chances that the light has to cross this interface is subsequently reduced.

7.3.2 Aging of Titanium Dioxide Cladding Extrusions

In an analogous fashion to the manufacturer-produced cladding coupons, we may also extract aging rates for the extrusion samples that were produced in-house. These samples were measured beginning on September 6, 2022. Since these extrusion samples were produced at FNAL and use the same specifications for the Mu2e CRV counters in terms of thickness of the cladding layer, composition of titanium dioxide to polystyrene in the cladding material, and co-extruded with FNAL-NiCADD scintillator, it is expected that these samples represent the condition of the Mu2e CRV cladding better than the manufacturer-produced coupons in the previous section. Recall the effects of excess scintillator on the cladding surface from the introduction of this chapter: excess scintillator or plastic polymer on the cladding surface reduces reflectance across the spectrum. In addition, recall that the scintillator includes dopants which fluoresce and produce light at certain

wavelengths. For these reasons, only the outer cladding surface is considered for the aging rate estimations in this section. An average plot of reflectance over time for the outer cladding surfaces of the extrusion samples at 390 nm is shown in Figure 7.18.




Figure 7.18: Extrusion sample reflectance over time at 390 nm. The reflectance of the sample surface is shown on the y-axis in percentage versus time in fraction of years on the x-axis. The x-axis extends from 9/6/22 until 6/23/23. The trendline is a linear fit where the slope of the fit indicates the aging rate at this particular wavelength.

Similar to the aging plots presented for the standards and the coupons, each point in Figure 7.18 is an average of the reflectance spectra collected for a given date. For the extrusion samples, five of the ten samples were chosen each week at random to measure reflectance data on, then the spectra collected on each of the five samples are averaged together rather than measuring one sample repeatedly. The error associated with each point is the standard deviation of the measurements on the five samples. The linear trendline on Figure 7.18 indicates an aging rate of $0.91\% \pm 0.13\%$ of surface reflectance per year at 390 nm. As was done for the coupons, analogous plots may be made for the extrusion samples for each wavelength in our wavelength range of interest, from 380 nm to 450 nm in increments of 10 nm. A summary plot for the extrusion sample set showing the aging rate for each wavelength extracted from the slope of the aging plot at each wavelength is shown in Figure 7.19.

Figure 7.19 shows that the aging rates that were extracted for each wavelength are both positive and negative, with the aging rate generally getting more positive as wavelength increases. This trend of smaller aging rates at larger wavelengths can also be seen in the summary plot for the USPN115101 coupon set, as well as loosely in the GoreDRP summary plot. The mechanism behind the trend for higher wavelengths to have more positive aging rates is unknown, but was noted by




Figure 7.19: Extrusion sample aging rate summary plot. Aging rates for the extrusion samples at different wavelengths of interest for the set of extrusion samples on the outer cladding side. Aging rates in loss of reflectance per year, extracted from the slope of the trendline of each aging plot, are shown on the y-axis. Wavelength, in nm, is shown on the x-axis.

the Scintillation R & D team. It is important to keep in mind that this data spans a shorter period of time than the data that was collected on the USPN115101 coupons, by roughly half a year, so with more data collected over a longer period of time, the rates may coalesce and the positive aging rates may begin to show negative aging trends after stabilizing over time. Though the reflectance data for the inner scintillator surface of the extrusion samples is not shown, the aging rates for the scintillator side followed a similar structure in the wavelength regime, shifted by roughly +1% in aging rate for each wavelength. Though the last three wavelength bins from 430-450 nm result in positive, unphysical aging rates, all wavelengths are used to find a mean aging rate for the extrusion samples. The outer cladding side of the extrusion samples that were produced in-house at FNAL age at a mean rate of $0.34\% \pm 0.75\%$ of surface reflectance per year, given by the mean aging rate across all wavelengths with the standard deviation of the distribution of aging rates as the error. This is another very small aging rate, but similar to the aging results from the cladding coupons, this small rate could possibly have a large effect on light yield within the Mu2e CRV counters for the same reason of reducing the mean number of bounces that scintillation light takes on the surface of the Mu2e CRV counters.

7.3.3 Studies on Other Counter Components

During this investigation, there have also been studies into isolating the aging effects from the other two major components of the Mu2e CRV: the scintillator and the wavelength shifting

fibers. Brian Leung studied the polystyrene base that is used during the scintillator extrusion process and measured polystyrene samples over time to look for signs of aging. By creating fresh extrusions of pure polystyrene samples, this study aimed to rule out aging in the polystyrene plastic base. Brian's data did not show significant signs of aging in pure polystyrene [7]. Matt Solt, postdoc from University of Virginia, lead a study into the wavelength shifting fibers used in the CRV counters. For part of the study, batches of wavelength-shifting fiber that were ordered from the manufacturer were measured with both a spectrometer and photodiode to determine the light yield in the fiber. Light yield data was collected on the same batches again after approximately one year and compared to provide a direct measurement of aging effects. These comparisons indicated a few percent decrease in light yield in the newer spools, but no appreciable aging effects in the benchmark spool [8]. The results of these two studies support that the polystyrene base and the wavelength-shifting fiber are not the main contributors to the overall aging rate that was observed in the Wideband data.

After this investigation concluded, further studies have taken place at Wideband to better estimate the overall Mu2e CRV aging rate. Postdoc Yongyi Wu has been continuing aging rate investigations using Wideband data where the analysis code and instrumental setup is upgraded as necessary. Recent aging rate studies on old and new data with updated analysis algorithms estimate closer to a 3-5% aging rate, in loss of light yield, in PEs, per year, rather than the 8% aging rate that I found during my original analysis that is documented in Chapter 6 [9]. Studies using the Wideband test bench will continue until the CRV is ready to be moved to the Mu2e hall. Collecting more cosmic data over time in Wideband will improve the aging rate estimation for the CRV.

7.4 Comparing with Simulations

One way to estimate the effect that surface aging of the titanium dioxide cladding will have on the Mu2e CRV overall is to use simulations of different types of events. Performing Monte Carlo simulations and studies of the Mu2e CRV counters can estimate the expected light yield within a counter in addition to simulating the entire CRV apparatus and detector response. Using the data collected on the cladding coupons and extrusion samples above, we can compare the light yield of simulations that use different cladding surface reflectivities and look for impacts of aging on the light yield of the CRV counters. Since the results for the extrusion samples suggested a wide range of surface reflectance aging rates, these comparison studies were performed assuming a 0.5% aging rate in the reflectance of the cladding surface. This number was close to the aging rate estimated from the two sets of samples analyzed for my study on titanium dioxide cladding in May 2023. After May 2023, a few more data points were collected by Brian and added to the aging plots, resulting in the rates discussed in the previous section. The aging rates estimated from the updated September plots are lower than the aging rates estimated from the May plots.

