


Proposal to Search for $\mu^- N \rightarrow e^- N$ with
a Single Event Sensitivity Below 10^{-16}

Mu2e Experiment



Fermilab
October 10, 2008

The *Mu2e* Collaboration

R.M. Carey, K.R. Lynch, J.P. Miller*, and B.L. Roberts
Boston University, Boston, Massachusetts

W.J. Marciano, Y. Semertzidis, and P. Yamin
Brookhaven National Laboratory, Upton, New York

Yu.G. Kolomensky
University of California, Berkeley, California

W. Molzon
University of California, Irvine, California

J.L. Popp
City University of New York, New York, New York

C.M. Ankenbrandt, R.H. Bernstein*, D. Bogert, S.J. Brice, D.R. Broemmelsiek, R. Coleman,
D.F. DeJongh, S. Geer, D.E. Johnson, R.K. Kutschke, M. Lamm, M.A. Martens, D.V. Neuffer,
M. Popovic, E.J. Prebys, R.E. Ray, M.J. Syphers, H.B. White, K. Yonehara, and C.Y. Yoshikawa
Fermilab, Batavia, Illinois

K.J. Keeter and E. Tatar
Idaho State University, Pocatello, Idaho

P.T. Debevec, G. Gollin, D.W. Hertzog, and P. Kammel
University of Illinois, Urbana-Champaign, Illinois

V. Lobashev
Institute for Nuclear Research, Moscow, Russia

D.M. Kawall and K.S. Kumar
University of Massachusetts, Amherst, Massachusetts

R.J. Abrams, M.A.C. Cummings, R.P. Johnson, S.A. Kahn, S.A. Korenev, T.J. Roberts, and
R.C. Sah
Muons, Inc., Batavia, Illinois

André de Gouvêa
Northwestern University

F. Cervelli, R. Carosi, M. Incagli, T. Lomtadze, L. Ristori, F. Scuri, and C. Vannini
Istituto Nazionale di Fisica Nucleare Pisa, Università Di Pisa, Pisa, Italy

M. Corcoran
Rice University, Houston, Texas

R.S. Holmes and P.A. Souder
Syracuse University, Syracuse, New York

M.A. Bychkov, E.C. Dukes, E. Frlez, R.J. Hirosky, A.J. Norman, K.D. Paschke, and D. Počanić
University of Virginia, Charlottesville, Virginia

J. Kane
College of William & Mary, Williamsburg, Virginia

*Contactperson

Contents

Executive Summary	xv
I <i>Mu2e</i> in Brief	1
1 Theoretical Motivation	3
1.1 Introduction	3
1.2 Comparison to Current HEP Program	5
1.2.1 Model Independent Analysis: $\mu N \rightarrow e N$ vs. $\mu \rightarrow e \gamma$	5
1.2.2 CLFV and new physics at the TeV scale	8
1.2.3 CLFV, neutrino masses, and the matter–antimatter asymmetry of the Universe	13
2 Overview of the Experiment	17
2.1 Experimental Sensitivity	20
2.2 Experimental Principles	20
2.3 Proton Beam	22
2.4 Solenoid System	24
2.4.1 Production Solenoid	24
2.4.2 Transport Solenoid	25
2.4.3 Detector Solenoid	26
2.5 Stopping Target	28
2.6 Detector	28
2.7 Backgrounds	29
2.7.1 Decay-in-Orbit Electrons	29
2.7.2 Nuclear Capture of Muons	30
2.7.3 Radiative pion capture	30
2.7.4 High energy beam electrons, (~ 100 MeV)	31
2.7.5 Late Arriving Particles	31
2.7.6 Cosmic Rays	32
2.8 Phase II: <i>Mu2e</i> at Project X Intensities	32

CONTENTS

3	Sensitivity and Backgrounds	35
3.1	Introduction	35
3.1.1	Underlying Processes and Calculations	36
3.2	Physics Background Sources	37
3.2.1	Electrons from Muon Decay in Orbit	39
3.2.2	Radiative μ Capture	40
3.2.3	Beam Electrons	43
3.2.4	Muon Decay in Flight	44
3.2.5	Pion Decay in Flight	45
3.2.6	Radiative π Capture	45
3.2.7	Antiproton Induced	46
3.2.8	Long Transit Time Backgrounds	50
3.2.9	Cosmic Rays	51
3.3	Physics Requirements	52
3.4	Background Level	52
3.5	Resolution Measurements	53
3.5.1	Decay-in-Orbit	53
3.5.2	Extinction Measurement	58
3.6	Conclusions on the Sensitivity	59
4	Cost Estimate and Schedule	61
4.1	Cost Estimate	61
4.2	Schedule	63
5	R&D Request	65
5.1	Introduction	65
5.2	Specific R&D Tasks	65
5.2.1	Solenoid System	65
5.2.2	Tracker	66
5.2.3	Extinction System and AC Dipole	67
5.2.4	Accelerator	67
5.2.5	Cosmic Ray Shield	69
5.2.6	Calorimeter	69
5.2.7	Spectrometer Calibration System	69
5.3	R&D Cost Estimate	70
II	<i>Mu2e</i> in Detail	71
6	Proton Beam	73
6.1	Overview	73
6.2	Delivering Protons to the Accumulator	73
6.3	Momentum Stacking and Rebunching	76

CONTENTS

6.4	Resonant Extraction	77
6.5	Proton Extinction	79
6.6	Total Proton Delivery	82
6.7	Radiation Safety	83
7	Production Target	85
7.1	Introduction	85
7.2	Muon Production	87
7.3	Stopped Muon Yield	91
7.4	Target Heating	93
7.5	Target Cooling	94
7.6	Heat and Radiation Shield	98
8	Muon Beamline	101
8.1	Collimators	102
8.2	Absorbers	103
8.3	Muon Beam Stop	104
8.4	Vacuum System	107
9	Production, Transport, and Detector Solenoids	117
10	Detector	125
10.1	Muon Stopping Target	126
10.2	Muon Capture Monitor	129
10.2.1	Purpose and Method	129
10.2.2	Location for the Germanium Detector	130
10.2.3	Calibration	131
10.2.4	Selection of Germanium Spectrometer System	132
10.3	The Tracking Detector	133
10.3.1	Physics Requirements	134
10.3.2	Tracking Detector Overview	135
10.3.3	Tracking Detector Performance	138
10.3.4	Mechanical Construction	150
10.3.5	Pad Readout	151
10.3.6	Drift Gas	154
10.3.7	An Alternate Design	154
10.4	Electromagnetic Calorimeter	155
10.4.1	Calorimeter Description	157
10.4.2	Crystal Choice	158
10.4.3	Calorimeter Readout	161
10.4.4	Calorimeter Cooling	162
10.4.5	Calorimeter Mechanical Support	163
10.4.6	Calorimeter Radiation Dose	164

CONTENTS

10.4.7	Calorimeter Performance	164
10.4.8	Calorimeter Calibration	164
11	Cosmic Ray Shield	167
11.1	Cosmic Ray Background Rate Calculation	167
11.2	Beam Induced Rates	171
11.3	Passive Cosmic Ray Shield	172
11.4	Active Cosmic Ray Shield	172
11.5	Scintillator Strips	175
11.6	Wavelength Shifting Fiber	175
11.7	Photomultiplier Tube and Signal Response	176
11.8	Assembly and Installation of a Three-Layer Module	178
11.9	Calibration of Active CR Shield	179
12	Trigger & Data Acquisition System	181
12.1	Overview	181
12.2	Tracker Front End Electronics	183
12.3	Calorimeter Digitizer Modules	186
12.4	Level-1 Trigger	188
12.5	Event Building and Processing	189
12.6	Slow Control and Monitoring	191
13	Offline Computing and Data Analysis	193
14	Detector Enclosure, Civil Construction, and Infrastructure	195
14.1	Beamline	195
14.2	Detector Enclosure	196
14.3	Electrical Power	199
14.4	Cryogenics	199
15	Environment, Safety and Health	203
15.1	Antiproton Accumulator and Debuncher Rings	203
15.2	Proton Beamline	204
15.3	Detector Enclosure	204
16	<i>Mu2e</i> in the Project X Era	205
16.1	<i>Mu2e</i> and Future Accelerator Upgrades	205
16.1.1	SNuMI	205
16.1.2	Project X	206
16.2	Detector upgrades	207
	References	208

List of Figures

1.1	Massive neutrino contribution to $\mu \rightarrow e\gamma$	4
1.2	Sensitivity of $\mu \rightarrow e$ conversion and $\mu \rightarrow e\gamma$ to new physics	6
1.3	Sensitivity of $C \equiv B(\mu \rightarrow e\gamma)/B(\mu \rightarrow e)$ to the sign of μ	7
1.4	$\mu \rightarrow e$ conversion rate for different nuclei	8
1.5	MSSM slepton–neutralino contribution to $\mu \rightarrow e\gamma$	9
1.6	$\mu \rightarrow e$ conversion rate for different SUSY–GUT scenarios	10
1.7	$\mu \rightarrow e$ versus $\mu \rightarrow e\gamma$ for different littlest Higgs scenarios	11
1.8	MSSM tree-level R-parity violating contribution to $\mu \rightarrow e$ conversion	12
1.9	Branching ratios for muon conversion as a function of $ U_{e3} \cos \delta$	15
2.1	History of charged lepton flavor violation searches	18
2.2	Layout of the <i>Mu2e</i> apparatus	19
2.3	The proton beam bunch scheme	23
2.4	Cutaway view of the Production Solenoid	25
2.5	Cutaway view of the Transport Solenoid	26
2.6	Cutaway view of the detector solenoid	27
2.7	Polar angle distribution of accepted conversion electrons	27
3.1	The muon stop time distribution	37
3.2	The distribution of DIO and signal events for a BR of 10^{-16}	41
3.3	Number of signal events vs. momentum threshold	41
3.4	The number of background events vs. momentum threshold	42
3.5	The figure of merit vs. lower cut on momentum	42
3.6	Form factor and cross section for electrons scattering on aluminum	43
3.7	Distribution in the π arrival time for late arriving pions	46
3.8	Antiproton differential cross section	48
3.9	Scatterplot of the \bar{p} annihilation positions in the transport system	49
3.10	The expected momentum distribution for decay-in-orbit events	54
3.11	The momentum for conversion events at the tracker	55
3.12	One side of the distribution of (reconstructed - true) momentum for decay in orbit events at the tracker.	56
3.13	Momentum distribution of DIO events for normal and 90% fields	56
3.14	Acceptance vs. momentum for DIO events	57

LIST OF FIGURES

3.15	Path of calibration electron through reconfigured calorimeter	59
4.1	<i>Mu2e</i> technically limited schedule for the superconducting solenoid	64
6.1	Boomerang scheme	74
6.2	Timeline of booster batches to NO ν A and <i>Mu2e</i>	75
6.3	Extraction area in the Recycler Ring	76
6.4	Momentum stacking	77
6.5	The hybrid rebunching scheme	78
6.6	Proposed location for resonant extraction system	79
6.7	Preliminary simulation of a third order resonance in Debuncher . .	79
6.8	The AC dipole scheme for beam extinction	81
6.9	Preliminary extinction channel design.	82
7.1	Production solenoid	86
7.2	Production target installed in PS	87
7.3	GEANT simulation of the Production Solenoid	88
7.4	Negative pion cross sections	90
7.5	Muon creation points in the Production Solenoid	91
7.6	Stopped muon yield per primary proton	92
7.7	Production target power distribution	93
7.8	Cross-sectional view of target cooling design	95
7.9	Target and coolant temperature	95
7.10	Conceptual design target cooling system	97
7.11	Cutaway view of the radiation shield	99
8.1	Particle arrival time at the muon stopping target.	101
8.2	Particle momentum at entrance to the Detector Solenoid.	101
8.3	Transport Solenoid collimators	103
8.4	The Proton and Neutron Shields in the Detector Solenoid	104
8.5	Late arrival time electron trajectories	106
8.6	Particle trajectories at the Muon Beam Stop	106
8.7	Muon Beam Stop design	107
8.8	Muon decay-in-orbit endpoint energy vs atomic number	108
8.9	DS vacuum closure elements	109
8.10	End view of DS showing cryo pumps	110
8.11	DS pressure during initial pump-down	111
8.12	DS pressure vs time after pump-down	111
8.13	Anti-proton stopping window	112
8.14	Vacuum closure with vacuum pumps	114
8.15	Drawing of the IFB	114
8.16	Production solenoid vacuum closure	115
9.1	Layout of <i>Mu2e</i> magnetic elements	118

LIST OF FIGURES

9.2	View of the assembled magnet system	118
9.3	Proton beam interface	119
9.4	Layout of the Solenoid Magnets.	123
10.1	The <i>Mu2e</i> detector in the Detector Solenoid	125
10.2	The stopped muon lifetime vs Z.	127
10.3	The muon decay-in-orbit endpoint energy vs Z.	127
10.4	The muon conversion rate vs. Z	128
10.5	The muon capture rate and decay-in-orbit rate vs Z.	128
10.6	Layout of Muon Beam Monitor	131
10.7	Pulse-height spectrum of ^{152}Eu . Energies in keV.	132
10.8	Acceptance as a function of σ	135
10.9	Cross section of the tracking detector showing electron trajectories	137
10.10	The differential and integral energy distributions for μ decay in orbit	140
10.11	The kinetic energy of protons originating from μ^- capture on Al	141
10.12	Simulated conversion electron trajectory.	144
10.13	Cluster number and pitch angle distributions	144
10.14	Detail of the L-tracker straws and pads	152
10.15	Sketch of a typical set of eight pads with the cables	152
10.16	Equivalent circuit to study the effect of cathode resistivity	153
10.17	Induced signal on the strips as a function of the cathode resistivity	154
10.18	Isometric view of the electromagnetic calorimeter	156
10.19	Calorimeter front view	157
10.20	Incident angle of electrons at the calorimeter	158
10.21	Calorimeter geometrical acceptance	158
10.22	Energy resolution for incident photons measured by a PWO-II array	160
10.23	Integral light yield from PWO-II crystals	160
10.24	PbWO ₄ crystal, 2 APDs, and a preamplifier	161
10.25	Carbon-epoxy square-cell honeycomb support	163
10.26	Front view of the calorimeter	163
10.27	Calorimeter trigger acceptance and trigger rate	165
11.1	End view of Cosmic Ray Shield	168
11.2	Top view of Cosmic Ray Shield	169
11.3	Flux of cosmic ray muons at the earth's surface	170
11.4	Cosmic ray flux vs. zenith angle at different depths	170
11.5	Sectional view of a full Cosmic Ray Shield module	173
11.6	Cross-sectional view of scintillator with embedded waveshifting fibers	174
11.7	Isometric view of front of the H8711 PMT assembly	178
11.8	Isometric view of back of H8711 PMT assembly	178
11.9	Optical readout of a full Cosmic Ray Shield module	179
12.1	Schematic overview of the DAQ system.	182

LIST OF FIGURES

12.2	Block diagram of the front-end readout electronics	185
12.3	A block diagram of the Calorimeter Digitizer Module.	187
12.4	Level-1 vane trigger	189
12.5	Schematic of the data flow in the Event Builder	190
12.6	Schematic of the data flow in the Processor Farm	191
14.1	A conceptual view of the site location.	196
14.2	Drawing of the site relative to the Accelerator Complex	197
14.3	Beamline from the Accumulator/Debuncher to the Detector Hall .	198
14.4	A plan view of the trench for the solenoid.	200

List of Tables

2.1	<i>Mu2e</i> parameters	21
2.2	Expected sensitivity for a two-year (2×10^7 s) run.	22
3.1	The \bar{p} induced backgrounds for different incident proton momenta. . .	49
3.2	Summary of the backgrounds from various sources	53
4.1	Preliminary estimate of the total project cost.	63
5.1	R&D request	70
6.1	Resonant extraction parameters	80
6.2	The most significant beam backgrounds	80
7.1	Optimization of stopped muon yield	92
7.2	Heat shield radiation heat loads	100
8.1	TS collimator parameters	102
8.2	Particle fluxes per beam proton	105
8.3	Neutron flux in the calorimeter from the Muon Beam Stop	107
8.4	Calorimeter IFB cabling	113
8.5	Tracker IFB cabling	113
9.1	Field specification characteristics	120
9.2	Coil Set Inductance, Operating Current and Minimum Charging Rate	121
10.1	Transition energies for muonic atoms: Aluminum and Titanium . . .	129
10.2	Peak detector rates at the beginning of the measurement interval . .	142
10.3	Selection criteria used in the electron momentum measurement. . . .	146
10.4	Pattern recognition results for two different detector lengths	148
10.5	Number of detector channels	150
10.6	Calorimeter parameters	156
10.7	Calorimeter crystal parameters	159
12.1	Design parameters for the ELEFANT and an upgraded A/D chip . .	184

Executive Summary

WE propose a new experiment, *Mu2e*, to search for charged lepton flavor violation with unprecedented sensitivity. We will measure the ratio of the coherent neutrinoless conversion in the field of a nucleus of a negatively charged muon into an electron to the muon capture process:

$$R_{\mu e} = \frac{\mu^- + A(Z, N) \rightarrow e^- + A(Z, N)}{\mu^- + A(Z, N) \rightarrow \nu_\mu + A(Z - 1, N)},$$

with a sensitivity $R_{\mu e} \leq 6 \times 10^{-17}$ at 90% CL. This is almost a four order-of-magnitude improvement over the existing limit. The observation of such a process would be unambiguous evidence of physics beyond the Standard Model.

Since the discovery of the muon in 1936, physicists have attempted to answer I.I. Rabi’s famous question: “Who ordered that?” Why is there a muon? What role does it play in the larger questions of why there are three families and flavors of quarks, leptons, and neutrinos? We know quarks mix through a mechanism described by the Cabbibo-Kobayashi-Maskawa matrix, which has been studied for forty years. Neutrino mixing has been observed in the last decade, but mixing among the family of *charged* leptons has never been seen. The current limits are of order $10^{-11} - 10^{-13}$ so the process is rare indeed.

Why is such an experiment important and timely? A major motivation for experiments at the Large Hadron Collider (LHC) is the possible observation of supersymmetric particles in the TeV mass range. Many of these supersymmetric models predict a μ - e conversion signal at $R_{\mu e} \sim 10^{-15}$. We propose to search for μ - e conversion at a sensitivity that exceeds this by more than an order of magnitude. The LHC may not be able to conclusively distinguish among supersymmetric models, so *Mu2e* will provide invaluable information should the LHC observe a signal. In the case where the LHC finds no evidence of supersymmetry, or other beyond-the-standard-model physics, *Mu2e* will probe for new physics at mass scales up to 10^4 TeV, far beyond the reach of any planned accelerator.

The competition for this physics is quite limited. The most relevant alternative process, $\mu \rightarrow e\gamma$, is roughly one order of magnitude less sensitive in a number of supersymmetric models and cannot probe the mass reach of *Mu2e* for new Z bosons, leptoquarks, or other non-supersymmetric processes. Muon-to-electron conversion is

unique and vital, both complementing the LHC and extending beyond it.

The importance of the search for charged lepton flavor violation through muon-to-electron conversion was emphasized by the P5 subpanel charged by HEPAP to present a strategic plan for the field for the next ten years. They state in their report, *US Particle Physics: Scientific Opportunities, A Strategic Plan for the Next Ten Years*:

A muon-to-electron conversion experiment at Fermilab could provide an advance in experimental sensitivity of four orders of magnitude. The experiment could go forward in the next decade with a modest evolution of the Fermilab accelerator complex. Such an experiment could be the first step in a world-leading muon-decay program eventually driven by a next-generation high-intensity proton source. Development of a muon-to-electron conversion experiment should be strongly encouraged in all budget scenarios considered by the panel.

The experiment will be mounted at Fermilab with only minor modifications to the existing accelerator complex. The Fermilab Booster will accelerate the 8 GeV protons needed to produce the intense muon beam for *Mu2e*. The protons will be transferred from the Booster via the Recycler to the Accumulator and Debuncher rings where they will be stacked, bunched, and extracted to a new experimental hall housing the *Mu2e* apparatus. The delivery of protons to the NO ν A neutrino production target will be unaffected. The apparatus design is taken from the MECO proposal. It employs an ingenious system of three superconducting solenoids to produce and transport the intense muon beam, and to detect the conversion electrons from the stopped muons. An advanced conceptual design for the solenoids exists.

Detailed cost estimates of the apparatus and solenoid system done by MECO, along with preliminary estimates of the cost of accelerator modifications, a new beam-line, and a new experimental hall indicate a project cost of approximately \$180M (FY09\$), including an overall 50% project contingency. About \$10M in R&D funds are also needed. The experiment could begin taking data in 2016 and would need two years running to achieve the proposed sensitivity goal.

Project X upgrades to Fermilab could allow an upgraded *Mu2e* to increase the sensitivity by as much as two orders of magnitude. Should we observe a signal in *Mu2e*, the greatly improved statistics provided by Project X could then allow detailed measurements of the source and magnitude of the effect.

The schedule is timely and will provide answers just as they are needed. This experiment is a unique opportunity to help unravel the mysteries of flavor physics and answer Rabi's question that has puzzled us for more than seventy years.

Part I

Mu2e in Brief

Theoretical Motivation

1.1 Introduction

WITH the discovery of neutrino masses and lepton mixing, the fact that individual lepton-flavor numbers — electron-number, muon-number, and tau-number — are not conserved has been established [1–13]. All such violating effects to date have been observed in the neutral lepton sector, through the phenomenon of neutrino oscillations. Charged-lepton flavor-violation (CLFV), on the other hand, has been the subject of intense experimental searching since the discovery of the muon but, to this date, no evidence for it has ever been uncovered.

The Standard Model augmented by new physics that leads to the experimentally observed neutrino masses predicts a non-zero rate for CLFV process, but expectations depend dramatically on the mechanism responsible for neutrino mass generation. For example, if the physics responsible for neutrino masses is very heavy (as in the case of a high mass scale seesaw mechanism [14–17]) or very weakly coupled (as in the case of Dirac neutrinos), expectations for CLFV processes are around forty orders of magnitude smaller than current experimental bounds. The reason for this is that the active neutrino contribution is GIM suppressed, such that the amplitude for CLFV is proportional to the tiny neutrino mass-squared differences [18]. For example, the massive neutrino contribution (Fig. 1.1) to $\mu \rightarrow e\gamma$ is

$$Br(\mu \rightarrow e\gamma) = \frac{3\alpha}{32\pi} \left| \sum_{i=2,3} U_{\mu i}^* U_{ei} \frac{\Delta m_{1i}^2}{M_W^2} \right|^2 < 10^{-54}, \quad (1.1)$$

where $U_{\alpha i}$, $\alpha = e, \mu, \tau$ and $i = 1, 2, 3$, are the elements of the neutrino mixing matrix and M_W is the W -boson mass.

On the other hand, certain neutrino mass generating mechanisms are already disfavored due to the fact that CLFV has yet to be observed. It is fair to say that searches for CLFV are bound to play a key role as far as uncovering the origin of neutrino masses. Moreover, like other flavor-changing neutral current processes, searches for CLFV are also among the most powerful and promising probes of new




Figure 1.1: Massive neutrino contribution to the charged lepton flavor-violating muon decay $\mu \rightarrow e\gamma$. ν_i are neutrino mass eigenstates, while $U_{\alpha k}$, $\alpha = e, \mu, \tau$ and $k = 1, 2, 3$, are the elements of the lepton mixing matrix.

physics at or even above the TeV scale, regardless of its connection to neutrino masses. Concrete examples will be discussed in the next subsection.

Among the different CLFV channels, three rare muon processes stand out, thanks in part to the muon’s small mass and long lifetime: $\mu \rightarrow e\gamma$, $\mu \rightarrow eee$ and $\mu \rightarrow e$ -conversion in nuclei. Current experiments have been able to rule out, at the 90% confidence level, $\mu^+ \rightarrow e^+\gamma$ with branching ratios above 1.2×10^{-11} [19] and $\mu^+ \rightarrow e^+e^-e^+$ with branching ratios above 1.0×10^{-12} [20]. The rate for $\mu^- + {}^{48}\text{Ti} \rightarrow e^- + {}^{48}\text{Ti}$ normalized to the capture rate ($\mu \rightarrow e$ conversion in titanium), is constrained, at the 90% confidence level, to be less than 4.3×10^{-12} [21], while that for $\mu \rightarrow e$ conversion in gold is constrained to be less than 7×10^{-13} [22, 23]. The concurrent exploration of all three rare muon processes is of the utmost importance given that the three process “feel” different types of new physics in distinct ways. If CLFV is observed in any one of these processes, results from other searches will play a fundamental role as far as establishing the nature of the lepton-flavor violating new physics.

Depending on the nature of the CLFV physics, one of the three bounds listed above turns out to be the most significant. For a particular class of models, including several of the standard supersymmetric ones, efforts to observe $\mu \rightarrow e\gamma$ prove to be most promising currently and in the immediate future. The MEG experiment [24], currently taking data at PSI, is ultimately aiming at being sensitive to $\mu \rightarrow e\gamma$ branching ratios larger than several times 10^{-14} . However, given the existence of very intense future muon sources, $\mu \rightarrow e$ -conversion will likely serve as the deepest probe of CLFV, superior to $\mu \rightarrow e\gamma$ in its new physics reach regardless of the nature of the new physics. Among other factors, it is this “feature,” which will be discussed in more detail in the next subsection, that drives us to concentrate on the CLFV process where a nuclear-captured muon converts into an electron — $\mu \rightarrow e$ -conversion in nuclei.

Negatively charged muons that stop in matter are quickly trapped and form muonic atoms, which undergo electromagnetic transitions until the muon is in the $1s$ orbital. Trapped muons either Michel-decay or convert into neutrinos in the field

Theoretical Motivation

of the nucleus:

$$\mu^- + (A, Z) \rightarrow \nu_\mu + (A, Z - 1), \quad (1.2)$$

where (A, Z) represents a nucleus with mass number A and atomic number Z .

Similarly, the $\mu \rightarrow e$ -conversion process is characterized by

$$\mu^- + (A, Z) \rightarrow e^- + (A, Z). \quad (1.3)$$

Instead of discussing the rate Γ for this muon and electron number violating process, it is convenient to define the normalized capture rate

$$B(\mu \rightarrow e - \text{conv}) \equiv \frac{\Gamma(\mu^- + (A, Z) \rightarrow e^- + (A, Z))}{\Gamma(\mu^- + (A, Z) \rightarrow \nu_\mu + (A, Z - 1))}, \quad (1.4)$$

to which we will often refer to as the ‘ $\mu \rightarrow e$ -conversion rate.’

1.2 Comparison to Current HEP Program

It is important to place searches for CLFV in general and $\mu \rightarrow e$ -conversion in particular in the larger context of the current and near future developments of the high energy physics program. We will assume that next-generation neutrino oscillation experiments will take place and provide nontrivial information regarding lepton mixing and other new “neutrino physics.” We will also assume that the LHC experiments will have taken enough data in order to provide a clearer picture of physics at around the TeV scale. Finally, we assume that one will either have observed $\mu \rightarrow e\gamma$ or constrained its branching ratio to be less than 10^{-14} . It appears unlikely that a future experiment will be able to significantly improve on this, regardless of whether very intense muon sources are available. This is *not* the case of future searches for $\mu \rightarrow e$ -conversion, as will be discussed later.

1.2.1 Model Independent Analysis: $\mu N \rightarrow e N$ vs. $\mu \rightarrow e\gamma$

One can estimate the sensitivity of CLFV processes to new physics in a model independent way by adding to the standard model effective operators that violate lepton flavor. For concreteness, consider the effect of the standard model augmented by the following CLFV effective Lagrangian:

$$\mathcal{L}_{\text{CLFV}} = \frac{m_\mu}{(\kappa + 1)\Lambda^2} \bar{\mu}_R \sigma_{\mu\nu} e_L F^{\mu\nu} + \frac{\kappa}{(1 + \kappa)\Lambda^2} \bar{\mu}_L \gamma_\mu e_L \left(\sum_{q=u,d} \bar{q}_L \gamma^\mu q_L \right), \quad (1.5)$$

where Λ is the scale of new physics and κ measures whether the dominant new physics contribution to CLFV comes in the form of a dimension-five, CLFV magnetic moment-type operator ($\kappa \ll 1$) or from a CLFV four-fermion interaction ($\kappa \gg 1$).




Figure 1.2: Sensitivity of a $\mu \rightarrow e$ conversion in ^{48}Ti experiment that can probe a normalized capture rate of 10^{-16} and 10^{-18} , and of a $\mu \rightarrow e\gamma$ search that is sensitive to a branching ratio of 10^{-13} and 10^{-14} , to the new physics scale Λ as a function of κ , as defined in Eq. (1.5). The dimensionless parameter κ interpolates between a flavor-violating magnetic moment-type operator ($\kappa \ll 1$) and a flavor-violating four-fermion operator ($\kappa \gg 1$). Also depicted is the currently excluded region of this parameter space.

The effective Lagrangian above will mediate both $\mu \rightarrow e\gamma$ and $\mu \rightarrow e$ -conversion (and $\mu \rightarrow eee$, which will not be discussed in any detail). While a handful of other effective operators can also contribute, the ones above contain qualitatively the predictions of most distinct new physics scenarios as far as $\mu \rightarrow e\gamma$ and $\mu \rightarrow e$ -conversion are concerned. The sensitivity of different CLFV probes to Λ as a function κ is depicted in Fig. 1.2. Note that, regardless of the value of κ , a $\mu \rightarrow e$ -conversion experiment sensitive to capture rates above 10^{-16} probes Λ values smaller than a few thousand TeV.

For $\kappa \ll 1$, the normalized $\mu \rightarrow e$ -conversion is around several times 10^{-3} of the branching ratio for $\mu \rightarrow e\gamma$, while for $\kappa \gg 1$ the normalized capture rate for $\mu \rightarrow e$ -conversion is many orders of magnitude larger than the branching ratio for $\mu \rightarrow e\gamma$. Hence, a $\mu \rightarrow e$ -conversion experiment sensitive to normalized rates above 10^{-16} is at

Theoretical Motivation




Figure 1.3: $C \equiv B(\mu \rightarrow e\gamma)/B(\mu \rightarrow e)$ in the MSSM with MSUGRA boundary conditions for the soft parameters and neutrino masses induced by the seesaw mechanism, as a function of $\tan\beta$ for different signs of the μ -parameter. See Ref. [25] for details.

least as sensitive to new physics as a $\mu \rightarrow e\gamma$ experiment sensitive to branching ratios above a few $\times 10^{-14}$, regardless of the nature of the new physics. It is important to emphasize that while we are using Eq. (1.5) to make this point, this conclusion is very generic and applies to most new physics scenarios that have been explored in the literature to date.

In the case of a positive CLFV signal in either $\mu \rightarrow e\gamma$ or $\mu \rightarrow e$ -conversion, combined results from different CLFV processes provide detailed information regarding the new physics. For example, should the world be properly described by Eq. (1.5), a measurement of $\mu \rightarrow e\gamma$ and $\mu \rightarrow e$ -conversion allows one to determine both Λ and κ independently, while a single measurement can only determine a combination of the two new physics parameters. In general, it is well known that a comparison of $B(\mu \rightarrow e - \text{conv})$ and $B(\mu \rightarrow e\gamma)$ helps distinguish among models or even measure the value of new physics parameters. A concrete example is depicted in Fig. 1.3, where the ratio of branching ratios $C \equiv B(\mu \rightarrow e\gamma)/B(\mu \rightarrow e - \text{conv})$ is plotted as a function of $\tan\beta$ in the case of the MSSM with MSUGRA boundary conditions for the soft SUSY breaking parameters [25]. One can see that a precise measurement of C can determine the sign of the MSSM μ -parameter, especially if $\tan\beta$ is not too large.

The effective Lagrangian that describes $\mu \rightarrow e$ -conversion and $\mu \rightarrow e\gamma$ contains, in general, several other dimension-six operators not in Eq. (1.5), including those with different muon and electron chiralities and scalar-scalar four-fermion operators [26]. Information regarding all the different parameters that describe CLFV can be obtained from the CLFV probes themselves. In the event of a positive signal for $\mu \rightarrow e$ -conversion, details of the effective Lagrangian can be obtained by comparing




Figure 1.4: $\mu \rightarrow e$ conversion rate for different nuclei, normalized to that for $\mu \rightarrow e$ conversion in aluminum. The different curves represent the contribution of different types of higher dimensional operators. See Ref. [27] for details.

the rate for $\mu \rightarrow e$ -conversion in different nuclei, since different nuclei are sensitive to new physics in distinct ways, as depicted in Fig. 1.4 [27]. This flexibility is not shared by $\mu \rightarrow e\gamma$ (where one can only hope to measure, in principle, the final state photon or electron polarizations [28]). In the case of a positive signal in $\mu \rightarrow eee$, some detailed information regarding the underlying physics can also be obtained by analyzing in detail the kinematics of the three final state leptons. See, for example [29–31].

1.2.2 CLFV and new physics at the TeV scale

By the end of 2009, we expect the LHC experiments to start accumulating data that will reveal the mechanism of electroweak symmetry breaking and explore the physics of the TeV scale. Several theoretically motivated scenarios predict the existence of new degrees of freedom with masses at or below 1 TeV and, if this is the case, one expects some of these new states to be discovered at the LHC.

New degrees of freedom at the TeV scale are expected to mediate CLFV processes. Expectations are model-dependent, but detailed computations in specific models lead to CLFV rates very close to current experimental bounds, as will be discussed in more detail shortly. We first conservatively assume that the new physics will predominantly induce flavor-violating magnetic-moment type effective interactions at the one-loop level. A concrete example is depicted in Fig. 1.5.