These Monte Carlo simulation studies were performed by graduate student George Iskander and postdoc Ralf Ehrlich of University of Virginia. These simulations use two types of sources to simulate cosmic ray particles: bismuth-207 and protons. These sources were chosen to compare to data that can be collected at Lab 6 or at the test beam facility at Fermilab. In Lab 6, there is a dark box that is used to collect data on short CRV counter pieces where a bismuth-207 source is placed at one end of a short CRV counter in the dark box, so the simulated bismuth source matches the energy and position of the Lab 6 bismuth source. During test beam runs, a 120 GeV proton beam is incident on CRV counters, so the simulated proton source matches the energy and direction of the test beam. For each type of source, there are different scintillator geometries to choose from, which include the cross-sectional dimensions of the scintillator, reflectivity of the cladding surface around the scintillator, the size and placement of fiber and fiber holes, and whether or not the fiber holes are filled with some sort of resin or epoxy to prevent effects of air gaps between the scintillator and the fibers. For the purposes of comparing these simulations to the data collected on Mu2e bar pieces, I selected the simulations closest to the Mu2e geometry for analysis. The bismuth simulations use the exact Mu2e geometry for the bars, including the correct elliptical fiber hole dimensions with no filling in the gaps in the fiber channels. The proton simulations are closer to cosmic ray muon particles in energy, but the geometries that are available to use with the proton simulations use slightly different fiber hole sizes than that of Mu2e production bars and the bismuth simulations. The fiber hole sizes that are simulated for the proton source simulations are circular, rather than elliptical, as a future project intends to use circular fiber channels.

The simulations then produce light yield histograms for different positions along the cross section of the dicounter. Within the simulations, the response of only one of the two fibers in the dicounter is simulated. The positions that are simulated in the histograms are relative to the fiber which is being simulated. The fiber that is being simulated is located at -13 mm (y), so the histograms representing the position at -10 mm (y) closest to the active fiber has a higher peak PE yield compared to the position at +10 mm (y) which is farther from the active fiber. A diagram showing the positions within the counter that are simulated is shown in Figure 7.20 [10].

The simulated light yield histograms for the bismuth simulations are shown in Figure 7.21 and the simulated light yield histograms for the proton simulations are shown in Figure 7.22. For both sets of these simulations, the current, nominal Mu2e titanium dioxide surface reflectivity simulations are shown on the top and the nominal surface reflectivity with 0.5% reflectivity loss is shown on the bottom.

From the plots in Figures 7.21 and 7.22, the light yield, number of reflections at the cladding coating surface, and track length in scintillator can be extracted and compared between the simulations representing the nominal titanium dioxide surface reflectivity with the simulations where the titanium dioxide surface reflectivity has been decreased by 0.5% from nominal. These




Figure 7.20: CRV counter diagram with MC position labels. Diagram of a single CRV counter extrusion, outlined in red, with schematics dimensions marked in blue, in mm. The positions that are simulated in the Monte Carlo study, in mm, are marked and labeled in black.

values are summarized in Table 7.1 for the bismuth simulations and Tables 7.2, 7.3, and 7.4 for the proton simulations. For the proton simulations, three different fiber hole sizes were simulated, 0.7mm, 0.9mm, and 1.2mm. The data from all three fiber hole sizes are summarized in tables, but only one set of proton simulation plots are shown in Figure 7.22 as an example.

Looking at the proton source simulations, the 0.7mm diameter fiber hole simulations have the highest light yield overall, but is not representative of the Mu2e geometry. Recall that the fibers used within the Mu2e dicounters are 0.7mm in diameter. Thus, a 0.7mm fiber hole diameter simulation assumes that there is no gap between the surface of the fiber and the surface of the scintillator. This artifact increases the light yield in these simulations. In Mu2e, the fiber holes are oblong and measure about 2.0mm \times 1.0mm in diameter, so there is a gap between the fiber and the scintillator in Mu2e dicounters. The 0.9mm and 1.2mm diameter fiber hole simulations do include a gap between the fiber surface and scintillator surface. The values simulated for light yield, number of reflections at cladding coating, and track length in scintillator are very close between the 0.9mm diameter fiber hole simulations and the 1.2mm diameter fiber hole simulations, so only the 0.9mm diameter fiber hole simulations are shown in Figure 7.22.

To analyze this data for general trends in response to the 0.5% decrease in titanium dioxide cladding surface reflectivity, averages may be calculated for the light yield, number of reflections at the cladding coating surface, and the track length in scintillator across the different simulated positions for each set of simulations. Using these averages, we can compare the trends between the simulations for current TiO₂ surface reflectivity to a reduction of 0.5% reflectivity across all

Position (y)	-10mm	-6mm	-3mm	0mm	3mm	6mm	10mm
Light Yield (PEs)							
Current TiO ₂	44	40	38	36	34	33	31
TiO ₂ - 0.5%	42	39	36	34	33	31	30
Average Reflections							
Current TiO ₂	6.4	6.7	6.8	6.9	6.8	6.7	6.4
TiO ₂ - 0.5%	6.0	6.4	6.5	6.5	6.5	6.3	6.1
Average Track length (mm)							
Current TiO ₂	186	196	201	202	201	196	186
TiO ₂ - 0.5%	178	188	192	194	192	187	178

Table 7.1: Light yield simulation results from bismuth source. Simulation results from bismuth source for light yield, number of reflections at the cladding coating surface, and average track length in scintillator. For each set of data, results using the current titanium dioxide surface reflectivity are shown in the top row and results using the current titanium dioxide surface reflectivity decreased by 0.5% are shown in the bottom row. Light yield is given in PEs and the average track length is given in mm.

Position (y)	-10mm	-6mm	-3mm	0mm	3mm	6mm	10mm
Light Yield (PEs)							
Current TiO ₂	110	99	94	90	87	85	80
TiO ₂ - 0.5%	106	97	92	88	84	81	77
Average Reflections							
Current TiO ₂	6.4	6.7	6.8	6.8	6.8	6.7	6.4
TiO ₂ - 0.5%	6.1	6.4	6.5	6.6	6.5	6.4	6.1
Average Track length (mm)							
Current TiO ₂	188	198	201	202	201	198	188
TiO ₂ - 0.5%	180	190	193	194	193	190	180

Table 7.2: Light yield simulation results from proton source with 0.7 mm fiber holes. Simulation results from proton source with 0.7 mm diameter fiber holes for light yield, number of reflections at the cladding coating surface, and average track length in scintillator. For each set of data, results using the current titanium dioxide surface reflectivity are shown in the top row and results using the current titanium dioxide surface reflectivity decreased by 0.5% are shown in the bottom row. Light yield is given in PEs and the average track length is given in mm.