In this case, CLFV is given by Eq. (1.5) (potentially augmented by similar oper-

Theoretical Motivation




Figure 1.5: MSSM slepton–neutralino contribution to $\mu \rightarrow e\gamma$. $\Delta m_{\tilde{\mu}\tilde{e}}^2$ stands for the insertion of an off-diagonal element of the slepton mass-matrix. From [26].

ators where the electron and muon chiralities are reversed) with $\kappa \ll 1$ and

$$\frac{1}{\Lambda^2} \sim \frac{eg^2}{16\pi^2} \frac{\theta_{e\mu}}{M_{\text{new}}^2}. \quad (1.6)$$

M_{new} are the masses of the new states that couple to standard model fields with coupling g and $\theta_{e\mu}$ is a flavor-violating factor, most likely inaccessible to the LHC. If one assumes g (henceforth assumed to be of order one) and M_{new} to be known, failure to observe CLFV translates into bounds on $\theta_{e\mu}$.

As a concrete example, consider the possibility that the currently observed discrepancy between the standard model prediction and the measurement of the muon anomalous magnetic moment [32] (a $> 3\sigma$ effect) is due to new electroweak scale physics. In this case, the new physics contribution to the anomalous magnetic moment of the muon is captured by a flavor-conserving version of the magnetic moment-type operator that mediates CLFV. Current data on the muon anomalous magnetic moment translates into a measurement of a combination of the new physics scale M_{new} and the new coupling g , in which case current bounds on CLFV are already quite severe and constrain $\theta_{e\mu} \lesssim 10^{-4}$ [33].

Similarly, if the LHC discovers new states with masses M_{new} around 1 TeV, current bounds from CLFV will already translate into $\theta_{e\mu} > \sim 10^{-3}$. In this case, $\mu \rightarrow e$ -conversion experiments sensitive to conversion rates greater than 10^{-16} will probe $\theta_{e\mu} < \sim 10^{-5}$.

What are the expected values for $\theta_{e\mu}$? The answer to this is model dependent, but one can identify general categories. Generic new physics models predict $\theta_{e\mu} \sim 1$, in which case searches for CLFV already rule out $M_{\text{new}} \sim 10$ TeV. Searches for CLFV, along with other flavor observables, already constrain any physics at the TeV scale to be flavor conserving at the leading order. For this reason, one often assumes that the only sources of lepton-flavor violation are the ones already present in the standard model, *i.e.*, the charged-lepton Yukawa couplings and the neutrino mass generating sector (for one example, see [34]). Under this assumption, $\theta_{e\mu}$ values can be computed on a case by case basis, and its value may or may not dependent on the unknown

Mu2e Proposal




Figure 1.6: The $\mu \rightarrow e$ conversion rate in Ti for different SUSY-GUT scenarios. The plots are obtained by scanning the LHC accessible parameter space. The horizontal lines are the present (SINDRUM II) bound and the planned (Future) sensitivity to the process both at the proposed PRIME experiment in JPARC and at the proposed Mu2e experiment at Fermilab. See Ref. [35] for details.

new physics responsible for neutrino masses and lepton mixing.

Several detailed analyses have been performed for different independently motivated new physics scenarios, including models with weak scale supersymmetry, models with flat and warped extra-dimensions, and little Higgs models. Some results depend on details of the physics responsible for neutrino masses, about which we will discuss more shortly, but all tend to lead to $\theta_{e\mu}$ values such that $\mu \rightarrow e$ -conversion is “guaranteed” to happen with rates above 10^{-17} or so as long as the new physics is observable at the LHC.

Two examples are depicted in Figs. 1.6 [35] and 1.7 [36]. Fig. 1.6 depicts the result of a scan of the MSSM parameter space for different SUSY-GUT scenarios where neutrino masses are generated via the seesaw mechanism. The GUT hypothesis fixes the values of the right-handed neutrino Majorana masses, while there remains the freedom to choose the off-diagonal structure of the neutrino Yukawa couplings. In [35] two different choices are made: the neutrino Yukawa coupling matrix is PMNS-like, *i.e.*, all its mixing angles are large, as in the physically observable lepton mixing matrix, or it is CKM-like, *i.e.*, all its mixing angles are small, as in the physically observable quark mixing matrix. While the different choices lead to $\mu \rightarrow e$ -gamma rates that vary by more than four orders of magnitude, it is clear that a $\mu \rightarrow e$ -conversion experiment sensitive to normalized rates above 10^{-17} or so should cover the majority of the LHC accessible parameter space.

Figure 1.6 depicts the result of a scan of the parameter space of the littlest Higgs model with T-parity [36]. The different colored (shaded) points refer to different ansatze for the structure of the mirror lepton mixing sector, not dissimilar from the




Figure 1.7: $\mu \rightarrow e$ conversion rate in Ti versus $\mu \rightarrow e\gamma$ branching ratio for different littlest Higgs scenarios. The light grey region is allowed by current searches for CLFV. The different shaded points represent different ansätze for the mirror fermion mixing matrix. From Ref. [36].

choice of neutrino Yukawa matrices made in the SUSY example discussed above. Also here, a $\mu \rightarrow e$ -conversion experiment sensitive to normalized conversion rates above 10^{-16} should cover the parameter space explored in the figure. This is also true for points in the parameter space where the branching ratio for $\mu \rightarrow e\gamma$ is less than 10^{-14} . Note that in this case results do not depend on the mechanism responsible for neutrino masses, but do depend on the unknown mirror fermion mixing matrix.

It is also important to discuss the case where CLFV is generated by new physics at the tree-level, *i.e.*, it is a consequence of the simple exchange of a heavy new physics particle. An example is depicted in Fig. 1.8 [37]. Other than SUSY with R-parity violation, depicted in Fig. 1.8, several well-motivated new physics scenarios lead to similar CLFV effects including the models with lepto-quarks, neutrino mass models with Higgs triplets [38], and models with extra Z' gauge bosons [39]. In the case of SUSY with R-parity violation, if the dominant relevant R-parity violating couplings are the ones depicted in Fig. 1.8, the $\mu \rightarrow e$ -conversion rate is expected to be 2×10^5 larger than the branching ratio for $\mu \rightarrow e\gamma$ or $\mu \rightarrow eee$ [37]. Other examples that span different theoretical possibilities are discussed in [37].

In this case, CLFV is described by Eq. (1.5) (potentially augmented by, say, scalar-scalar four-fermion operators) with $\kappa \gg 1$ and

$$\frac{1}{\Lambda^2} \sim \frac{g^2 \theta_{e\mu}}{M_{\text{new}}^2}. \quad (1.7)$$

Here, if M_{new} is measured at the LHC, current bounds from $\mu \rightarrow e$ -conversion con-




Figure 1.8: MSSM tree-level R -parity violating contribution to $\mu \rightarrow e$ conversion [37].

strain $g^2\theta_{e\mu}$ to be tiny. In the example depicted in Fig. 1.8, $g^2\theta_{e\mu} \sim \lambda'_{221}\lambda'_{121}$, which is, not surprisingly, most severely constrained by searches for $\mu \rightarrow e$ -conversion in nuclei.

To summarize the discussion so far: if the LHC discovers new states at the TeV scale, several distinct new physics scenarios predict CLFV violating processes to occur with rates that are close to current experimental bounds. In this case, a positive $\mu \rightarrow e$ -conversion result (which may or may not be accompanied by a positive $\mu \rightarrow e\gamma$ or $\mu \rightarrow eee$ result) will tell us about the flavor structure of the new physics sector, and may help distinguish among different new physics scenarios. It is important to emphasize that the information one will extract from the CLFV sector is complementary to the information one can hope to extract from LHC data.

Similarly, in the advent of new physics at the LHC, a negative $\mu \rightarrow e$ -conversion result will also contribute to our understanding of the new TeV scale physics. It would reveal that (i) the new physics is indeed intrinsically lepton-flavor conserving and (ii) the flavor breaking effects induced by the known sources of flavor violation are smaller than naive expectations. Different physics may explain (ii). For example, in the case of SUSY, low-energy gauge-mediated scenarios [40] usually lead to suppressed flavor-violating effects. Furthermore, as will be discussed in more detail shortly, some of the estimates above depend on the physics responsible for neutrino masses, which may be such that TeV-scale CLFV effects are smaller than naive expectations.

There remains, of course, the possibility that the LHC discovers no new degrees of freedom other than the standard model Higgs boson. This will indicate that the new physics beyond the standard model is significantly heavier than the TeV scale. Under these circumstances, searches for CLFV violation remain extremely valuable, especially searches for $\mu \rightarrow e$ -conversion (and, perhaps, searches for $\mu \rightarrow eee$). The reason is that if the LHC fails to discover any beyond-the-standard-model effect, the gauge hierarchy problem will, most likely, prove to be a poor indicator for the new physics scale. In this case, progress in fundamental particle physics, including the answers for the open questions, will have to rely, at least for some time, on indirect new physics probes, *i.e.*, on the “intensity frontier.” These include precision studies

of neutrinos and their properties, precision measurements of well-known standard model quantities (like the anomalous magnetic moment of the muon), and searches for forbidden or extremely suppressed processes, like CLFV muon processes. Of the probes listed above, CLFV is the one capable of “directly” reaching out to several thousand TeV, especially — in the case of $\mu \rightarrow e$ -conversion — if new flavor-violating effects are large (in the sense $\theta_{e\mu} \sim 1$) and strongly coupled ($g \sim 1$) and occur at the tree-level (see Eq. (1.7) and Fig. 1.2). While there are, by no means, any guarantees that this is the case, there are reasons to believe in such a scenario. Among other plausibility arguments is the fact that the only direct evidence for new physics is the existence of non-zero neutrino masses. The tiny neutrino masses are potentially due to new, heavy physics and the large lepton mixing angles seem to indicate that flavor-numbers are not conserved in the neutrino sector. Hence, as long as the neutrino mass scale is not too heavy (*i.e.*, below 10^4 TeV), there remains the possibility that $\mu \rightarrow e$ -conversion will directly teach us about the physics responsible for non-zero neutrino masses.

1.2.3 CLFV, neutrino masses, and the matter–antimatter asymmetry of the Universe

As discussed in the introduction, massive neutrinos and lepton mixing imply that CLFV *must* occur at some level. As already discussed, while the active neutrino contribution to CLFV is known to be tiny, there remains the possibility that the new physics responsible for neutrino masses will induce CLFV at observable levels. Several examples of this have been mentioned in passing above.

Here we will discuss in more detail the MSSM case in order to underline the potential importance of CLFV to understanding neutrino masses and other related phenomena, including leptogenesis [41]. In the MSSM with MSUGRA boundary conditions large CLFV is generated if neutrino masses are a consequence of the seesaw mechanism. In more detail, a non-zero off-diagonal slepton mass-squared (see Fig. 1.5)

$$\Delta m_{\mu\bar{e}}^2 \simeq \sum_k \frac{M_{\text{SUSY}}^2}{16\pi^2} y_{\mu k}^* y_{ek} \log \left(\frac{M_{\text{GUT}}}{M_k} \right), \quad (1.8)$$

is generated (for a detailed discussion see, for example, [42]). Here $y_{\alpha k}$, $\alpha = e, \mu, \tau$ and $k = 1, 2, 3$, are the neutrino Yukawa couplings, M_k are the right-handed neutrino Majorana masses, M_{SUSY} is a typical supersymmetric mass and M_{GUT} is a typical ultraviolet cut-off, often equated with the GUT scale. In this case, the rate for muon CLFV is proportional to $(\Delta m_{\mu\bar{e}}^2)^2$. Using the language of Eq. (1.6), $\theta_{e\mu} \sim \Delta m_{\mu\bar{e}}^2 / M_{\text{SUSY}}^2$. Hence, CLFV processes probe some combination of the neutrino Yukawa couplings and hence provides non-trivial information regarding the neutrino mass sector (this statement is very dependent on the physics of SUSY breaking, which must be well understood).

Knowledge of the neutrino Yukawa couplings is also fundamental when it comes to determining whether leptogenesis is the mechanism responsible for the matter–antimatter asymmetry of the Universe. The amount of baryon number generated depends on a linear combination of neutrino Yukawa couplings and right-handed Majorana neutrino masses different from the one above ($\propto \text{Im}[(y^\dagger y)_{j1}^2]$ in the case of thermal leptogenesis [43–45]), while a third linear combination determines the observed active neutrino masses:

$$m_{\alpha\beta} = \sum_k y_{\alpha k} \frac{v^2}{M_k} y_{\beta k}. \quad (1.9)$$

It has been shown that (see, for example, [46]), given the right circumstances, neutrino oscillation measurements combined with positive results from CLFV, positive results from searches for neutrinoless double-beta decay, and information regarding the low-energy SUSY and SUSY breaking, can provide enough information in order to test leptogenesis. Hence, if thermal leptogenesis is ever to be tested experimentally, CLFV will certainly play a fundamental role.

Negative results from CLFV, combined with the discovery of SUSY at the LHC, may prove as important as positive results. The reason is as follows. In the MSSM, standard thermal leptogenesis requires the lightest right-handed neutrino Majorana mass to be larger than 10^9 GeV or so. This translates into a rough lower bound on the neutrino Yukawa couplings, using Eq. (1.9): $y^2 < \sim 10^{-6}$. This, in turn, implies that $\theta_{e\mu} \sim 10^{-7}$. Hence, if CLFV experiments can rule out $\theta_{e\mu} > 10^{-7}$, standard thermal leptogenesis would be severely disfavored. Such a sensitivity (which corresponds to $\Lambda \sim 10^4$ TeV, $\kappa \ll 1$ in Fig. 1.2) can only be obtained in searches for $\mu \rightarrow e$ -conversion fed by very intense muon sources.

In the case of low energy SUSY with a high scale seesaw (discussed above), the relation between neutrino mixing parameters and CLFV is indirect. This is easy to see, as the rates for CLFV and the neutrino Majorana mass matrix depend on different combinations of the parameters of the seesaw Lagrangian (Yukawa couplings and right-handed neutrino Majorana masses). The same is true of the combination of parameters that determines the matter–antimatter asymmetry of the universe via leptogenesis. On the other hand, there are scenarios where the neutrino masses and the lepton mixing angles can be directly related to the rates of several CLFV processes. Examples include models of large extra-dimensions where the small neutrino Yukawa couplings are a consequence of the fact that right-handed neutrinos are allowed to propagate in all space dimensions, and models where neutrino Majorana masses are a consequence of the existence of $SU(2)$ triplet Higgs fields. Figure 1.9 depicts the rate for different muon CLFV processes as a function of the unknown oscillation parameters θ_{13} , for different hypothesis regarding the neutrino mass ordering (normal or inverted) [38]. The overall expectation for the transition rates depends on parameters external to the neutrino mass matrix, like the triplet mass and vacuum

Theoretical Motivation




Figure 1.9: The branching ratios B for $\mu \rightarrow e\gamma$ (solid line) and $\mu \rightarrow eee$ (dashed line), and the normalized capture rate B for $\mu \rightarrow e$ -conversion in Ti (dotted line) as a function of $|U_{e3}| \cos \delta$ in a scenario where neutrino masses arise as a consequence of the presence of a triplet Higgs field with a small vacuum expectation value. The lightest neutrino mass is assumed to be negligible while the neutrino mass hierarchy is assumed to be normal (left-hand side) and inverted (right-hand side). See Ref. [38] for details.

expectation value. The combination of data from neutrino oscillation experiments, high energy collider experiments (like the LHC) and CLFV should ultimately allow one to thoroughly test particular Higgs triplet models and, if these turn out to be correct, unambiguously reveal the physics behind neutrino masses. Note that the interference between “solar” and “atmospheric” parameters can lead to cancellations and a tiny rate for CLFV muon processes.

Overview of the Experiment

WE propose to measure the ratio of the rate of the neutrinoless, coherent conversion of muons into electrons in the field of a nucleus, relative to the rate of ordinary muon capture on the nucleus:

$$R_{\mu e} = \frac{\mu^- + A(Z, N) \rightarrow e^- + A(Z, N)}{\mu^- + A(Z, N) \rightarrow \nu_\mu + A(Z-1, N)}.$$

The conversion process is an example of charged lepton flavor violation (CLFV). Thus far no CLFV interaction has been observed experimentally; the current best experimental limit on muon-to-electron conversion, $R_{\mu e} < 6.1 \times 10^{-13}$ (90% CL), is from the SINDRUM II experiment, using a gold target [47]. Figure 2.1 shows the history of charged lepton flavor violation searches in muon decays. *Mu2e* intends to probe four orders of magnitude beyond the SINDRUM II sensitivity, measuring $R_{\mu e}$ with a sensitivity of 6×10^{-17} at 90% CL.

When a muon is stopped in a target it rapidly (10^{-16} s) cascades down to the 1S state. Several processes can then occur, the most likely being muon decay-in-orbit, $\mu^- N_{A,Z} \rightarrow e^- \nu_\mu \bar{\nu}_e N_{A,Z}$, and muon capture, $\mu^- N_{A,Z} \rightarrow \nu_\mu N_{A,Z-1}$, the former dominant for light nuclei and the latter dominant for heavy nuclei. Neutrinoless conversion of a muon into an electron, $\mu^- N_{A,Z} \rightarrow e^- N_{A,Z}$, results in an electron with an energy slightly less than that of the rest mass energy of the muon:

$$E_{\text{CE}} = m_\mu c^2 - B_\mu(Z) - C(A),$$

where Z and A are the number of protons and nucleons in the nucleus, B_μ is the atomic binding energy of the muon, and $C(A)$ is the nuclear recoil energy. In the case of muonic aluminum, $B_{\mu,1S}(13) = 0.48$ MeV, $C(27) = 0.21$ MeV, and $E_{\text{CE}}(\text{Al}) = 104.97$ MeV. This energy is well above the bulk of the dominant background, electrons from muon decay-in-orbit. The tail of the muon decay-in-orbit electron energy spectrum falls off as $(E - E_e)^5$ near the endpoint: only $\sim 3 \times 10^{-13}$ of the electrons have an energy within 3 MeV of the endpoint energy. This gives muon-to-electron conversion a major experimental advantage over most other CLFV experiments, and makes it potentially the most sensitive of all.




Figure 2.1: History of charged lepton flavor violation searches in muon decays. The MEG and Mu2e sensitivity goals are shown.

In order to achieve a sensitivity of $\sim 10^{-17}$, in excess of 10^{18} muons must be stopped; roughly 50 *billion* per second in a two-year running period. A continuous muon beam, with its attendant pions and electrons, would produce unacceptably high detector rates and backgrounds. To mitigate this, in *Mu2e* the muon beam is bunched with a separation of roughly twice the muon lifetime of 864 ns in Al. The detector is turned off for some 700 ns after the intense bunched beam impacts muon stopping target because of the large flux of particles emanating from the target, and then turned on until the next bunch arrives. The experimental signature is an isolated 105 MeV energy electron exiting the stopping target no less than ~ 700 ns after the bunched beam has left.

The *Mu2e* apparatus, shown in Fig. 2.2 is identical to the MECO design proposed for the AGS at Brookhaven [48]. It consists of three solenoidal magnets: the Production, Transport, and Detector solenoids, with fields ranging from 5.0 T at the far end of the Production Solenoid to 1.0 T at the opposite end of the detector solenoid. The bunched muon beam is produced as follows. The 8 GeV proton bunches from the Debuncher ring enter the Production Solenoid at a slight angle to its axis and impact a water-cooled Au target. Pions are produced, which decay into muons, and some

Overview of the Experiment

are captured in the graded field of the Production Solenoid and reverse direction, spiraling into the Transport Solenoid. The Transport Solenoid is curved in order to prevent line of sight transport of gammas and neutrons from the production target to the stopping target and to separate the positively and negatively charged particles. A collimator at the midpoint of the of the Transport Solenoid removes the positively charged particles. Approximately 0.25% of the protons impacting the production target produce muons that stop in one of the 17 0.2-mm thick Al foils that form the stopping target. Since the stopping target is in a graded magnetic field, particles emitted backward from the stopping target have their directions reversed.




Figure 2.2: Layout of the Mu2e apparatus. The proton beam enters the Production Solenoid from the left, and muons are collected and transported via the Transport Solenoid to the Detector Solenoid.

The detector consists of a magnetic spectrometer, which is the primary energy measuring device, and an electromagnetic calorimeter, which is used for triggering and to confirm the energy and position measurements of the magnetic spectrometer. The *Mu2e* detector has been designed: (1) to ensure that the intense non-stopped beam from the production target is transported with minimal interactions to a beam stop at the end of the Detector Solenoid, (2) to minimize the acceptance for electrons from backgrounds such as muon decay-in-orbit, and (3) to have excellent momentum resolution for $\sim 100 \text{ MeV}/c$ electrons. The detector sits in an evacuated vessel in the 1 T field of the Detector Solenoid. It must withstand instantaneous rates up to 200 kHz in individual detector elements.

Two stages of *Mu2e* could be envisaged. A first stage that we are proposing now will produce low energy muons using 8 GeV protons from a modestly upgraded Booster, Debuncher, and Accumulator rings, improving on the SINDRUM II limit by

a factor of 10^4 . A second stage, not part of this proposal, would employ the larger proton flux from Project X plus upgrades to the *Mu2e* apparatus to achieve up to another two orders of magnitude improvement in sensitivity in $R_{\mu e}$. If we were to observe a signal in Stage I, we could use this increased rate to explore the conversion process on a variety of nuclei.

2.1 Experimental Sensitivity

The experimental parameters on which the expected sensitivity is based are listed in Table 2.1. What we experimentally measure is the ratio of the muon-to-electron conversion rate to the muon capture rate in Al:

$$R_{\mu e} = \frac{\Gamma(\mu N \rightarrow e N)}{\Gamma(\mu N \rightarrow \nu_{\mu} N^*)} = \frac{N_{\mu e}/N_s \cdot 1/\epsilon_{\mu e}}{\Lambda_{\mu\nu}/\Lambda_{\text{tot}} (= 0.609)},$$

where: $N_{\mu e}$ is the number of events in the signal window $E_e > 103.6 \text{ MeV}$; N_s is the total number of muon stops; $\epsilon_{\mu e}$ is the overall detector acceptance (0.076), which is the product of the calorimeter trigger efficiency (0.80), the fraction of muon captures in the time window (0.51), and the reconstruction and selection efficiency (0.19); $\Lambda_{\mu\nu}$ is the muon capture rate; and Λ_{tot} is the total decay rate (in Al).

The expected single-event sensitivity for a two-year run ($2.0 \times 10^7 \text{ s}$) is 2.3×10^{-17} ; the 90% CL sensitivity is 5.7×10^{-17} . The factors entering into the sensitivity calculation are summarized in Table 2.2. Note that running time taken for special calibration runs, accidental cosmic ray vetoes, dead-time losses, experiment downtime, etc., is assumed to be 50% of the total running time. Backgrounds, which are described in detail in Chapter 3, are expected to be on the order of 0.4 events for a two-year run. Hence, if $R_{\mu e} = 1 \times 10^{-16}$, then the signal-to-background will be about ten.

2.2 Experimental Principles

In order to operate at the proposed rates to achieve a 10^4 improvement in sensitivity in $R_{\mu e}$ several significant innovations over the experimental approaches of the previous SINDRUM II [47] and TRIUMF [23] experiments are needed. We highlight three of them.

The first is a pulsed beam, which was used before in an early PSI experiment [49]. At the muon intensities required to achieve the planned *Mu2e* sensitivity, there is a large flux of particles other than muons that would overwhelm any attempt to perform pattern recognition. With a pulsed beam, we simply wait about 700 ns until these “prompt” backgrounds die away. The measurement period then begins and continues until the next muon injection. The pulse spacing at Fermilab is $1.7 \mu\text{sec}$, well matched to the muon lifetime of 864 ns in aluminum and allows a 700 ns detector-off time after the beam flash, followed by a measurement period of 900 ns.

Overview of the Experiment

Table 2.1: Mu2e experimental parameters. A two-year run is assumed with 2.0×10^7 s per year of beam delivered to the experiment, of which half is used for normal data taking.

Beam kinetic energy	8.0 GeV
Running time per year	2.0×10^7 s/yr
Total running time	2 years
Protons per year	3.6×10^{20}
Average proton beam power	23.1 kW
Instantaneous proton beam power	25.6 kW
Booster beam	4.0×10^{12} protons/batch
Booster batches	6
Total protons per spill	2.4×10^{13}
Normal data taking fraction	0.50
Spill period	1.33 s
Spill length	1.20 s
Microscopic duty factor	0.90
Microbunch period	1,700 ns
Microbunch FWHM	100 ns
Live gate length	900 ns
Total live time per year	4.77×10^6 s
Protons per second	1.80×10^{13} s ⁻¹
Protons per spill second	2.00×10^{13} s ⁻¹
Protons per microbunch	3.40×10^7
Probability per proton of μ^- stopping	0.0025
μ^- stops per spill second	5.00×10^{10} s ⁻¹
μ^- stops per microbunch	8.48×10^4
μ^- stops per year	4.50×10^{17}
Capture rate/decay rate (Al)	0.61
μ^- captures per spill second	3.04×10^{10} s ⁻¹
μ^- captures per microbunch	5.17×10^{13}
μ^- captures per year	2.74×10^{17}
Muon captures in time window	0.51
Reconstruction and selection efficiency	0.193
Calorimeter trigger efficiency	0.80
Overall efficiency	7.6%

Table 2.2: Expected sensitivity for a two-year (2×10^7 s) run.

Running time (s)	2.0×10^7
Proton flux (Hz)	1.8×10^{13}
Probability of μ/p transported and stopped in target	0.0025
μ capture probability	0.61
Fraction of μ which capture in time window	0.51
Electron trigger efficiency	0.80
Fitting and selection criteria	0.19
Single-event sensitivity	2.3×10^{-17}
90% CL sensitivity	5.7×10^{-17}
Events per 1×10^{-16} BR	4.3

A second improvement is the use of curved solenoids [48, 50]. These allow us to greatly reduce prompt backgrounds, and to separate μ^- from μ^+ and other extraneous particles. The two work together to produce a 10^3 increase in the muon rate possible with the 8 GeV accelerator complex at Fermilab.

The third improvement is to embed the muon production target in a graded solenoidal field, which significantly improves the stopped muon rate [50]. Similarly, the aluminum stopping target is placed in a graded field, which substantially improves the acceptance of conversion electrons in the detector and also helps suppress cosmic ray and other backgrounds.

Much of this work rests on the experience of previous muon-to-electron conversion experiments or proposals. The solenoidal muon collection scheme was first conceived by the MELC collaboration [50] (a muon-to-electron conversion experiment proposed in Russia) and was developed further by the MECO collaboration. The experiment we propose is largely the MECO experiment moved to Fermilab to take advantage of the time-structured proton beam that can be delivered by the Booster-Debuncher-Accumulator complex. This document relies heavily on the MECO technical design report, proposal, and internal documents [48].

We now begin a highly abbreviated overview of the experiment from the Fermilab Booster through to the final beam absorber. Details are presented in individual chapters.

2.3 Proton Beam

The Fermilab accelerator complex is ideally suited for producing the *Mu2e* muon beam, since high intensity proton beams of no more than 8 GeV energy are needed (in order to limit antiproton production) with a bunch spacing on the order of the stopped muon lifetime (1–2 μ s). The availability of the Debuncher and Accumulator

Overview of the Experiment

rings allows a proton beam with the desired parameters to be produced without any impact on the beam required for the NO ν A neutrino program, as is described in detail in Chapter 6.

The Booster accelerates 4×10^{12} protons per batch to 8 GeV every 1/15 s. A total of 12 of these batches are stacked in the Recycler ring, then sent to the Main Injector where they are accelerated to 120 GeV before being extracted to the NO ν A production target. The process is repeated at the Main Injector cycle time of 1.33 s. While the NO ν A batches are being accelerated in the Main Injector, 8 Booster batches are available for other physics; *Mu2e* intends to use 6 of them. The *Mu2e* Booster batches are sent to the Debuncher ring via the Recycler ring (in what we call the “Boomerang Scheme”) where they are stacked, 3 batches at a time, and bunched into ~ 100 ns wide bunches separated by the $1.7 \mu\text{s}$ period of the Debuncher (see Fig. 2.3). The bunches are then transferred to the Accumulator ring from which they are slow extracted to the *Mu2e* experimental hall. An average of 1.8×10^{13} protons/s are delivered to the muon production target in microbunches of 3.4×10^7 protons, with a duty factor of 90%.




Figure 2.3: The proton beam bunch scheme. Every $1.7 \mu\text{s}$ a bunch of 3.4×10^7 protons is delivered to the muon production target. For 700 ns after the end of the 100 ns wide bunch the detector is off.

In order to control an important class of backgrounds, the flux of protons striking the primary target between beam pulses must be suppressed by a large factor. While these protons are headed toward the production target, a pair of high-frequency dipoles will sweep all but the desired beam out of the beam path. We show later that it is extremely important that for every proton in the beam pulse, there are fewer than 10^{-9} out-of-time protons. We refer to this suppression as the “extinction” requirement and will refer to it many times throughout this Proposal. The extinction scheme is described in Section 6.5.

2.4 Solenoid System

The superconducting solenoid is the heart of the experiment. An isometric view is provided in Fig. 2.2. This shows the essence of the technique. The solenoid is divided into three sections: Production, Transport, and Detector Solenoids. The 8 GeV protons enter the Production Solenoid and strike a target inside the solenoid, creating pions that decay into muons. The muons travel through the Transport Solenoid along an S-shaped section, and then enter the Detector Solenoid. Inside the Detector Solenoid they are brought to rest in a stopping target and $\mu N \rightarrow eN$ conversion electrons are then tracked and momentum analyzed. An electromagnetic calorimeter provides the trigger. The remaining beam from the Transport Solenoid is stopped in a downstream beam dump that also serves as a flux return for the magnet.

There are four essential features of this design:

1. The protons are targeted opposite to the direction of the outgoing muons. They enter the solenoid through a small port and exit in the opposite direction from the experiment. This protects the downstream solenoids from damage from the beam and reduces the of secondary particles entering the Transport Solenoid.
2. The muon production target is immersed in a high-gradient high-field solenoidal magnetic field with a high pion and muon collection efficiency.
3. The S-shaped Transport Solenoid serves two purposes. The bend removes any lines-of-sight neutral particles from the production target impacting the muon stopping target. Furthermore, the curved portions of the solenoid possess field gradients. As a result the first curve separates positive from negative particles in the vertical direction. The second curve re-centers particles onto the solenoid axis. Collimators are used to select the desired low momentum μ^- and suppress unwanted positive particles and high momentum negative particles.
4. The stopping target is placed in a gradient field in order to reduce cosmic ray backgrounds and rates in the detectors. The gradient reflects some upstream-going conversion electron candidates back downstream, improving the detector acceptance by a factor of two.

2.4.1 Production Solenoid

The production solenoid is shown in Fig. 2.4. Protons enter through a small port on the upstream side of the production solenoid and strike a production target; the remnant proton beam and forward-produced secondary particles exit at the upstream end of the solenoid. The inner bore of the solenoid is heavily shielded to protect the cold magnet coils from radiation damage and excessive radiation heating. The magnetic field is graded from approximately 5.0 Tesla on the upstream side down to 2.5 Tesla at the entrance to the Transport Solenoid. This graded field captures

Overview of the Experiment

pions, which spiral around in the field and decay into muons. The graded field serves to “accelerate” muons into the Transport Solenoid, and also enhances muon flux by reflecting some upstream-going particles back downstream.




Figure 2.4: Cutaway view of the Production Solenoid. The initial proton beam enters the production solenoid from the right, and muons are collected and transported back to the right into the Transport Solenoid.

2.4.2 Transport Solenoid

The Transport Solenoid, shown in Fig. 2.5, takes the negatively charged muons from the Production Solenoid and transports them to the Detector Solenoid, while eliminating many of the unwanted particles coming from the production target. Collimators at the entrance and exit of the Transport Solenoid eliminate high-momentum charged particles coming directly from the production target and from pion decays. Antiprotons, which if allowed to continue on to the muon stopping target would produce a serious background, are absorbed in a thin window (not shown) halfway down the Transport Solenoid. Neutrons and photons are absorbed in the walls of the solenoid as there is no line of sight from its entrance to exit. Positively charged particles are eliminated as described below. In the upstream curved solenoid portion, as shown in Fig. 2.5, the spiraling positive (blue) and negative (red) muons are deflected downwards and upwards respectively, by amounts which depend on their momenta. An asymmetric collimator allows passage of the low momentum negative particles, including the desired low momentum negative muons and strongly suppresses positives. The particle trajectories are re-centered on the solenoid axis by the second curved section.




Figure 2.5: Cutaway view of the Transport Solenoid in the region of the asymmetric collimator. The upper spiraling negative muons (red) pass through, the lower positive muons (blue) are stopped.

2.4.3 Detector Solenoid

The Detector Solenoid is shown in Fig. 2.6. The solenoid field decreases from 2 Tesla at the entrance down to a constant 1 Tesla in the detector region. About half of the incident muons stop in the stopping target. If a muon interaction produces an electron with sufficiently high energy — for example a conversion electron at 105 MeV — the electron is then tracked and analyzed in the octagonal tracker. The electromagnetic calorimeter confirms the momentum measurement of the tracker and provides the trigger.