Position (y)	-10mm	-6mm	-3mm	0mm	3mm	6mm	10mm
Light Yield (PEs)							
Current TiO ₂	61	56	54	52	50	49	47
TiO ₂ - 0.5%	59	54	52	50	48	46	45
Average Reflections							
Current TiO ₂	7.8	8.1	8.2	8.2	8.2	8.1	7.8
TiO ₂ - 0.5%	7.4	7.7	7.8	7.8	7.8	7.7	7.4
Average Track length (mm)							
Current TiO ₂	227	236	239	240	239	236	227
TiO ₂ - 0.5%	215	224	227	228	227	224	215

Table 7.3: Light yield simulation results from proton source with 0.9 mm fiber holes. Simulation results from proton source with 0.9 mm diameter fiber holes for light yield, number of reflections at the cladding coating surface, and average track length in scintillator. For each set of data, results using the current titanium dioxide surface reflectivity are shown in the top row and results using the current titanium dioxide surface reflectivity decreased by 0.5% are shown in the bottom row. Light yield is given in PEs and the average track length is given in mm.

Position (y)	-10mm	-6mm	-3mm	0mm	3mm	6mm	10mm
Light Yield (PEs)							
Current TiO ₂	61	56	54	52	50	49	48
TiO ₂ - 0.5%	59	54	52	50	48	46	45
Average Reflections							
Current TiO ₂	7.8	8.1	8.2	8.2	8.2	8.1	7.8
TiO ₂ - 0.5%	7.4	7.7	7.8	7.8	7.8	7.7	7.4
Average Track length (mm)							
Current TiO ₂	226	235	238	240	238	235	226
TiO ₂ - 0.5%	215	224	227	228	227	224	215

Table 7.4: Light yield simulation results from proton source with 1.2 mm fiber holes. Simulation results from proton source with 1.2 mm diameter fiber holes for light yield, number of reflections at the cladding coating surface, and average track length in scintillator. For each set of data, results using the current titanium dioxide surface reflectivity are shown in the top row and results using the current titanium dioxide surface reflectivity decreased by 0.5% are shown in the bottom row. Light yield is given in PEs and the average track length is given in mm.

	Avg Light Yield (PEs)	Avg Reflections	Avg Track Length (mm)
Current TiO ₂	36.6	6.67	195
TiO ₂ - 0.5%	35.0	6.33	187

Table 7.5: Average values from bismuth source light yield simulations. Average simulated values of light yield, in PEs, reflections at the cladding coating surface, and track length, in mm, from the bismuth source simulations. The averages are taken across the seven different cross-sectional positions that are simulated. The top row shows the current TiO₂ surface reflectivity simulated averages and the bottom row shows the simulated averages for the current TiO₂ cladding surface reflectivity reduced by 0.5%.

0.7mm fiber holes	Avg Light Yield (PEs)	Avg Reflections	Avg Track Length (mm)
Current TiO ₂	92.1	6.66	197
TiO ₂ - 0.5%	89.3	6.37	189
0.9mm fiber holes	Avg Light Yield (PEs)	Avg Reflections	Avg Track Length (mm)
Current TiO ₂	52.7	8.1	235
TiO ₂ - 0.5%	50.6	7.7	223
1.2mm fiber holes	Avg Light Yield (PEs)	Avg Reflections	Avg Track Length (mm)
Current TiO ₂	52.9	8.1	234
TiO ₂ - 0.5%	50.6	7.7	223

Table 7.6: Average values from proton source light yield simulations. Average simulated values of light yield, in PEs, reflections at the cladding coating surface, and track length, in mm, from the proton source simulations. The averages are taken across the seven different cross-sectional positions that are simulated. The table is split into three sections that show averages for the 0.7mm, 0.9mm, and 1.2mm diameter fiber hole sizes from top to bottom, respectively. For each set of data, the top row shows the current TiO₂ surface reflectivity simulated averages and the bottom row shows the simulated averages for the current TiO₂ cladding surface reflectivity reduced by 0.5%.

wavelengths. Then, we can also compare the trends found in the bismuth and proton sources and look for similarities between both sets of simulations. If the same trend is seen in both types of sources, it is likely that this trend will materialize in cosmic ray data. The average simulated values are collected in Table 7.5 for the bismuth source simulations and Table 7.6 for the proton source simulations.

Comparing the averages from the current titanium dioxide cladding surface reflectivity simulations with the averages of the simulations where the current titanium dioxide cladding surface reflectivity is reduced by 0.5% in Tables 7.5 and 7.6 reveals that light yield, the number of reflections at the coating surface, and the track length in scintillator all decrease in response to a loss in reflectivity of the cladding coating surface. For the bismuth source simulations in Table 7.5, the average light yield drops from 36.6 PEs to 35.0 PEs. This is a reduction of about 4.3% in light yield. Similarly, the number of reflections at the cladding coating decreases by about 5.1% and the track length in scintillator decreases by about 4.3%.

For the proton source simulations in Table 7.6, similar trends are observed in all three

simulations of different fiber hole sizes. For the 0.7mm diameter fiber hole simulations, the average light yield drops from 92.1 PEs to 89.3 PEs. This is a reduction of about 3.1% in light yield. Similarly, the number of reflections at the cladding coating decreases by about 4.3% and the track length in scintillator decreases by about 4.1%. For the 0.9mm diameter fiber hole simulations, the average light yield drops from 52.7 PEs to 50.6 PEs. This is a reduction of about 4.1% in light yield. Similarly, the number of reflections at the cladding coating decreases by about 5.0% and the track length in scintillator decreases by about 5.1%. For the 1.2mm diameter fiber hole simulations, the average light yield drops from 52.9 PEs to 50.6 PEs. This is a reduction of about 4.3% in light yield. Similarly, the number of reflections at the cladding coating decreases by about 5.0% and the track length in scintillator decreases by about 4.8%.

The trends in comparing the average simulated values for light yield, number of reflections at the cladding coating surface, and track length in scintillator suggest that all three of these event characteristics are impacted by a reduction of reflectivity in the TiO₂ cladding surface on the CRV dicounters. It appears that a 0.5% decrease in the reflectivity of the TiO₂ cladding surface results in roughly a 4% reduction in light yield and 4-5% reductions in the number of reflections at the cladding coating surface and track length in the scintillator. These trends were observed similarly in the bismuth source and the proton source simulations. A short discussion about the uncertainty in the simulations took place to ensure confidence in these simulations. Although the uncertainty in the simulations is not quantified, it was decided that the uncertainty within the simulations is much smaller than the uncertainty in the data that was collected with the spectrophotometer that we are comparing the simulations with.

7.5 Conclusions

In conclusion, this investigation into the aging of titanium dioxide cladding has revealed that the reflective TiO₂ cladding used in the Mu2e CRV counters ages slowly on the order of 0.3%-0.5% of surface reflectance lost per year. If we compare this 0.5% loss in reflectance aging rate of the titanium dioxide cladding to the aging rate that was estimated overall for the CRV in Chapter 6, it appears this small aging rate of the titanium dioxide cladding does not completely account for the high overall aging rate that was estimated from Wideband data. However, analyzing simulated datasets revealed that a 0.5% decrease in surface reflectivity of titanium dioxide results in roughly a 4% decrease in light yield. Comparing 4% decrease in light yield, in PEs, to the roughly 8% decrease in light yield, in PEs, that was observed at Wideband may suggest that small changes in the titanium dioxide cladding surface have a larger impact on light yield aging than expected.

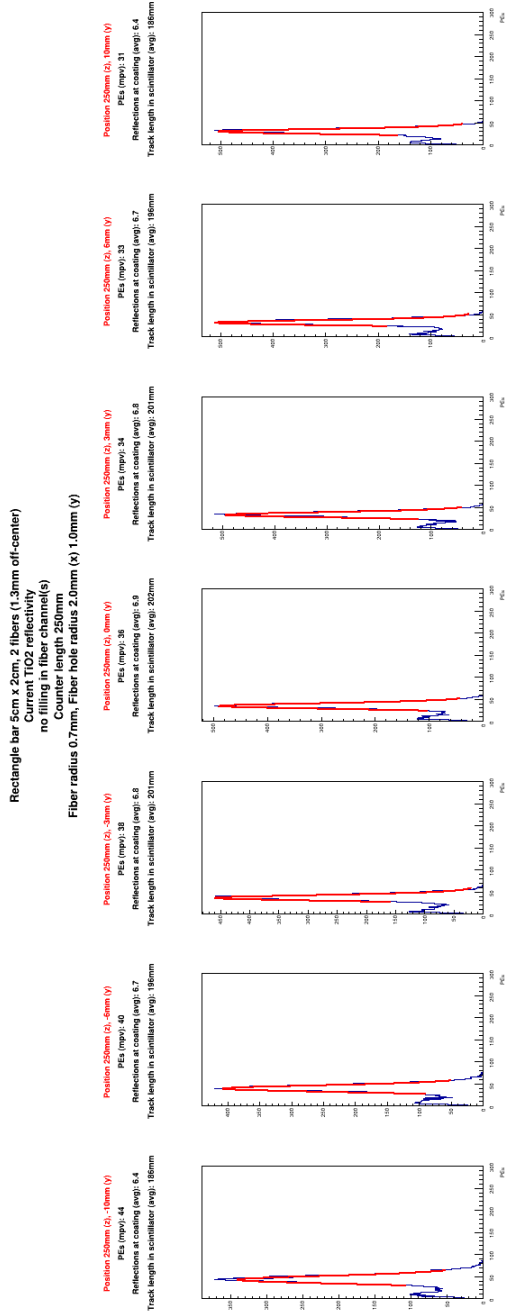
In order to determine just how much of an impact the 0.5% titanium dioxide reflectance aging rate has on the overall CRV light yield aging rate requires more studies with freshly produced

samples of CRV counters. As a result of this study, different Mu2e counter samples that were saved during extrusion production runs were discovered in the Lab 6 extrusion facility. When some of these samples were measured for outer surface reflectance, the reflectance spectra were not consistent between samples [5]. Any of the early changes in reflectance or light yield that occurred in these samples was not measured or documented, so the early aging effects in Mu2e counters are not well understood. This investigation serves as the first investigation into the early aging effects of TiO₂. There are plans to produce new Mu2e CRV counters with and without titanium dioxide cladding coextruded around the outside of the scintillator for this purpose. In these studies, the light yield of the samples will be measured beginning from production, where they can be monitored for rapid initial aging, and the samples with and without titanium dioxide cladding can be compared for differences in light yield to decouple the effects of scintillator aging from cladding aging. If the sample with cladding shows faster decreases in light yield compared to the sample without cladding, it would have interesting implications for not only Mu2e, but any other present or future experiments using scintillator-based detectors with reflective claddings on the scintillator surface [5].

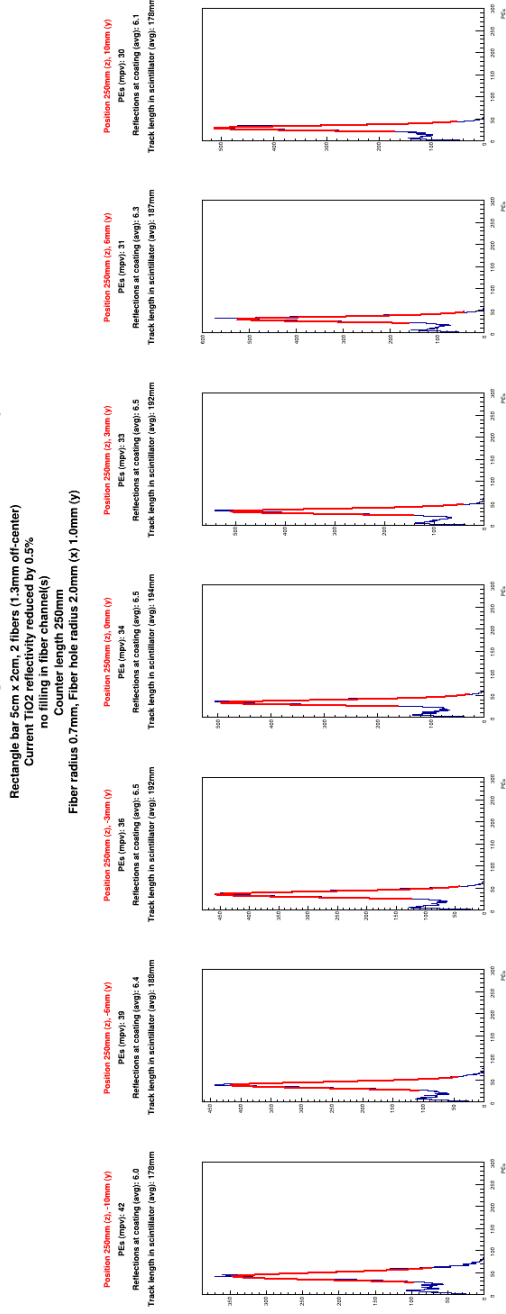
References

- [1] D. Beznosko, A. Bross, A. Dyshkant, A. Pla-Dalmau, and V. Rykalin, *FNAL-NICADD Extruded Scintillator*, Tech. Rep. (Fermilab, 2005) [FERMILAB-PUB-05-344].
- [2] Hunter Associates Laboratory, *UltraScan VIS Supplemental Manual for EasyMatchQC*, [Online] Available at:
https://hunterlab-production.s3.amazonaws.com/media/documents/Supplemental_Manual_for_EZMQC_UltraScan_VIS.pdf.
- [3] Labsphere Inc., “Spectrafect 97% Diffuse Reflectance Coating,” (2019), datasheet, Available at: <https://www.labsphere.com/wp-content/uploads/2021/07/Spectrafect.pdf>.
- [4] A. Pla, (private communication).
- [5] A. Bross, (private communication).
- [6] Kuraray, “Wavelength shifting fibers datasheet,” , datasheet, Available at: <http://kuraraypsf.jp/psf/ws.html>.