A critical feature of this arrangement is that the detector surrounds the muon beam with a radius such that most of the background decay-in-orbit (DIO) electrons, which are produced at a rate of 25×10^9 per second, have a radius less than that of the inner elements of the detector. The vast majority of the DIO spiral harmlessly through the empty center of the tracker and calorimeter into the muon beam dump.

The stopping target is embedded in a graded field that extends from the end of the Transport Solenoid to the beginning of the tracker. Conversion events are emitted isotropically from the stopping target, and so half the events are traveling backwards and would be lost. The graded field, analogous to the gradient in the production solenoid, “reflects” these muons, nearly doubling the detector acceptance. Figure 2.7 shows the effect for conversion events.

Further, beam electrons above 100 MeV emerging from the Transport Solenoid

Overview of the Experiment




Figure 2.6: Cutaway view of the detector solenoid. The green curves are the trajectories of conversion electrons.




Figure 2.7: The polar angle distribution of accepted conversion electrons. Electrons emitted from the target at $\theta = 0^\circ$ are parallel with the solenoid axis. Those at $\theta > 90^\circ$ are emitted backwards and are reflected in the graded field of the Detector Solenoid.

have their pitch increased so much by the field gradient that they can be distinguished from > 100 MeV electrons from the target. High energy electrons from cosmic ray decays produced in the Transport Solenoid can be distinguished in the same way.

The detectors are displaced downstream of the stopping target where the field is uniform. The distance reduces the backgrounds from neutrons and photons emanating from the stopping target. Thin absorbers are placed around the target and between the target and the detectors to absorb a portion of low energy protons from muon captures which would otherwise have enough momentum to intercept the detectors.

2.5 Stopping Target

The Transport Solenoid directs negatively charged muons to the stopping target consisting of seventeen 0.02 cm Al foils. Aluminum has been selected after considering trade-offs between the stopping rate, backgrounds, lifetime, and conversion electron energy degradation due to energy straggling in the target. Titanium is also a viable candidate. About half of the muons stop in the target, 0.0025 per proton incident on the muon production target.

After the muon is stopped it will either decay-in-orbit into an electron and two neutrinos, be captured by the nucleus, or undergo the sought for conversion process. In muonic aluminum, approximately 60% of the muons undergo ordinary capture on the nucleus, which is essentially an interaction with a proton in the nucleus, $\mu^- + p \rightarrow n + \nu_\mu$. The remaining 40% of the muons decay in orbit. Muons that are captured serve as the denominator in $R_{\mu e}$ since the wavefunction overlaps in the conversion and capture processes can cancel. The lifetime of a muon bound in atomic orbit in aluminum is 864 ns, well-suited to the time structure of the Fermilab proton beam and the muon decay lifetime [51].

2.6 Detector

A straw-tube tracking detector downstream of the stopping target measures the momentum of the electrons. The magnetic field is a highly uniform 1.00 Tesla. There is no material inside a cylindrical volume of radius 38 cm, and the Detector Solenoid is evacuated, allowing the large flux of low transverse momentum ($p_\perp < 55$ MeV/c) charged particles to spiral into the downstream muon beam stop. These particles are mainly: (1) DIO electrons from the stopping target, (2) protons, neutrons, and gammas from the muon capture process, (3) beam electrons, or (4) muons which do not stop in the target. The muon beam stop is carefully designed to suppress albedo of secondary particles back to the detectors.

The tracking detector is made from straw tubes nearly oriented along the axis of the solenoid in an octagonal geometry, with 8 vanes extending radially outward as

Overview of the Experiment

shown in Fig. 2.6. It has a geometric acceptance of about 50% for conversion electrons [48].

The tracker momentum resolution for 100 MeV electrons is about 120 keV and is dominated by multiple scattering in the tracker. The overall conversion electron energy resolution, 900 keV FWHM, is asymmetric; the low side has a long tail which is dominated by the stochastic energy loss in the stopping target and proton absorbers, and the high side is closer to gaussian in shape and has contributions from multiple scattering in the tracker as well as energy losses.

The tracking detector is followed by the electromagnetic calorimeter. The baseline calorimeter consists of four vanes of arrays of lead-tungstate crystals. It provides a measurement of the conversion electron energy of about $\sigma(E) = 5 - 8\%$, and a ~ 1 cm position resolution, which assists in background suppression. The calorimeter is used as a trigger for the experiment, which then initiates the data transfer from the tracker and other elements of the experiment to the data acquisition system.

2.7 Backgrounds

Much of the design of the experiment is driven by the need to control backgrounds in the detectors. We discuss the most important background sources below. They are described in more detail in the following chapter.

2.7.1 Decay-in-Orbit Electrons

The 40% of the muons that decay in orbit (DIO):

$$\mu^- + A(Z, N) \rightarrow A(Z, N) + e^- + \bar{\nu}_e + \nu_\mu,$$

contribute about half the total background in the signal energy region. The bulk of electrons from muon decay-in-orbit have energies below $m_\mu/2 \approx 53$ MeV, which is the kinematic limit for the decay of a free muon. However, in the decay of bound muons, the outgoing electron can recoil off the nucleus, and in the limit that the neutrinos have zero energy the electron will have exactly the conversion energy of $E_{\text{CE}} = 104.97$ MeV. The rate is quite small and falls very rapidly as $(E_{\text{CE}} - E_{\text{electron}})^5$ near the endpoint energy. Only $\sim 3 \times 10^{-13}$ of the decay electrons are > 100 MeV. The expected resolution on the tracker momentum for decay-in-orbit events smearing into the conversion signal region is approximately 110 keV (see Chapter 3). This is combined with energy straggling to give about 250 keV resolution on the high side of the energy distribution, which is sufficient to keep the DIO background level at 0.22 counts in the signal region for a run of 2×10^7 s.

It is possible for background tracker hits to be included in a lower energy DIO electron track, creating a false conversion electron near 105 MeV. Extensive studies by MECO show that the resultant degradation of the tracker resolution is small at

rates of up to 500 kHz in the straw tubes, the highest expected in MECO. *Mu2e* plans to operate at about one-third the MECO rates because of an improved duty factor (90% vs 50% for MECO) and a longer running time. Hence these backgrounds are expected to be even smaller.

2.7.2 Nuclear Capture of Muons

The muon can undergo capture on the nucleus:

$$\mu^- + A(Z, N) \rightarrow \nu_\mu + A'(Z', N') + an + bp + c\gamma.$$

This occurs about 60% of the time for muonic aluminum, and on average, $a \approx 2$, $b \approx 0.1$, $c \approx 2$. Most of the secondary protons, neutrons and gammas have low energies. They can produce high rates of background in the detectors, with the potential to cause errors in reconstructed tracks, which could create false conversion electrons by boosting low energy DIO electrons up to the conversion electron energy. Thin absorbers placed between the stopping target and the tracker remove many of the low energy protons, and the detectors are displaced downstream from the stopping target to help reduce neutron and photon background.

The muon can undergo radiative capture in the aluminum target:

$$\mu^- + \text{Al}(13, 27) \rightarrow \text{Mg}(12, 27) + \nu_\mu + \gamma$$

with a branching ratio of a few $\times 10^{-5}$, where the photon in the above reaction is typically above about 55 MeV up to the maximum kinematic limit. This kinematic limit depends on the rest energies of the nuclear fragments. If the nuclear fragments have a total rest mass which is equal to or less than the stopping target nucleus, then the radiative photon can have an energy equal to or greater than the conversion electron. The photon can undergo e^+e^- pair production in or nearby the stopping target. If the electron is close to 105 MeV, it can be mistaken for a conversion electron. This background can be reduced with proper selection of the stopping target isotope, namely, the daughter nucleus should have a rest energy above that of the parent nucleus. In the case of an aluminum stopping target the radiative capture photon kinematic endpoint is 2.6 MeV below the conversion electron energy.

2.7.3 Radiative pion capture

Those negative pions surviving decay in the beam line can stop in the aluminum target, forming pionic atoms. The pion, like the muon, is captured into atomic orbit. As it cascades to lower orbits, the overlap between the nuclear and pion wavefunctions increases, and the pion rapidly undergoes a strong interaction with the nucleus. The strong interaction is so large that the pion does not make it to the 1S state — it interacts mid-cascade. About 2% of the time, a high energy photon with an average

Overview of the Experiment

energy of about 110 MeV and a maximum energy up to ≈ 140 MeV is produced, in reactions such as

$$\pi^- + A(Z, N) \rightarrow A'(Z - 1, N + 1) + \gamma$$

The photon can undergo pair production in the target or other nearby material, with a small probability of producing an electron near 105 MeV, thus faking a conversion electron. This is an example of a prompt background, since the background is produced in time coincidence with an incoming pion. It is potentially the most dangerous beam-related background. In the SINDRUM II experiment, remnant pion background was eliminated by rejecting electron events having a time correlation with incoming beam particles. As explained in more detail below, the *Mu2e* experiment plans to use a pulsed muon beam. The beginning of the measurement period is delayed by 700 ns after the proton pulse — after that period the pions have decayed and other prompt backgrounds produced by the proton pulse have dissipated sufficiently. We note that the proton intensity must be suppressed by a factor of $\sim 10^9$ between pulses in order to avoid stopped pions and other beam related background arriving at the stopping target during the measurement period.

2.7.4 High energy beam electrons, (~ 100 MeV)

High energy beam electrons can scatter from the stopping target. If they arrive during the measurement period they can appear to be a conversion electron. A high factor of between-pulse proton beam extinction must be sustained and transmission of these electrons to the stopping target during the measurement period are strongly suppressed by the design of the beamline and collimation.

2.7.5 Late Arriving Particles

A beam muon that has sufficiently large momentum (>77 MeV/c) can decay into a ≈ 105 MeV electron. If this occurs near the stopping target during the measurement period, the electron could be mistaken for a conversion electron. Muons can arrive at the stopping target during the measurement period if they move slowly through the beam line, or if they are produced by between-pulse protons striking the primary target. The latter are suppressed by collimators, particularly the slowly moving ones.

Muons and other beam particles can traverse the beam line slowly, arriving at the stopping target during the measurement period. Long transit times occur when the particle has a small helical pitch angle, or for heavy particles like antiprotons. The muon beam line must be designed to minimize or eliminate slow-moving particles, especially pions and antiprotons. The solenoids are designed with negative field gradients which tend to increase the pitch of particles traversing the solenoids, thus reducing the flux of late arriving particles. The length of the beam line eliminates, by decay, pions which have small pitch angles. Particles with very small pitches

are deflected large vertical distances in the curved solenoids and are eliminated with collimators.

Antiproton production is small at 8 GeV proton energies. Those that are transmitted in the Transport Solenoid have very low energies and can be stopped by a thin window placed half way down the Transport Solenoid line. This window has little effect on the flux of muons.

2.7.6 Cosmic Rays

Electrons of ≈ 105 MeV can be produced near the target by cosmic rays through several processes, among them being an electron scattered from the target by a traversing cosmic ray muon, decay by in-flight cosmic ray muons, cosmic ray electrons directly entering the detector region, and cosmic ray photons which can produce energetic electrons by interaction with the target or detector. These backgrounds are eliminated with passive shielding and an active veto counter surrounding the detector region. Indeed, it is planned to site the experiment underground in order to increase the passive shielding. Further suppression can be obtained by requiring hypothetical conversion electrons to originate in the stopping target, and for their helical pitches to fall in the narrow range of values expected for electrons born in the target.

2.8 Phase II: *Mu2e* at Project X Intensities

The *Mu2e* Stage II goal is to reach a sensitivity of up to $R_{\mu e} = 10^{-18}$. The ultimate sensitivity will depend on the future proton flux and upgrades to the muon beam line and detectors. It is estimated that 200 kW of proton beam power, or about 10 times more beam, will be available for *Mu2e* in the Project X upgrade. There also may be additional improvements in the number of stopped muons per proton, based on ongoing developments in muon collider research.

Upgrades to the Stage I version of the *Mu2e* experiment will be studied. At the much increased proton intensity in Stage II, the shielding of the production solenoid may have to be increased and the target improved to handle the higher heat and radiation load. New muon sources will be considered, such as helical cooling channels [52].

At the much increased muon stopping intensity, the background rate of protons, neutrons and gamma rays emitted from the stopping target, due to capture of stopped muons, will be much larger than in Stage I. Some experiment modifications may be necessary to handle it, and a range of detector upgrades will be considered. The proton background from the stopping target could be reduced by placing a series of absorbers between the stopping target and the tracker. They would have a geometry which exploits the opposite helical rotation of protons compared to conversion electrons by preferentially absorbing the protons. The neutral backgrounds can be reduced by increasing the distance between the stopping target and the tracker [48].

Overview of the Experiment

It may be possible to suppress backgrounds by installing a “C” [53] or “S” shaped solenoid channel between the stopping target and the tracker

At the same time, better electron energy resolution may be required to further suppress the background from high energy DIO electrons. Energy straggling in the stopping target and proton absorbers are major contributors to the degradation of energy resolution. The target can be made thinner at the expense of the rate of stopped muons. The stopping rate can be restored if the incident muon momentum range can be made narrower, for example by using helical cooling of the muon beam. The proton absorbers could be eliminated if a curved detector solenoid were installed.

The upgrade path for the detector in Stage II will rely heavily on lessons learned in Stage I. If a small conversion signal at $R_{\mu e} \sim 10^{-16}$ were detected in Stage I, Stage II would confirm the signal and measure the branching ratio with greater precision. In this case, the experiment could have the unique capability of being able to distinguish the underlying effective operators responsible for CLFV by observing the conversion rate on high- Z nuclear targets (e.g., gold) [54]. With the presence of a signal in Stage I, higher backgrounds could be tolerated and it may not be necessary to reduce the DIO background through improvements to the detector energy resolution. We cannot provide definitive answers at this time but both as the experiment develops and we determine its result we will develop and refine our upgrade paths.

Sensitivity and Backgrounds

3.1 Introduction

AN experiment searching for a branching ratio of $\sim 10^{-17}$ requires an exquisite understanding of resolution and backgrounds. We will show for our run of 3.6×10^{20} protons-on-target there are roughly 0.5 background events, 0.25 from decay-in-orbit and 0.25 from other prompt sources. Improving one of either the resolution or reducing the prompt background will not significantly improve the experiment. Key features of the experimental design are summarized below.

1. Since the proton beam is directed away from the detector, extra particles and hadronic debris that could cause pattern recognition errors are minimized.
2. The S-shaped solenoid prevents there from being a line-of-sight from the production region to the detector. The transport solenoid also adds a radially-graded field which deflects particles in the vertical direction according to their charges and momenta, providing sign-separation and removing positives (pions and protons) that could contribute to unwanted processes and extra hits.
3. A key design feature of the detector is that it *surrounds* the muon beam, greatly suppressing the acceptance for decay-in-orbit electrons. Only about 10^{-3} of the decay-in-orbit electrons in the solenoidal magnetic field make it to even the innermost elements of the straw tracker.

Therefore only “pathological” DIO events which combine lower momentum electrons with extra hits causing the pattern recognition to fail contribute to the DIO background. We intend to measure the resolution and its tails as described in the last section of this chapter. Of course the actual DIO spectrum extends out to the endpoint in the case that $\mu \rightarrow e\bar{\nu}_e\nu_\mu$ recoiling off the nucleus yields neutrinos whose summed momentum is less than a couple of MeV. We can calculate this background. The background itself does have some theoretical uncertainties but a modern calculation can be done with the uncertainty having no effective impact [55]. This chapter will examine the additional prompt sources of backgrounds as well.

3.1.1 Underlying Processes and Calculations

We first describe the processes that occur after the muon comes to rest in the stopping target. This will set the context for our discussion of the decay-in-orbit background and sensitivity calculation. This subsection is a simple introduction to orient the reader.

A stopping muon is captured by an Al nucleus and within picoseconds falls into the 1s state. The muon has a Bohr radius of ≈ 20 fm compared to the ~ 4 fm size of the Al nucleus. Hence there is a significant overlap of the muonic and nuclear wavefunctions. We can calculate the fraction of muons that are captured and those that decay-in-orbit. We know the total disappearance rate over all processes for muons on Al is 864 ns [51]. Given that the free muon lifetime is $2.197 \mu\text{s}$ we can trivially calculate:

$$\begin{aligned}\frac{dN_{\text{DIO}}}{dN_{\text{total}}} &= \frac{-\lambda_{\text{DIO}} N dt}{-\lambda_{\text{DIO}} N dt + \lambda_{\text{capture}} N dt} \\ &= -\frac{\lambda_{\text{DIO}}}{\lambda_{\text{DIO}} + \lambda_{\text{capture}}}\end{aligned}$$

and then

$$\text{DIO Fraction} = \frac{dN_{\text{DIO}}}{dN_{\text{total}}} = \frac{1/2.1}{1/0.864} = 0.393$$

for decay-in-orbits and $1 - 0.393 = 0.606$ for the captures.

Reference [51] calculated the *total* disappearance rate, which includes both capture and decay-in-orbit. Therefore the pure capture lifetime is given by (units of $1/\mu\text{sec}$):

$$\begin{aligned}\lambda_{\text{capture}} &= \lambda_{\text{total}} - \lambda_{\text{DIO}} \\ &= \frac{1}{0.864} - \frac{1}{2.197} \\ &= 0.70224\end{aligned}$$

and then

$$\begin{aligned}\tau_{\text{capture}} &= \frac{1}{\lambda_{\text{capture}}} \\ &= 1.424 \mu\text{sec}\end{aligned}$$

There are several other factors to be calculated. What fraction of captured muons are captured in the live gate? The muon arrival at the stopping target, relative to the proton arrival time at the production target, has a time distribution shown in Fig. 3.1. We start the gate 700 ns after the arrival of the protons at the production

target, well after the bulk of arrivals. However, now that we know the muons have a $1.4 \mu\text{s}$ lifetime we convolute that with the time structure of the pulsed beam as well as the distribution of muon stopping times. We estimate this fraction for the scenario in which the pulse spacing is 1700 ns, the time required for one revolution around the Debuncher; the gate is taken to start at 700 ns and to end at 1600 ns, just before the next pulse. The calculation yields 51%.




Figure 3.1: The distribution in the μ^- stop time relative to the time when the proton beam strikes the production target.

3.2 Physics Background Sources

In this section we give the results of calculations of the expected background levels based on the performance of the baseline beam and detector. Details of the beam and detector and of the method by which the background calculations were done are in many cases deferred to later sections. The primary sources of physics backgrounds are:

1. Muon decay in a Coulomb bound state (DIO).
2. Radiative muon capture on a nucleus (RMC).
3. Beam electrons that scatter in the stopping target.
4. Muon decay in flight.
5. Pion decay in flight.

6. Radiative pion capture on a nucleus.
7. Backgrounds induced by anti-protons.
8. Backgrounds induced by other late arriving particles.
9. Cosmic-ray induced electrons.

The decay-in-orbit (DIO) and radiative muon capture (RMC) backgrounds are intrinsic to muons stopped in the target and therefore cannot be eliminated. However, with proper detector design, they can be minimized. The endpoint of the spectrum of DIO electrons is equal to the energy of conversion electrons when the neutrinos have zero kinetic energy; energy and momentum are conserved by nuclear recoil. The spectrum falls very steeply near the endpoint conversion energy E_{CE} , $\propto (E_{CE} - E)^5$, and the number of DIO electrons that have an energy consistent with that of conversion electrons can be made negligible by a sufficiently precise measurement of the energy. The situation with RMC is similar; the spectrum of RMC electrons has an endpoint well below the signal and is also eliminated with a precise energy measurement.

Most of the remaining backgrounds derive from prompt processes, with the electron detected close in time to the arrival of a secondary beam particle at the stopping target. The pulsed beam of *Mu2e*, a critical improvement over SINDRUM, is necessary to eliminate this background. The idea is that conversion electrons are detected only during periods when the flux of particles in the secondary muon beam is extremely low; a similar technique was used in an earlier conversion experiment [49]. A first estimate of the degree to which secondary beam particles must be suppressed during the measurement time can be obtained by reference to the SINDRUM II data. Without the beam veto there is one prompt background at a sensitivity of about 10^{-10} . To get an expected background below 0.01 events at a branching fraction of 10^{-16} requires a beam extinction of 10^{-9} , which is a factor of 10^8 greater extinction than in the SINDRUM II experiment. A detailed simulation of prompt background processes leads us to conclude that a beam extinction (defined as the fraction of the total proton flux that hits the production target during the detection time) of $\leq 10^{-9}$ is required.

Even with a pulsed beam, backgrounds may result from particles that spend a very long time in the muon transport system and hence arrive at the stopping target late, even though they are produced by protons in the main pulse. These backgrounds have been calculated with a variety of Monte Carlo and semi-analytic techniques. The flux of slow moving particles is greatly reduced by the gradient fields in the straight sections of the solenoids, which tend to increase the helical pitches of particles traveling downstream, and by their large vertical deflection in the curved sections of the solenoids.

Background from antiprotons is a special case in that the background is from prompt processes for which the ultimate source is not reduced by a pulsed beam.

It results from very slow \bar{p} 's that have a transit time in the muon beam line long compared to the time between pulses. The \bar{p} 's eventually stop in material, annihilate, and the resulting electrons and pions produce background. Hence, antiprotons must not reach the stopping target and transport of their annihilation products there must be suppressed.

Finally, cosmic ray background arises from a variety of processes. This background is unique in that it scales with running time, not sensitivity. Only modest improvement in the cosmic ray background rejection with respect to that of earlier experiments is required, appropriate to the somewhat longer running time proposed.

3.2.1 Electrons from Muon Decay in Orbit

The rate for production of electrons from μ decay in Coulomb bound orbit is approximately proportional to $(E_{\text{CE}} - E_e)^5$ near the endpoint, where E_{CE} is the conversion, and the endpoint, energy. Both the absolute normalization and energy spectrum have been calculated [56] in various calculations. The approximately 25% discrepancy in the calculations can be traced to different approximations for nuclear recoil effects and relativistic corrections to the muon wave function. As discussed earlier, a new calculation is underway [55].

Because the spectrum rises very steeply from the endpoint, the signal to background ratio is extremely sensitive to electron energy resolution. To reduce the background, the width of the resolution function must be minimized and any high energy tails must be highly suppressed. Requirements on the central part of the resolution function and the high energy tails are discussed in Chapter 10.3. In principle, one needs to convolve the DIO energy spectrum with the energy resolution function to determine if the backgrounds level is acceptable. Qualitatively, the FWHM of the resolution function should be less than ~ 1 MeV, and the level of high energy tails more than ΔE above the electron's true energy should be less than $\sim 0.2 \times (\Delta E)^6$, with E in MeV. For example, the high energy tails above 5 MeV should be less than $\sim 10^{-5}$.

The main contributions to the central part of the resolution function are energy loss and straggling in the stopping target and multiple scattering in the magnetic spectrometer. Smaller contributions arise from energy loss in the spectrometer and spatial resolution in the spectrometer detectors. The largest potential contribution to high energy tails is pattern recognition errors (associating noise hits in the tracking detector with a putative particle trajectory) in the noisy environment around the stopping target and spectrometer. Straggling (large energy loss) in the stopping target and detector reduces acceptance when energy loss is sufficiently large that the event cannot be distinguished from background (typically more than 1.4 MeV loss), but does not contribute to background since DIO events are smeared *downward* to lower energy away from the signal region. Hence dE/dx is a more of a problem because of loss of acceptance rather than from smearing. Nonetheless in order to

predict the observed spectrum we must understand it as well and special runs are being investigated.

A simulation of all processes contributing to energy mismeasurement and an analysis that includes selection criteria designed to remove badly measured events was done and is discussed in Chapter 10.3 and Ref. [57]. The most probable detected energy is less than 105 MeV due to energy loss in the target, the proton absorbers, and the tracking detector. By accepting events between 103.6 MeV and 105.1 MeV, the noise to signal ratio is below 0.05 for $R_{\mu e} = 10^{-16}$ with selection criteria that give $\sim 20\%$ acceptance. The resolution function has FWHM of about 900 keV and very little high energy tail. There are additional potential backgrounds resulting from pattern recognition errors; these are discussed in detail in Section 10.3.3. The background level depends sensitively on the detector accidental rates. It is calculated to be < 0.002 events at the expected detector rates.

Tails in the resolution function arising from pattern recognition errors can be studied by relaxing the selection criteria. For example, the requirement that the fitted trajectory project to the energy deposited in the trigger calorimeter can be removed and the number and energy distribution of background excluded by that cut compared to simulations. Other examples of essentially independent selection criteria that can be relaxed are the requirement that there be no missing hits on the fitted trajectory and the requirement that no low energy track be reconstructed in the same event. In this way, the efficacy of each selection criterion can be measured from the data and compared with the calculated efficacy. Finally, tails in the resolution function are extremely sensitive to detector rates, and the background rate with relaxed selection criteria will be studied as a function of the muon stopping rate. We have begun to examine an ensemble of resolution-measuring runs, described in Sec. 3.5.

We conclude with a number of plots showing the behavior of the DIO background. The final expected distribution of signal and DIO background for a branching ratio of 10^{-16} is given in Fig. 3.2. This plot shows the momentum reconstructed at the tracker.

We will demonstrate in this chapter that we expect roughly 0.25 decay-in-orbit and 0.25 “prompt” event backgrounds. Using the expected distributions, we can plot the integrated signal and the integrated signal divided by the square root of the background. This quantity S/\sqrt{B} is a proxy for the number of standard deviations above background and is a quick “figure-of-merit” (FOM) we can use for estimating where to cut on observed momentum to optimize the sensitivity. If we cut at a reconstructed momentum at the tracker at 103 MeV/c we obtain (for a BR = 10^{-16}), 4 signal events to 0.4 background, with a FOM of 5.5 σ , which is above the 5 σ “discovery” criterion.

3.2.2 Radiative μ Capture

Electrons result from radiative muon capture (RMC), $\mu^- \text{Al} \rightarrow \gamma \nu_\mu \text{Mg}$. The γ end-point energy is 102.5 MeV and the probability per μ capture of producing a photon

Sensitivity and Backgrounds




Figure 3.2: The final distribution of decay-in-orbit and signal events for a BR of 10^{-16} . The x -axis is the reconstructed momentum at the tracker.




Figure 3.3: The integrated number of signal events above the indicated momentum at the tracker.




Figure 3.4: The integrated number of background events (both DIO and prompt) above the indicated momentum at the tracker for a BR of 10^{-16} .




Figure 3.5: The quantity $FOM = S/\sqrt{B}$ vs. lower cut on momentum, from the previous two plots, for a BR of 10^{-16} . We see for a cut at 103.25 MeV the FOM is maximized at 5.5.

Sensitivity and Backgrounds

with energy exceeding 100.5 MeV is $\sim 4 \times 10^{-9}$ [58]. The pair production probability in the target is ~ 0.005 , and the probability that the electron energy exceeds 100 MeV is ~ 0.005 . Thus, the probability of producing an electron above 100 MeV is $\sim 10^{-13}$.

These electrons are all less than 102 MeV and for an electron to be considered signal, its measured energy must exceed ≈ 103.25 MeV as we saw above. The integral of the high energy tail in the resolution function above 1.6 MeV is $< 10^{-6}$. Hence, the probability of getting an electron above 103.25 MeV from radiative μ capture is $< 10^{-19}$ or a background to signal ratio of < 0.001 for $R_{\mu e} = 10^{-16}$.

This background is not distinguished from DIO electrons. However, the measured energy distribution near the endpoint can be fit to a combination of DIO and RMC electrons to verify the respective contributions.

3.2.3 Beam Electrons

Beam electrons may cause background if they are produced in the production or transport solenoid region and then scatter in the stopping target. Independent of the transverse momentum of the electron as it exits the transport solenoid, the transverse momentum at the tracking detector is below 75 MeV/c unless it scatters in the stopping target, by design of the detector solenoid field.

The rate for electrons scattering at ~ 100 MeV is defined by the Mott cross section multiplied by a nuclear form factor for the target material. The experimentally determined [59] form factor for aluminum is shown in Fig. 3.6. Figure 3.6 shows the scattering cross section on aluminum for Mott scattering with and without the form factor included.




Figure 3.6: The form factor for electrons scattering on aluminum is shown on the left. The plot on the right shows the electron scattering cross section as a function of scattering angle for 100 MeV electrons. Cross sections are shown for the Mott formula, and with the inclusion of the nuclear form factor.

The collimator system is designed to suppress high energy electrons. A GEANT simulation of the production of electrons and their transport to the detector solenoid yielded no transmitted electrons above 100 MeV for 10^7 incident protons. We approximate the expected number of events above 100 MeV by fitting the energy spectrum between 70 and 90 MeV to an exponential and extrapolating to the region above 100 MeV. We take the transverse momentum distribution of the events to be that of electrons with energy in the interval 70–90 MeV (essentially flat from 0–20 MeV/c), and use that distribution to calculate the probability of scattering in the target to a transverse momentum exceeding 90 MeV/c. Including the solid angle acceptance, the probability that electrons of 100 MeV scatter to $p_t > 90$ MeV/c is about 10^{-5} . With a run time of 10^7 s, a proton intensity of 4×10^{13} p/s, and a beam extinction of 10^{-9} , the expected background is 0.04 events in a 1.5 MeV region around the signal.

Background from beam electrons can be distinguished from other prompt backgrounds by the energy distribution of these events, which is falling very steeply with energy due to the beam acceptance, and by their transverse momentum distribution, which is also strongly forward peaked. This is contrasted with the transverse momentum distribution of electrons from muon DIO and radiative pion capture. The energy distribution is also very different from that of electrons from radiative pion capture. These events will also have a very different time distribution than those from muon DIO.

3.2.4 Muon Decay in Flight

Muons decaying in flight can produce energetic electrons that either have sufficient transverse momentum to fake signal or that scatter in the stopping target and fake signal. In order for the electron to have energy above 102 MeV, the μ^- momentum must exceed 77 MeV/c. Electrons produced by μ^- decays before and within the transport solenoid are included in the beam electron background calculation. Background from decays in the detector solenoid are calculated using a GEANT beam simulation. The yield of muons with $p_\mu > 77$ MeV/c passing the transport solenoid is $\sim 0.5 \times 10^{-4}$ per incident proton, the decay probability upstream of the tracking detector is 1.2×10^{-2} , and the probability of having $103 \text{ MeV} < E_e < 105 \text{ MeV}$ and $p_t > 90 \text{ MeV/c}$ is less than 1.2×10^{-7} with no events seen in a larger energy interval. With a beam extinction of 10^{-9} , the total background in a one year run is less than 0.03. This estimate can be improved with increased statistics in the simulation.

A second background source is electrons from μ^- decay that scatter in the stopping target. A GEANT simulation was used to calculate the energy distribution of electrons from muons that passed the transport solenoid and decayed in the region between the entrance to the production solenoid and the tracking detector. The kinematic distribution of these electrons was then used as input to a simple simulation of the probability that an electron of the appropriate energy scattered in the stopping target with a resulting transverse momentum exceeding 90 MeV/c. The probability

per proton of getting an electron with $103 \text{ MeV} < E_e < 105 \text{ MeV}$ from a μ^- decay is 0.5×10^{-8} and the probability of scattering to $p_t > 90 \text{ MeV}/c$ is 2×10^{-5} , resulting in an expected background of 0.04 events for an extinction of 10^{-9} .

This background is very similar in kinematic and time distributions to that from beam electrons and can be distinguished from muon DIO and radiative pion decay in the same way.

3.2.5 Pion Decay in Flight

Beam pions decaying to electrons with $E_e > 102 \text{ MeV}$ and $p_t > 90 \text{ MeV}/c$ are also a potential source of background. The π momentum must exceed $60 \text{ MeV}/c$ for this background process. A GEANT simulation was used to calculate the probability of a proton producing a beam π with $p_\pi > 54 \text{ MeV}/c$ passing the transport solenoid; it is 2.0×10^{-6} . The probability for a π to decay into an electron after the transport solenoid and before the tracking detector is 1×10^{-4} and the probability of the decay electron to have $E_e > 102 \text{ MeV}$ and $p_t > 90 \text{ MeV}/c$ is 5×10^{-6} . The background from this source is $< 10^{-3}$ events for an extinction of 10^{-9} .

A second background mechanism is π decay electrons that scatter in the stopping target. This background was calculated in much the same way as the similar process for μ^- decay. The number of electrons from π decay with $103 \text{ MeV} < E_e < 105 \text{ MeV}$ per proton is 0.8×10^{-11} and the probability of scattering to $p_t > 90 \text{ MeV}/c$ is 4×10^{-5} , also resulting in an expected background of less than 10^{-3} events for an extinction of 10^{-9} .

3.2.6 Radiative π Capture

Pions are immediately captured by a nucleus after they stop in the target; about 2% of the captures result in the emission of a photon [60] without significant nuclear excitation. The photon energy spectrum has a peak at 110 MeV and endpoint at 140 MeV. The probability of photon conversion in the Al target, with a conversion electron in a 1.5 MeV energy interval around 104 MeV is 3.5×10^{-5} , as calculated in a GEANT simulation. The acceptance for electrons from photon conversion is large, ~ 0.8 , since the path length for conversion is largest for photons emitted at 90° . The yield of π 's that pass the transport solenoid and stop in the target is $\sim 3 \times 10^{-7}$ per proton. With a beam extinction of 10^{-9} , the background is estimated to be 0.06 events from pions produced by protons impinging on the target between beam pulses.

The energy spectrum of these events extends to above 130 MeV. They are distinguished from beam electrons and muon decay in flight backgrounds by this spectrum. They also are more strongly peaked to higher transverse momentum and can be distinguished in this way. As with all prompt processes, the fluxes of these can be increased by orders of magnitude by decreasing the extinction.