- [7] B. Leung, “Lab 6 Polystyrene Polymer Tests – IP2I Extrusion,” (2022-10-07), <https://indico.fnal.gov/event/56488/>.
- [8] M. Solt, D. Coveyou, E. C. Dukes, R. C. Group, Y. Oksuzian, and S. Roberts, “Performance of the wavelength-shifting fiber upgrade for the Mu2e cosmic-ray veto detector,” (2023), arXiv:2302.09172 [hep-ex].
- [9] Y. Wu, “Mu2e CRV Module Aging Rate Studies,” (Mu2e Biweekly Meeting 2023-09-14), mu2e-doc-46508-v1.
- [10] S. Boi, Ph.D. thesis, University of Virginia (2021).
- [11] V. Senchishin, V. Koba, O. Korneeva, V. Seminozhenko, V. Kovtun, I. Zalubovsky, I. Chirikov-Zorin, J. Budagov, F. Markley, and G. Bellettini, *New Radiation Stable and Long-Lived Plastic Scintillator for the SSC*, Tech. Rep. (Fermilab, 1993) FERMILAB-TM-1866.
- [12] L. Aliaga, *et. al.* (The MINERvA Collaboration), Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip. **743**, 130 (2014).
- [13] K. Abe, *et. al.* (The T2K Collaboration), “Scintillator Ageing of the T2K Near Detectors from 2010 to 2021,” (2022), arXiv:2207.12982 [ins-det].



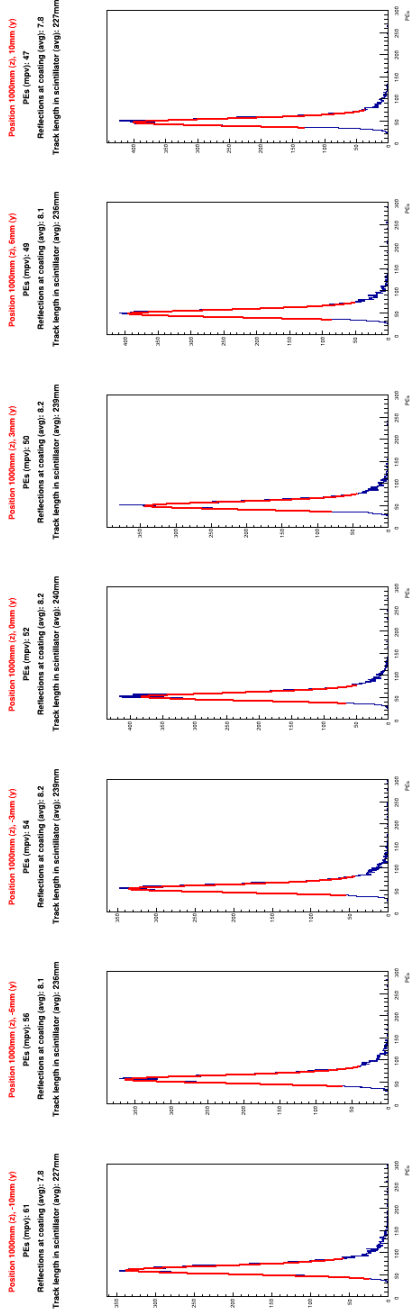
(a) Simulations using nominal TiO₂ reflectivity



(b) Simulations using - 0.5% TiO₂ reflectivity

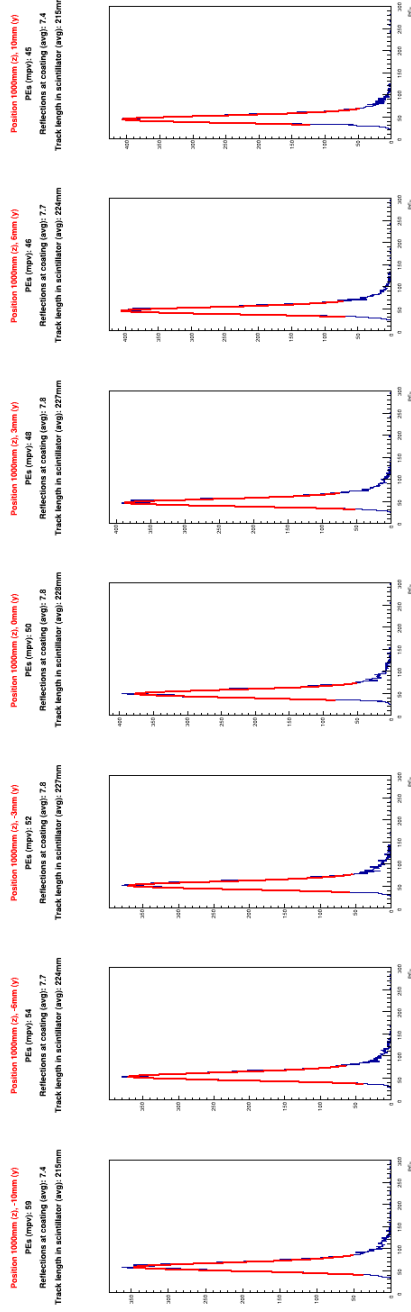
Figure 7.2.1: Simulated histograms of the light yield within a CRV dicounter at different cross-sectional positions using a bismuth-207 source. Counts are on the y-axes with light yield, in PEs, on the x-axes. Figure a) shows light yield simulations with the current, nominal Mu2e TiO₂ cladding surface reflectivity and Figure b) shows light yield simulations with a - 0.5% drop in TiO₂ surface reflectivity compared to nominal.

Rectangle bar 5cm x 2cm, 2 fibers (1.3mm off-center)
 Current TiO₂ reflectivity
 no filling in fiber channel(s)
 Counter length 3000mm
 Fiber radius 0.7mm, Fiber hole radius 0.3mm (x) 0.9mm (y)



(a) Simulations using nominal TiO₂ reflectivity

Rectangle bar 5cm x 2cm, 2 fibers (1.3mm off-center)
 Current TiO₂ reflectivity reduced by 0.5%
 no filling in fiber channel(s)
 Counter length 3000mm
 Fiber radius 0.7mm, Fiber hole radius 0.9mm (x) 0.9mm (y)



(b) Simulations using - 0.5% TiO₂ reflectivity

Figure 7.22: Simulated histograms of the light yield within a CRV dicounter at different cross-sectional positions using a proton source. Counts are on the y-axes with light yield, in PEs, on the x-axes. Figure a) shows light yield simulations with the current, nominal Mu2e TiO₂ cladding surface reflectivity and Figure b) shows light yield simulations with a - 0.5% drop in TiO₂ surface reflectivity compared to nominal.

Chapter 8 Mu2e Code Management

In 2020, I started working with the Mu2e Computing and Software group to take over some code validation and infrastructure management duties. In this role, I have been responsible for updating external products, internal products, and compiler versions in addition to maintaining Mu2e code to ensure compatibility after various updates take place.