A second contribution to radiative pion capture is that due to pions that take a very long time to traverse the production and transport solenoid and arrive at the stopping target. For these events, the suppression factor from the beam extinction is absent. However, since our detection window starts 700 ns after the proton pulse, the pions must live approximately that long and must either move slowly or follow a trajectory in the transport solenoid that results in a flight time of 700 ns in order to be a source of background. This background is estimated as follows. Protons are caused to impinge on the production target. The momentum, position, and time coordinates are recorded for pions that reach the entrance of the transport solenoid. These events are then transported to the stopping target without allowing them to decay. Figure 3.7 shows the distribution in the arrival time at the stopping target,




Figure 3.7: Distribution in the π arrival time for late arriving pions, weighted by the survival probability in the transport solenoid.

weighted by the survival probability. Based on the time distribution in Fig. 3.7, we take an accepted time window starting 700 ns after the proton pulse (the minimum flight time to the target is ~ 50 ns, giving a ratio of late arriving pions per proton of 0.4×10^{-17} . The probability of making a background electron is 5.6×10^{-7} , as in the preceding paragraph. Hence, the expected background is 0.001 events.

This background is easily calibrated from the data due to the very strong time dependence. By measuring the number of energetic electrons as a function of time during the pulse, this source can be directly normalized and an appropriate starting time for the detection window chosen.

3.2.7 Antiproton Induced

Another potential source of background is due to anti-protons. Only low momentum, < 100 MeV/c, antiprotons can propagate down the transport solenoid; they have

very low kinetic energy and velocity and will take a very long time to transit the transport solenoid. For this reason, they are not suppressed by the beam extinction and arrive at the stopping target essentially continuously. Since the proton energy is near \bar{p} production threshold, the production cross section is small and strongly depends on the Fermi momentum in the nucleus, particularly for low momentum \bar{p} 's. This cross section is not well measured, nor is it well modeled in the GEANT code.

In general, the \bar{p} -induced backgrounds can be suppressed in the following ways:

- Reduce the rate of \bar{p} production by decreasing the proton beam momentum below production threshold.
- Absorption in a thin absorber somewhere in the muon transport.
- Sweeping antiprotons from the beam with crossed E and B fields.

We have studied the backgrounds arising from antiprotons produced at the production target for proton beam momenta in the interval 5–8 GeV/c. The studies included the appropriate nuclear model for nucleon momentum and energy distributions, the \bar{p} production mechanisms, and the \bar{p} annihilation mechanisms. The \bar{p} flux was calculated, and from that the flux of pions and electrons from annihilation. It was found that a thin absorber in the transport system is sufficient to reduce the background to below 10^{-3} events. This is accomplished without introducing additional sources of backgrounds or changing the muon yield. The calculations are discussed in Ref. [61].

Figure 3.8 shows the differential cross section for \bar{p} production at several angles for 7 GeV/c protons incident on W, calculated using a phase space model. The kinematics of nucleons in the heavy nucleus were modeled [62] including a dispersion relationship between energy and momentum. At low momentum, antiprotons are produced isotropically, as expected from phase space considerations. At high momentum, the production cross section is strongly forward peaked. Recall that for \bar{p} production on a proton at threshold, the momentum of each of the four baryons is one fourth the beam momentum. The \bar{p} momentum is peaked and slow-varying between 1-2 GeV/c, and the production is suppressed for momenta below 100 MeV/c. The differential cross section was normalized to the measurements of Cork, et al. [63], for incident proton kinetic energy of 6 GeV on a beryllium target and \bar{p} momentum of 1.41 GeV/c at 6 degrees exit angle, corrected for the A dependence of the production cross section [64, 65]. This is essentially at the peak of the differential cross section. Other measurements [66] (also on Be) give similar normalization. There is also differential production cross section data for 12.3 GeV protons incident on copper at \bar{p} momenta around 1 GeV/c in the forward direction [67]. As a comparison, the production cross section was also calculated using the same phase space method for 12.3 GeV incident protons. The calculated values were smaller than the measurements by a factor of 4, presumably because the phase space method under-estimates the production at higher energy where more production channels open up. Even if

the discrepancy between data and theory persisted near threshold for heavy targets, the \bar{p} induced background would not be limiting.




Figure 3.8: Antiproton differential cross section at production angles of 0, 90 and 180 degrees, respectively.

The calculated production cross sections were used to generate \bar{p} 's within the *Mu2e* production target; their motion and interactions were then studied using GEANT. It was found that most of the low energy antiprotons that entered the transport system were produced at the peak of the production cross section (forward and with ~ 1.5 GeV/c) and were shifted down in momentum by dE/dx energy loss and nuclear interactions in the target. Hence, the predicted \bar{p} flux is reliably normalized to experimentally determined production cross sections in the relevant kinematic region.

These simulations showed that an unacceptable flux of antiprotons reached the stopping target with the nominal transport. The typical kinetic energy was very low, and they could be completely absorbed with a thin window at the center of the transport solenoid. Figure 3.9 shows the \bar{p} annihilation positions in the transport with the absorber in place. The horizontal line at the position (0,0) is formed by the many annihilations in the absorber.

To calculate expected fluxes of electrons and pions in the detector solenoid from \bar{p} annihilations, experimentally determined annihilation cross sections were used to generate appropriate numbers and kinematic distributions of charged and neutral pions at the \bar{p} annihilation positions. These annihilation products were then tracked with GEANT and particle fluxes at the stopping target determined. Using previously determined probabilities that pions and electrons produce background, the expected \bar{p} -induced background was calculated. The resulting low level of background primarily resulted from radiative pion capture with a smaller contribution from electrons that scatter in the stopping target. Table 3.1 summarizes results of \bar{p} induced backgrounds for different incident proton momenta; the details of the calculations are discussed in reference [61].

Sensitivity and Backgrounds




Figure 3.9: A scatterplot of the \bar{p} annihilation positions in the transport system. Outlines of the production solenoid, the first 90° bend in the transport solenoid and the thin window that stops all antiprotons that would otherwise reach the detector solenoid are formed by the dense accumulation of annihilation positions.

Table 3.1: The \bar{p} induced backgrounds for different incident proton momenta.

p momentum (GeV/c)	Number of \bar{p}/p produced	Number of \bar{p}/p entering transport	Background events
5	3.9×10^{-10}	7.4×10^{-15}	7×10^{-7}
6	5.3×10^{-8}	8.0×10^{-13}	8×10^{-5}
7	1.4×10^{-6}	1.2×10^{-11}	1.2×10^{-3}
8	8.5×10^{-6}	6.8×10^{-11}	7×10^{-3}

Anti-proton induced backgrounds are very sensitive to the incident proton energy and in principle can be studied by changing proton beam energy. They are also the only beam associated background that has a time distribution with no correlation with the RF structure in the beam. Backgrounds without time structure above the level of that due to cosmic rays would indicate \bar{p} contamination.

3.2.8 Long Transit Time Backgrounds

We have discussed two sources of backgrounds from late arriving particles: pions that arrive at the stopping target with long transport times and anti-protons. Late arriving pions are easily calculated with high sensitivity because we can turn off decays and study them with high effective luminosity and then multiply by the very small probability that they have not decayed. Anti-proton induced backgrounds can be calculated because we can independently calculate the stopping probability at different places in the transport and the probability that annihilation product produced at the stopping position causes background. Both are small numbers, and they can be multiplied to give the total background probability.

Backgrounds from other processes with long transit times are not so easily calculated because large suppression factors cannot be independently calculated and then multiplied to give a total suppression. The largest contribution to this category is pions that decay in a region of nearly uniform field (in the first straight section of the transport solenoid, for example) yielding a trajectory with very little longitudinal momentum. Without a muon decay, such a particle would be swept out of the beam acceptance as it traversed the first curved section. However, with a muon decay and the electron subsequently scattering, it could in principle reach the detector solenoid. There is no simple set of multiplicative suppression factors that can be calculated to describe this situation.

To study those backgrounds in a complete Monte Carlo technique to the level of 10^{-3} events would require simulating about 10^{17} protons. Simulating that many events is not possible given available computing resources. We have adopted a semi-analytic technique that breaks this generic type of background into many classes (depending, e.g., on where the pion and muon decay) and then calculating either multiplicative suppression factors for these many classes or showing that a particular class cannot result in a particle surviving the transport.

A full discussion of the calculations is given in reference [68]. The potential for long transit time in both the production and the transport region is considered. In the transport region, maximum possible transit times in each of the straight and curved sections is calculated. We considered all possible decay modes of $\mu \rightarrow e$, $\pi \rightarrow e$ and $\pi \rightarrow \mu \rightarrow e$. We also take into account the effects of scattering on the collimators and the production target.

Some important features of the *Mu2e* design allows significant suppression of the late arriving particle backgrounds. First an axial B field with a negative gradient

tends to increase the longitudinal momentum of charged particles moving through it in helical orbits. As a result of these background calculations, we have removed all constant field regions of the transport, introducing gradients in each of the three straight sections of the transport solenoid and removing the constant field region in the production solenoid. Another important feature is that particles moving in helical orbits in a torus drift in a direction perpendicular to the plane of the torus. This is because the B field has a $1/r$ dependence which gives rise to a radial gradient of the field. The amount of drift is approximately proportional to the inverse of longitudinal momentum. Therefore, a charged particle with longitudinal momentum small enough to have a long transit time in the curved portion of the solenoid (produced with that trajectory by either scattering or decay) will drift to the wall and be absorbed in the curved section.

It is further discussed in [69] the requirements on the B field imposed by the above considerations in a quantitative way. It is found [68, 69] that the number of muons or electrons with sufficient energy to cause backgrounds is negligible when $dB_s/ds < -0.02$ T/m at any place where $|dB_s/dr| < 0.2$ T/m. The symbol s is position along the path of the global axis of the transport solenoid, and r is in the direction perpendicular to s .

3.2.9 Cosmic Rays

Cosmic ray induced electrons are potentially a limiting background and we have studied it using a GEANT simulation [70] of the detector and shielding. The details of the cosmic ray veto system required to reduce the background to a negligible level are discussed in Chapter 11. The conclusion of these studies is that cosmic ray background can be reduced to a negligible level with a combination of active and passive shielding and detection of extra particles in the tracking detector. The rejection techniques consist of the following:

- A passive shield of modest thickness (2 m of concrete and 0.5 m of steel).
- Three layers of scintillator veto counter surrounding the detector, with a combined efficiency for charged particles of 99.99% (1% inefficiency per layer).
- Selection criteria that eliminate events having significant evidence of extra particles in either the tracking or calorimetric trigger detector in time with the electron candidate.

With this suppression, the expected background from cosmic rays in 2×10^7 s is estimated to be ~ 0.016 events. An acceptable and somewhat higher background can be permitted by slightly relaxing either the passive or active suppression benchmarks — an optimization to be studied.

3.3 Physics Requirements

Based on the above discussions, we summarize the physics requirements below.

- In order to suppress the backgrounds from decay-in-orbit and radiative muon capture, the FWHM of the energy resolution function should be less than ~ 1 MeV with no significant high energy tails.
- Suppression of the prompt backgrounds from beam electrons, muon decay in flight, pion decay in flight and radiative pion capture calls for a pulsed beam structure with beam extinction of 10^{-9} .
- A thin absorber is needed at the center of the second straight section of the transport solenoid to absorb \bar{p} 's.
- Suppression of long transit time backgrounds puts stringent requirements on the B field inside the transport solenoid. At any place where $|dB_s/dr| < 0.2$ T/m, we require that $dB_s/ds < -0.02$ T/m, where s is the global axis along the path of the transport solenoid, and r is in the direction perpendicular to that of s .
- Suppression of cosmic rays requires a passive shield of modest thickness (2 m of concrete and 0.5 m of steel) and three layers of scintillator veto counter surrounding the detector with a combined efficiency for charged particles of 99.99%.

3.4 Background Level

Table 3.2 summarizes the expected background rates corresponding to the signal sensitivity given in Section 2.1. The backgrounds scale in different ways. We tabulate the backgrounds with the following assumptions. For μ^- DIO and radiative μ^- capture, the background scales with the number of captured μ 's, and hence is directly proportional to the sensitivity. The background from pattern recognition errors is a strong function of the noise rate, which is most sensitive to the rate of muon capture. It decreases rapidly with lower proton beam intensity. For cosmic ray backgrounds, we assume the sensitivity can be achieved in 2×10^7 seconds of data collection. This background is proportional to the total data collection time.

The background is dominated by electrons from muon DIO, which contribute a background/signal of 0.05 assuming $R_{\mu e} = 10^{-16}$. Substantial improvements in discrimination against this source of background can be made with modest loss in acceptance, as shown in Fig. 3.14. For example, the background/signal ratio can be decreased from 0.05 to 0.02 with a relative loss in sensitivity of less than 10%.

The next three largest contributions are due to radiative π^- capture, beam electrons, and μ^- decay in flight. These sources (and others identified with an asterisk

Sensitivity and Backgrounds

Table 3.2: The backgrounds from various sources, calculated for the sensitivity given in the previous table, and with scaling as discussed in the text. Backgrounds identified with an asterisk are proportional to the beam extinction and the numbers in the table assume 10^{-9} extinction. The number of background events corresponds to a 2×10^7 second data collection period, yielding a sensitivity of 4 events for $R_{\mu e} = 10^{-16}$.

Source	Events	Comment
μ decay in orbit	0.225	signal/noise = 20 for $R_{\mu e} = 10^{-16}$
Pattern recognition errors	< 0.002	
Radiative μ capture	< 0.002	
Beam electrons*	0.036	
μ decay in flight*	< 0.027	without scatter in target
μ decay in flight*	0.036	with scatter in target
π decay in flight*	< 0.001	
Radiative π^- capture*	0.063	from protons during detection time
Radiative π^- capture	0.001	from late arriving π^-
Anti-proton induced	0.006	
Cosmic ray induced	0.016	assuming 10^{-4} CR veto inefficiency
Total background	0.41	

in Table 3.2) are proportional to the proton beam extinction and we have assumed a value of 10^{-9} for this parameter in calculating the backgrounds shown.

3.5 Resolution Measurements

The experiment needs to demonstrate that it understands the momentum resolution of the apparatus, in particular both the gaussian and non-gaussian tails of the resolution function. Without a compelling measurement of the resolution function it will be difficult to convince anyone that a putative conversion electron signal of a few events is not simply decay-in-orbit events smeared into the signal region. Therefore we must have a program to measure the resolution of the detector *in situ*. This section describes several means by which the resolution function can be measured.

3.5.1 Decay-in-Orbit

The expected energy distribution of DIO events is shown in Fig. 3.10. This distribution plots the number of events vs. reconstructed momentum of the electron *at the tracker*. We make no attempt here to add back in energy loss in the foils or proton absorber. This is the distribution we will discuss in this Section, although eventually

we will use knowledge of the expected dE/dx as a cross-check, and one such check is described here.




Figure 3.10: The expected momentum distribution for decay-in-orbit events. The x -axis is the reconstructed momentum at the tracker.

It is useful to show the actual resolution function to be measured. The expected resolution for the straw-tube tracker with our current algorithm is shown in Fig. 3.11. The plot shows the value of the reconstructed momentum minus the value of the true momentum at the tracker. We used only conversion events with a true momentum of 104.97 MeV/c at birth for this plot. We see several interesting features:

1. The resolution function is asymmetric; the plot shows that because of dE/dx events are systematically reconstructed to lower momentum, supporting our claim that nearby DIO events are smeared away from the signal.
2. The resolution is approximately 120 keV on the high side where background smears into signal, and there is no significant tail.

This plot does not reflect the addition of extra hits from prompt sources causing catastrophic misreconstructions.

This resolution function is a critical distribution we must measure and reproduce in the simulations to prove we understand any signal. The special runs described below will be used to prove we understand the distributions in Fig. 3.10 and Fig. 3.11. However, given the limited statistics, we must gather more data. The full reconstructed distribution from 90 MeV on is shown in Fig. 3.13 with the acceptance immediately below in Fig. 3.14. This high-statistics distribution is accumulated over the course of the experiment and must agree with Monte Carlo.

We would not want to perform the experiment and then find we could not reproduce this distribution and the acceptance. Therefore we are exploring a sequence




Figure 3.11: The distribution of (reconstructed - true) momentum for conversion events at the tracker, in MeV/c. The low-side tail demonstrates that DIO events with momenta near the conversion momentum are smeared away from the signal region. The sum of all events is normalized to unity.

of resolution measuring runs. A simple way to increase the number of detected DIO events is to lower the magnetic field. Lowering the field allows events at lower momentum to pass through the octagon and since the spectrum is falling as $(E_{\text{max}} - E_{\text{electron}})^5$ we can greatly increase the statistics with a relatively small change in the field. As a criterion we imagine no test run should take more than a week. We have determined that lowering the field by 10% will yield 2.4×10^6 reconstructed DIO events/one-year run or 46K/week above 90 MeV/c, as shown in Fig. 3.13. We believe this will be sufficient to perform the required measurements. This is effectively close to measuring the acceptance above 100 MeV in the normal configuration at full field.