8.1 Code Maintenance and File Storage Systems

The Mu2e source code is primarily maintained on Github and lives on various Fermilab disks, as well as on internationally distributed disk systems. The Mu2e experiment has its own organization on Github where any relevant Mu2e code is stored and maintained [1]. The primary repository that is used for experimental simulations and analysis is the Mu2e Offline repository. The disk space that distributes Mu2e code and relevant products is the CERN Virtual Machine File System (CVMFS) [2]. CVMFS connects the experiment's code and libraries to interactive nodes and grids worldwide; it is used by most experiments at FNAL. Any code changes or product upgrades are completed by pushing the relevant changes to Github or adding new content to CVMFS. When a code manager makes changes to the CVMFS disk space, a transaction is opened. The changes are made or products are added, then the transaction is published to CVMFS. The process of publishing archives any action on the code in a database for future reference.

CVMFS is only efficient for distributing code and small data files, so larger data files are not handled by this process, and are instead distributed by dCache using the /pnfs/ file system [3]. dCache is a distributed storage system for scientific data. In the Mu2e space, dCache functions as a system of many disks aggregated across dozens of Linux disk servers. The system lets all of these disks look like "one big disk" to the user and hides all of the details of exactly where the files are and how they are being transferred. The entire system is designed to be high-bandwidth so it can serve data files to thousands of grid nodes simultaneously. This allows for load balancing and optimization where possible across nodes [4]. There are several different "flavors" of dCache areas where Mu2e data is stored: scratch, persistent, resilient, and tape-backed. Scratch is a temporary output area for any Mu2e user producing simulation or analysis data where data files are automatically deleted according to a least-recently-used algorithm as space is needed. The other three areas are not for

general use by the collaboration and are intended for long-term storage of production files used in simulation campaigns [4].

8.2 UPS Products

External products are packaged via Unix Product Support (UPS) and distributed via Unix Product Distribution (UPD). UPS/UPD are tools developed and supported by Fermilab to manage external software products in a Unix environment [5]. From the Mu2e point of view, ‘external’ means any product that does not come from any of our Mu2e Github repositories, including key products like art, Geant4, ROOT, and so on. UPS supports access to multiple different versions of a product and multiple builds per version, and switching between different versions and builds as needed.

Within the scope of my code management for Mu2e, I often interact with UPS via the Scisoft FNAL server. The Scisoft server is located at scisoft.fnal.gov and is only accessible within the Fermilab network [6]. This server contains many UPS products as standalone ready-to-install product tarballs and as product manifests which include lists of product tarballs that are required for the upgrade of a single product. Each UPS product has a series of qualifiers that indicate the specific product versions that are available. The UPS tarballs that are generated for each product are clearly labeled with relevant qualifiers for selecting appropriate product versions. There are a few scripts included as Scisoft tools that can be run from the command line within the `/cvmfs/` disk space and allow for products to be pulled from the Scisoft servers. To update a single product, like ROOT or CLHEP, I would typically navigate to the Scisoft web page, find the product with the appropriate qualifiers that I need, copy the single ready-to-install tarball link, and install the product individually using `wget` and `tar` commands onto the CVMFS disk space for Mu2e external products. To update a product that uses a manifest to maintain its dependencies, like art or Geant4, I would typically use the Pull Products script that is provided as a Scisoft tool to install multiple products at once onto the CVMFS disk space for Mu2e external products. The Pull Products script takes product qualifiers as arguments to select and install the appropriate product versions from the command line. After installing products individually or with a manifest, I check the corresponding product directories within the Mu2e external product disk space to ensure that the new product version numbers appear within the appropriate directories.

8.3 Internal Products

Products that are generated by the collaboration internally typically use a repository within the Mu2e Github organization for source code storage and the builds are handled using a code building tool called Jenkins. Jenkins is a freeware Java code build system that is maintained by the

FNAL Scientific Computing Division [7]. In addition to building products, Jenkins includes many management features; the Jenkins system is browsable for different products and also produces reports and archives each time a product build is attempted. Jenkins build scripts are written in bash and executed within the Jenkins Buildmaster web page that serves as a GUI to Jenkins. Access to the Jenkins system is restricted to Fermilab users and using Jenkins off-site requires a Fermilab VPN connection as well as other permissions from the Computing group. Like UPS products, Mu2e internal products that are built using Jenkins also have qualifiers that specify their version number and dependencies. These qualifiers are entered into boxes and the build is started using a button on the web page. After a build is launched within the Buildmaster, progress may be monitored using a progress bar that appears on the screen. For successful builds, a log file and a ready-to-install tarball file is produced that can be installed onto CVMFS as an individual product. For unsuccessful builds, a log file is produced that contains error messages related to the failed build. Some of the products that are produced by Mu2e Github repositories using Jenkins are tracking tools like KinKal and its predecessor BTrk, compiling tools like SCONS, and experiment-specific art peripherals like `artdaq_core_mu2e` and `mu2e_pcie_utils`.

8.4 Integrating Updated Products

After installing products onto CVMFS, there are still several steps to complete before updated products are integrated into the head of the Mu2e Offline repository and made widely available to Mu2e users. The Mu2e Offline code depends on many different internal and external products, and thus it is useful to archive different releases of the Mu2e Offline code with specific dependencies so that they may be used for later reference. There are two ways which Mu2e accomplishes this: building Offline as a ready-to-install product with Jenkins that can live on CVMFS and using a UPS product called Muse to catalog minimal Offline releases with environmental setup files. With respect to the first method, Mu2e Offline is a project within Jenkins that may be built and produces an Offline tarball with a unique version number that is installed onto CVMFS. The second method involving Muse does not use Jenkins, but is nonetheless a powerful tool for the Mu2e collaboration.

The Muse system is a set of scripts developed inside Mu2e to provide a bridge between a minimal, up-to-date release of Offline and small, user-developed analysis repositories. Muse allows for these smaller user code repositories to be built together with Offline or against an existing Offline build. Minimal, stable Offline builds with a specific set of product dependencies can be captured and stored into a file called an envset, short for environmental setup. Envset files are stored on both Github and CVMFS in a Mu2e directory containing small data files. The naming conventions for envset files are to name official Mu2e envsets using three numbers and the letter p for 'profile', pNNN, for example p001. When envsets are being developed and tested, or a user

needs a custom set of products that is non-standard, a user-defined envset is named uNNN, for example u004. Once new products are validated, the user-defined uNNN envsets may be renamed to their corresponding official release names, pNNN, and stored appropriately. After new products are installed to CVMFS, a new envset is typically required to use updated products with Offline.

After a new envset is created, it may be used to build against the existing Mu2e Offline code. Building Mu2e Offline code is initiated with a muse build command and linking the appropriate envset file. Typically, building is done interactively, so any errors that occur during the build are printed to the terminal and the build is halted. There is also a flag that may be used with the muse build command that forces compilation despite errors. In this case, it is useful to write any terminal output to a log file so error messages may be easily searched for. Some products require changes in the code to make new versions compatible. Oftentimes, product developers within Fermilab will write documentation along with new product versions if there are inherent code changes required in order to use new products. Art and art peripheral products often require code changes as the syntax for art commands evolves over time. Any code that is obsolete and requires changes will be flagged as an error by the compiler, and an error message will be written to the terminal as described above.

Once the Mu2e Offline build proceeds without errors, the new products need to be validated. There is a standard validation suite that runs nightly and emails the results of this short validation to a group of Mu2e code managers. For Mu2e internal products, like KinKal, the manager of the product will verify any new versions and make sure the updated products function as expected. For larger product upgrades, like updating art or Geant4, a more rigorous validation process takes place. The larger validation process that occurs here involves running a set of simulations using both the updated version of the Mu2e Offline code and the last stable version before the changes and comparing the output of these simulation jobs. The typical simulation jobs for this type of validation include conversion electron events, protons on target events, and muon stop events. After these jobs are run on the grid, a set of standard validation scripts are used to create meaningful histograms from the simulation data. The simulation samples from new and old environments are then compared using a standard script. The result of the comparison script is a set of overlaid plots that indicate how closely the new simulations match the old ones. In addition to comparing the results of the physics simulations, scripts are used to extract the file sizes and speeds of different parts of the simulations to compare between the new and old environments. The averages of total file size, memory used, total run time, and Geant4-only run time are the four characteristics that are compared between new and old simulation files. If the differences in the results of the new and old simulations are within tolerance, the new envset may be officially released to the wider collaboration. If the validation results point to large differences between the new and old environments, further simulations may be needed to investigate the behavior of new product versions. Results from these larger validation campaigns are documented using the Mu2e wiki web page under the entry 'GeantChecklist' [8].

After completing successful validation, the envset is saved as an official release in pNNN format and then it is stored on both Github and CVMFS. Within CVMFS, there is the notion of a ‘current’ envset and Offline release. The current link points any user attempting to use Muse to set up their Offline environment to this most up-to-date official release. The envset version number is also included into the Muse setup file within Offline, and the head of the Mu2e Offline Github repository is updated with a pull request to include the latest changes to the code base [1].

References

- [1] Mu2e Collaboration, “Mu2e Github Organization,” (2009), [Online] Available at: <https://github.com/Mu2e>.
- [2] J. Blomer, B. Bockelman, D.-F. Dosaru, D. Dykstra, N. Hardi, N. Hazekamp, R. Meusel, S. Mosciatti, R. Popescu, and L. Promberger, “Welcome to CernVM-FS’s Documentation,” (2022), [Online] Available at: <https://cvmfs.readthedocs.io/en/stable/>.
- [3] P. Fuhrmann, “dCache, the Overview,” (2006), [Online] Available at: <https://www.dcache.org/docs/dcache-whitepaper-light.pdf>.
- [4] R. Culbertson, L. Goodenough, R. Kutschke, and Y. Oksuzian, “dCache,” (2017), [Online] Available at: <https://mu2ewiki.fnal.gov/wiki/Dcache>.
- [5] R. Culbertson and R. Kutschke, “UPS,” (2017), [Online] Available at: <https://mu2ewiki.fnal.gov/wiki/UPS>.
- [6] Fermilab, “Scientific Software for Relocatable UPS,” (2014), [Online] Available at: <https://scisoft.fnal.gov/>.
- [7] R. Culbertson and H. Casler, “Jenkins,” (2017), [Online] Available at: <https://mu2ewiki.fnal.gov/wiki/Jenkins>.
- [8] R. Culbertson, K. Genser, M. Devilbiss, and R. Kutschke, “GeantChecklist,” (2018), [Online] Available at: <https://mu2ewiki.fnal.gov/wiki/GeantChecklist>.

Chapter 9 Conclusion

The Mu2e experiment at Fermilab will search for the neutrinoless conversion of a muon to an electron in the presence of a nucleus. Mu2e is currently under construction and will soon be in the detector commissioning phase with an expected first physics run of data in 2026. The first run will improve the current limit on $R_{\mu e}$ by a factor of 10^3 . The full dataset is expected by the end of the decade which will further improve upon the $R_{\mu e}$ limit by another order of magnitude. The experiment uses superconducting solenoids to steer an intense muon beam onto a stopping target, where the muons captured by the metallic nuclei of the target and may convert to electrons. Downstream detectors precisely track particles that are ejected from the stopping target to search for conversion electron event candidates.

In order for Mu2e to achieve its expected sensitivity of 3×10^{-17} , the experiment employs rigorous background mitigation techniques to suppress the total measurement background count to a sub-event level throughout the duration of the experiment. The most dominant background is from cosmic ray-induced events, so the CRV has been developed to surround the Mu2e DS and part of the TS as a dedicated, active shield against cosmic rays. The CRV is required to operate at an efficiency at least 99.99% to sufficiently control cosmic ray-induced backgrounds, which drives the design of the detector. Using four layers of scintillating counters paired with a dead-time veto window upon cosmic detection, the CRV delivers on its goal to mitigate cosmic rays.

Testing individual CRV counters using the test bench in Wideband at Fermilab is an invaluable resource for understanding the CRV's response to cosmic rays before the detector is in full operation. Long-term monitoring at Wideband revealed that the detector is aging and the data that has been collected is being used to estimate the aging rate of the CRV as a whole. A data collection campaign to monitor the light yield of a CRV counter over time is active and the aging rate estimation conducted in 2022 estimated that the CRV was aging at a rate of $8.06\% \pm 0.33\%$ per year in light yield lost per channel, in PEs. Further investigation into determining whether or not an individual component of the CRV counters, namely the titanium dioxide cladding, was responsible for a bulk of the aging rate estimated from the data at Wideband. By monitoring the reflectance of the surfaces of different titanium dioxide cladding samples over time, it was estimated that the surface of the titanium dioxide cladding decreases in reflectance by roughly 0.3% to 0.5% per year on average. Simulations indicate that a 0.5% change in reflectance can result in a 3-4% drop in

light yield, in PEs. The studies detailed in this dissertation have generated discussion about CRV detector technology and future investigations are planned to continue probing the effects of titanium dioxide and scintillator aging on detector response.

Studies will continue to determine the rate of aging on the Mu2e CRV as well as uncover the underlying cause for the loss of efficiency over time. The Wideband test bench remains in operation and various improvements have been made to improve the stability and monitoring of the environmental temperature, purity of the trigger using shielding blocks in addition to the trigger paddles, and orientation and setup of the front-end electronics for the modules to replicate how they will be set up during operation. In addition, studies continue to determine how temperature and other environmental and electronic fluctuations can impact the response of the readout. It has been found that several more corrections need to be added to the data to properly remove effects from things like temperature of the electronics or change in bias voltage on different FEB components. The Scintillation R & D team has also been continuing investigations into what kind of impact cladding material has on the overall aging rate of the CRV. Dr. Alan Bross plans to produce extrusion samples of Mu2e scintillator with and without cladding co-extruded to decouple scintillator aging effects from cladding aging effects. The story is not over when it comes to scintillator aging for Mu2e and other future experiments proposing scintillator-based detector design.

Regardless of whether or not Mu2e observes a conversion event, there are future upgrades planned for the Mu2e experiment in the form of Mu2e-II. Mu2e and other precision searches into CLFV phenomena provide strong search windows for BSM physics. This area of research will be an exciting space to check up on in the coming decade for any signs of new physics and for deeper understanding of our universe on the smallest scale.

Appendix A Propagation Time within CRV Bars

A.1 System of Equations for Propagation Time

In Mu2e, it is important to note that each of our CRV dicounters are identifiable with a unique index, and that each counter in a dicounter is outfitted with two SiPMs on each end (each labelled as SiPM 0-3 per counter). SiPMs 0 and 2 are on one side, and SiPMs 1 and 3 are on the other. Each CRV counter contains two fibers: one fiber connects SiPM 0 to SiPM 1 and the other connects SiPM 2 to SiPM 3.

A typical pulse block for a cosmic event contains on the order of 20 reco pulses per event. Each reco pulse is labeled by a dicounter index and SiPM number, such that ideally each counter in the pulse block has four corresponding reco pulses: one for each SiPM. Using these measured leading times for each reco pulse (t_0, t_1, t_2, t_3 ; labelled by SiPM number), the length of the bar (L), and the speed of light in the scintillator (v), we may solve for a corrected time (t^*) which accounts for propagation in the CRV counter. Naturally, we may also solve for a corrected position of origin for each reco pulse along the x-axis (x^*).

First, the leading times of the reco pulses detected on the same side of the bar are averaged.

$$\frac{1}{2}(t_0 + t_2) = t_{0_{avg}}$$

$$\frac{1}{2}(t_1 + t_3) = t_{1_{avg}}$$

The system of equations to solve is as follows:

$$t_{0_{avg}} = t^* + x^* * v^{-1}$$

$$t_{1_{avg}} = t^* + (L - x^*) * v^{-1}$$

Which we may solve for t^* and x^* :

$$t^* = \frac{1}{2}(t_{1_{avg}} + t_{0_{avg}} - (L * v^{-1}))$$

$$x^* = \frac{1}{2}(L - (t_{1_{avg}} - t_{0_{avg}}) * v)$$

For a given event, we can calculate one corrected time, t^* , per bar. We may then average the corrected times together to find the overall corrected time which will be associated with the $T_{start}(CRV)$ of the event.

Appendix B Determining Optimal PMT Operating Voltage

The Wideband test bench uses scintillating trigger paddles connected to PMTs as the trigger system for cosmic ray events through the CRV module that is being measured at the test bench. The trigger system in Wideband consists of three scintillating trigger paddles, a coincidence module, and a NIM gate generator. Each trigger paddle is attached to a PMT. For some older paddles used with module 127, the paddle efficiencies and optimal operating voltages were determined at University of Virginia. For newer paddles that arrived at the beginning of 2022, I determined the paddle efficiencies and optimal operating voltages myself at Wideband. All of the trigger paddles that are used in Wideband were assembled at University of Virginia and later shipped to Fermilab.

B.1 Optimal Operating Voltage Analysis

Analyzing plots of efficiency versus operating voltage is the best way to extract the optimal operating voltage of each trigger paddle. To make these plots, we can form a three-fold coincidence module using three trigger paddles stacked on top of each other. In this configuration, the signals from the outer two paddles are funneled into a coincidence module. Then, the output of the coincidence of the outer two paddles is used as an input along with the signal from the middle paddle to form a three-fold coincidence. By constructing the three-fold coincidence in this way, we are effectively measuring the efficiency of the middle paddle. We can then use two counters to keep track of the number of times the outer paddles are hit simultaneously and the number of times that all three paddles are hit simultaneously using the outputs of the two coincidence modules that are monitoring our trigger paddles. The effective efficiency of the middle paddle is then the ratio of the number of events that are seen by all three paddles to the number of events that are seen by the outer two paddles.

The outer two paddles are operated at a voltage that is sufficiently high enough that it is beyond or around the estimated optimal operating voltages for this specific model of PMTs. This ensures that the outer paddles are maximally efficient for counting any cosmic events that may pass through them. For these investigations, the outer paddles were operated at 1450V. The middle paddle is then operated in a range of voltages to determine which operating voltage is ideal. The sweeps that I did with the middle paddles ranged from 1200V to 1500V. The high voltage (HV)

power supply in Wideband is an older HV generator that uses an array of holes and pegs to choose the amount of HV supplied to particular channels. Each of the three trigger paddles in this optimal operating voltage investigation was powered using separate channels.

B.2 Data Collection

To collect data, the middle paddle was set to a particular operating voltage for five minute intervals. At the beginning and end of each five minute interval, the values of the two aforementioned counters were recorded. Then, the difference between the initial and final values was taken to determine how many events were recorded on the outer paddles and on all three paddles during this time. The ratio of counts is taken between the counts on the outer panels and the counts on all three panels to determine the efficiency of the middle paddle. Now, the data may be filled into a scatter plot where the ratio of counts, or efficiency, is plotted along the y-axis and the operating voltage of the middle paddle at the time is plotted along the x-axis.


Each plot then resembles an ‘S’ shape curve where the ratio of inner to outer event counts begins at 0 for low operating voltages and approaches 1 for high operating voltages. When the paddle is operating below its optimal voltage, many events are not detected, and thus the efficiency is low. As the operating voltage for the middle paddle increases, more events are detected until the middle paddle is detecting roughly the same number of events as the outer paddles, where the efficiency nears 100%. The optimal operating voltages for each paddle are determined by looking at the Ratio of Counts vs Voltage plots and finding the ‘shoulder’ of the plot where the Ratio of Counts approaches 1. The optimal operating voltage is 25-50V above this shoulder.

B.3 Results


Plots of efficiency versus operating voltage were created for each of the six trigger paddles obtained from University of Virginia by placing each paddle in the middle of two other paddles and collected data as described above. These plots are shown in Figure B.1. The optimal operating voltages determined were determined for each paddle by analyzing the plots and finding the voltage at the ‘shoulder’ of the plot for each paddle. The optimal operating voltages are summarized in Table B.1. The three trigger paddles that were chosen to be used with module 127 were paddles 2, 3, and 6.

Paddle 1	1450V
Paddle 2	1400V
Paddle 3	1400V
Paddle 4	1425V
Paddle 5	1350V
Paddle 6	1450V


Table B.1: Optimal operating voltages for the PMTs used on the scintillating trigger paddles in Wideband.




(a) Paddle 1 efficiency vs operating voltage




(b) Paddle 2 efficiency vs operating voltage




(c) Paddle 3 efficiency vs operating voltage



(d) Paddle 4 efficiency vs operating voltage



(e) Paddle 5 efficiency vs operating voltage



(f) Paddle 6 efficiency vs operating voltage

Figure B.1: Efficiency vs operating voltage for trigger paddles. Efficiency, in ratio of counts collected by all three paddles to the counts collected by only the outer paddles, versus operating voltage of the middle trigger paddle PMT, in volts, from the data collected in Wideband in January 2022. The six different figures show the six different trigger paddles: figures a) through f) show data collected using paddles 1 through 6 in the middle, respectively.