Individual Foil Measurements

One useful check is to use exactly one of the seventeen target foils at a time. This allows us to check geometric acceptances. Further, an “individual foil” run removes dE/dx losses from any downstream foils. We would therefore run each foil individually for one week; dividing the 46K above by seventeen would then yield 2700 reconstructed events above 90 MeV for a single-foil run, sufficient to determine the acceptance to good accuracy.




Figure 3.12: We examine decay-in-orbit events with an initial momentum at the foil greater than 90 MeV to gain in statistics. Here we take the side where decay in orbit events could smear into the signal and fit to a Gaussian, showing a resolution of 115 keV.




Figure 3.13: The distribution of momentum for DIO events > 90 MeV/c at the tracker, in MeV, at normal field (red) and with the field lowered by 10% (black).




Figure 3.14: The acceptance as a function of true momentum at the foil for DIO events > 90 MeV/c in MeV.

Multiple Foil Measurements

After we know the response of the detector to individual foils we can then check our knowledge of dE/dx and its effects by adding a second foil. The effect of dE/dx is largest for two consecutive foils since this would maximize the number of events produced in the first foil that would then pass through the second foil; we would propose two-foil runs and runs with the first, eighth, and final foils (obviously more study may change this precise plan.)

Annular and Central Foils

It will be useful to check the acceptance of the detector over the geometric surface of the foils. We would therefore propose a run with a small radius (4 cm) foil and a second run with an annular foil missing the central 4 cm. There would be no other foils for these runs.

Movable Target

We are also considering changing both the level and transverse distribution of decay-in-orbit events by moving the target along the x and y axes in an “Etch A Sketch”™ arrangement. Decay-in-orbit events could then pass through the octagon at lower momenta, greatly increasing the rate. Further, we could check the geometric acceptance as a function of azimuthal angle with such an arrangement.

$\pi \rightarrow e\nu_e$ Decay

Positive pions can be stopped and the momentum of electrons from $\pi^+ \rightarrow e^+\nu_e$, which should be half the charged pion mass, 69.7 MeV/c, can be measured with reduced magnetic field. Measuring the high-side tail could be invaluable in understanding smearing of DIO events to higher momentum and therefore this option must be studied.

There are a several disadvantages. The magnetic field would have to be reversed (to select positively charged particles) and its magnitude lowered by 30% (to place the electrons within the radial acceptance of the detector). As a result, the environment of the detector would be quite different: rates, particle mix, spectra, etc. would all change. Nonetheless a measurement of a clean Jacobian peak and a successful simulation of the smearing toward higher momenta would greatly increase the confidence in an observed signal. Another possible approach in which the field might not have to be reduced, would be to offset the position of the pion beam (or muons, in the case where DIO electrons are used for calibration) from the center of the Detector Solenoid, in order to get better detector coverage from the 69.7 MeV/c decay electrons. The pions could be moved off-center using dipole steering magnets in the Transport Solenoid.

Calibrated 100 MeV Source

An extremely interesting idea has come to light during the writing of this proposal. ILC cryomodules are capable of producing 100 MeV/c electrons with a small footprint. If we were to fire these electrons into the spectrometer from the *downstream* end, these electrons could spiral through the detector. Some fraction of them would reflect off the graded field at the beginning of the detector solenoid and then return through the spectrometer with the opposite helicity. These electrons would be an invaluable source of information about the detector response. They could be fired throughout the entire 1.7 μ s time window, not just the live gate. Therefore we could measure the change in resolution as a function of time (as a proxy for the prompt background) and extrapolate that result into the live window, as well as continuing to measure the result during the window itself. In addition if we had a fast trigger designed to look for candidate events, after each such event we could be sure to fire an electron into the same region of the spectrometer as an immediate check of the apparatus. A sketch of the arrangement is shown in Fig. 3.15. Note that the calorimeter would need to be reconfigured or the electrons will strike it.

3.5.2 Extinction Measurement

Section 6.5 described a scheme for achieving the required 10^{-9} extinction. How can we measure this?

In the AC-dipole scheme, we propose to measure the extinction factor throughout the 1.7 μ s cycle. There are approximately $2 \times 10^{13} \times 10^{-9} = 2000$ protons/s to be




Figure 3.15: The path of a 105 MeV calibration electron fired from a cryomodule at the downstream end of the detector. Note the reconfiguration of the calorimeter.

measured downstream of the extinction system, so just a few spills will suffice for the integral.

First, we would advance the phase of the AC-dipole by 90° so that the normally transmitted beam would strike the collimators and the normally extinguished beam would be transmitted. We would also perform a more subtle measurement by varying the phase of the dipole relative to the gate to “sweep out” the extinction rate as a function of time; this would allow us to extrapolate the measured extinction factor into the live gate. For this set of measurements we would need to insert a measuring device into the beam. We are collaborating with the Japanese through the US-Japan agreement to examine using a gas Cerenkov detector for this purpose (this agreement also covers the AC-dipole itself) [71].

We would also want to measure the environment in the tracker, at which time we would remove the measuring device and simply record the tracker environment as we vary the dipole phase.

3.6 Conclusions on the Sensitivity

We conclude this section with some general comments about the proposed sensitivity. First, the muon yield is now rather closely tied to experimental measurements of pion production cross sections. The technical issues involved in getting the design beam intensity have not been proved by example, but involve relatively modest extrapolations in accelerator performance. We rely on calculations of the muon beam parameters, and to the extent possible, these use realistic beam parameters. Nonetheless, there is always the possibility that the design intensity may not be realized. The consequence of not reaching the design beam intensity is that we would achieve a worse sensitivity

Mu2e Proposal

or we would run longer to achieve the proposal goals. The backgrounds scale with the sensitivity and are independent of the rate of data collection, with two exceptions. Backgrounds from pattern recognition errors scale as a power of the instantaneous intensity and will be smaller if beam fluxes are lower and the data are collected over a longer time. Cosmic ray backgrounds are proportional to the running time and are the only background source that would be adversely affected by longer running time to achieve the design sensitivity.

At the proposed sensitivity the experiment will not be limited by background. If backgrounds are higher than expected, many of them will be learned about early in the running when time is available to react. For example, if the detector rates are higher than expected, the background that is most sensitive to this (that from pattern recognition errors) can be substantially reduced with small changes in beam intensity since they scale as a power of the intensity.

Finally, while we believe that the proposed sensitivity can be achieved in the requested running time, even a result of a factor 2–5 less than our goal is an extremely sensitive test of models that allow lepton flavor violation and would represent a tremendous improvement over current experiments.

Cost Estimate and Schedule

4.1 Cost Estimate

The sources for the project cost estimate are itemized below.

1. Solenoid and Detector

We use the reviewed MECO costs at the end of RSVP to form a basis for our cost estimates [72]. These values are in actual year dollars assuming the MECO cost profile. We then escalate at 3.5%/yr from FY2005 to FY2009.

2. Cryogenics

Jay Theilacker and Arkadiy Klebaner of Fermilab's Technical Division have completed a preliminary cost estimate based on the design of Chapter 14.4 that can be found in Ref. [73].

3. Civil Construction

Fermilab's Facilities Engineering Services Section (FESS) has prepared a Project Definition Report (PDR) for all civil construction associated with the experiment [74]. This PDR documents one possible solution for the conventional facilities. It was prepared after extensive discussion with the collaboration.

This estimate includes the construction necessary to extract the beam from the Accumulator/Debuncher rings, transport it to a new experimental hall, and then build the hall and supply any needed physical infrastructure to the experiment. This cost should be regarded as preliminary but indicative; having identified the sources of costs we are prepared to make choices between the technical advantages of a given design and its costs and schedule.

The schedule shows roughly five years, but neither we nor FESS have studied this carefully and for now we will assume it is *not* on the critical path.

FESS follows a set of defined rules according to the level of understanding of the project, detailed in the Project Definition Report. They also include EDIA, overhead, profit, and a contingency based on the level of the data presented to them.

4. Fermilab Beamline

The Accelerator Division has prepared a preliminary cost estimate. A document supporting this estimate is in preparation.

5. Project Management

The then-year dollar value for Project Management for NO ν A was estimated at \$7.0M as presented at the CD-2 review. We use the identical number in the absence of a better method. Since this number is in then-year dollars for NO ν A's larger project it may be a slight over-estimate.

6. Contingency

Apparatus and Solenoids

The apparatus and solenoid magnet system described here come from the much reviewed MECO design, whose cost estimates were quite mature and backed by an extensive set of basis of estimates. Nevertheless, the Wojcicki panel, in one of the final RSVP reviews, recommended an overall contingency of 45% for RSVP [75]. Although in some cases this is significantly more than the MECO estimates, for example the Wojcicki report states that for the solenoids “the 45% top-down cost contingency seems to be an overestimate,” we have chosen at this time to apply a conservative overall contingency of 50% to the apparatus and solenoids. As we move forward we will develop a bottoms-up resource-loaded cost and schedule.

Project Management

For the Project Management cost we apply the NO ν A contingency of 35%.

Civil Construction

For the Civil Construction we apply the contingency of 31% recommended by FESS.

Beamline

For the beamline we apply 100% contingency (which is consistent with NO ν A methodology) given the conceptual status of the design.

Our estimate of the project cost is itemized in Table 4.1. The base cost is \$122M and the total cost, including a 48% overall contingency, is \$180M. This estimate includes design and construction and is in FY09\$. Additional funds required for R&D are itemized in Chapter 5.

There are many opportunities available for value engineering as we move ahead. The solenoids and the civil construction are the highest cost items and will receive the most attention. Many ideas exist to optimize these costs. It should also be noted our cost estimate is not the typical cost estimate that accompanies most proposals. Most proposals quote a cost that is unsupported by any basis of estimate and the cost from proposal to CD-2 typically experiences significant escalation as a result. *Mu2e* is in the unique position of having the MECO Work Breakdown Schedule (WBS) and

Cost Estimate and Schedule

a preliminary design document from FESS for the civil construction. In addition, we have included what we believe to be a conservative contingency that accurately captures the current level of uncertainty. As a result, we do not expect substantial cost increases as we move forward.

Table 4.1: Preliminary estimate of the total project cost.

Item	M&S) (k\$)	Labor (k\$)	Base Cost		Contingency		Total (k\$)
			FY05 (k\$)	FY09 (k\$)	Rate	(k\$)	
Production target							
and shield	\$2,490	\$219	\$2,709	\$3,109	50%	\$1,554	\$4,663
Muon beamline	\$1,209	\$1,265	\$2,474	\$2,839	50%	\$1,419	\$4,258
Straw tracker	\$2,280	\$998	\$3,278	\$3,762	50%	\$1,881	\$5,642
Calorimeter	\$3,466	\$1,177	\$4,643	\$5,328	50%	\$2,664	\$7,992
Cosmic ray veto	\$1,003	\$306	\$1,309	\$1,502	50%	\$751	\$2,253
Trigger and DAQ	\$884	\$584	\$1,468	\$1,685	50%	\$842	\$2,527
Installation	\$125	\$1,266	\$1,391	\$1,596	50%	\$798	\$2,394
Solenoids	\$35,061	\$12,115	\$47,176	\$54,136	50%	\$27,068	\$81,203
Cryogenic							
infrastructure	\$1,556	\$1,988		\$3,544	50%	\$1,772	\$5,316
Accelerator/beamline	\$4,747	\$3,793		\$8,540	100%	\$8,540	\$17,080
Civil construction	\$25,661	\$5,456		\$31,117	31%	\$9,643	\$40,760
Project office	\$556	\$4,535		\$5,091	35%	\$1,782	\$6,873
Total:	\$79,038	\$33,702		\$122,248	48%	\$58,715	\$180,962

4.2 Schedule

We assume a technically limited schedule that requires funds and resources to be available when needed. Assuming approval of this proposal and CD-0 in February 2009, we need to understand (1) the time required to advance to a final design, and (2) the time required for construction.

The solenoid system is the critical path for the project. We estimate that it will take approximately one year for Fermilab's Technical Division to assimilate and update the MIT conceptual design report. The MECO WBS and project schedule (GANTT chart) indicate approximately two years for the final design, and five years to construct, install, test and commission the system [75, 76]. This brings us from February 2009 for CD-0 to early 2016 for Project completion.

Mu2e Proposal



Figure 4.1: Mu2e technically limited schedule for the critical path item, the superconducting solenoid.

R&D Request

5.1 Introduction

A SIGNIFICANT amount of R&D will be required to advance to a final design between now and the beginning of the construction phase of the project. Some R&D tasks may overlap the beginning of construction. R&D is required to help decide between competing ideas, to optimize designs, maximize performance and understand costs. We assume contributions of effort from Fermilab and the *Mu2e* collaborating institutions to accomplish this R&D.

In this chapter we outline the broad *Mu2e* R&D program while acknowledging that our program is still immature. As we move forward we will develop a more comprehensive and detailed plan.

5.2 Specific R&D Tasks

5.2.1 Solenoid System

The magnet system for the *Mu2e* experiment is arguably the most critical element in technical risk, cost and schedule. The baseline proposal for the *Mu2e* experiment is to use the magnet design for the Brookhaven-based MECO experiment as much as possible. This is a very reasonable approach, since the magnet system for MECO has been favorably reviewed several times from 2001–2005, advancing to a stage that is roughly equivalent to a magnet conceptual design. In *Mu2e* it is natural for the Fermilab Magnet Program to take the lead role in the management of this magnet project. We have an accomplished staff of magnet scientists and engineers that have considerable experience in the design, tooling development, manufacturing, testing and analysis on numerous magnet projects.

The unique history of the *Mu2e*/MECO magnet design calls for a plan that makes the most efficient use of existing documentation and expertise. A conceptual design report was published on June 6, 2002 [77, 78]. Since then, there have been several incremental improvements to the conceptual design, documented in technical notes and drawing updates but the CDR documents were never officially updated. Most

Mu2e Proposal

of the conceptual design work was performed by a team of engineers and scientists in the MIT Plasma Science and Fusion Center. Essentially all of the team members have left MIT.

We must take a number of steps in order to confidently use the existing design effort. Our goal would lead to a revised Conceptual Design Report for the magnet system of *Mu2e* and then to Project CD-2/3a:

1. A complete compilation of existing up documentation: documents and drawings for the most up-to-date status of the *Mu2e* magnet.
2. A thorough review of the existing documentation, with changes to the existing magnet system design if warranted. We should then re-perform all of the magnetic, electrical, quench and mechanical calculations to validate the design, focusing on
 - (a) magnetic field and tracking studies to verify magnetic design
 - (b) updates in conductor and materials, in light of any recent technological advances
 - (c) identifying areas where the conceptual design function or interface is incomplete on the CDR level
3. Propose a plan to bring the magnet design toward an engineering design, with sufficient detail to complete a successful CD-2/3a project.

We propose that Fermilab hire outside contractors to assist in the timely completion of Items 1–2 of this R&D proposal. Several key members of the MIT team are at General Atomics (GA) and therefore GA might be a natural choice. However, we must certainly consider cost, schedule, and technical expertise before awarding a contract and all variables must be considered. For Item 1, members of the original design time are in the best position to assemble all documentation. The deliverables would be a written report accompanied by face-to-face presentations with the *Mu2e* collaboration and the Fermilab Magnet Department. For Item 2, contractors would be available to debrief collaboration on design choices, calculations, and the interpretation of data. In some cases the contractors may be asked to re-perform calculations with local *Mu2e*/Fermilab oversight.

Depending on the results of this first phase, a plan will be developed which could include further contract participation. The first phase effort would require 2 FTE contractors for roughly six months, which we expect would cost approximately \$250,000.

5.2.2 Tracker

Significant, but small-scale, prototyping has been performed for the L-tracker [79]. Only nine straws 0.5 meter long straws were used in this test; questions of electronic

noise and mechanical stability at the full size must be answered. It is imperative to construct a full scale prototype of a three-layer resistive straw system as in the L-tracker. A preliminary estimate of the cost and manpower for the L-tracker has been developed in consultation with PPD [80]. This plan will likely evolve as we try to understand the details more clearly for this critical piece of our R&D.

5.2.3 Extinction System and AC Dipole

The intensity of the proton beam must be suppressed by a factor of 10^9 between pulses in order to prevent pions and other beam related background from arriving at the stopping target during the measurement period (recall that pion radiative capture on a nucleus is a potentially limiting background.) This is the so-called extinction requirement, described in Chapter 6.

A key component of the extinction scheme will be a pair of AC dipoles synchronized with the beam such that in-time beam will pass through a narrow channel in a collimator designed to absorb all out-of-time protons. The AC dipoles must be continuously powered by a sine wave resulting in large power dissipations in the magnet yoke and driving conductor, so the choice of materials is strongly coupled with the magnet design and performance. In addition, there is a high voltage applied to the conductor and the electrical insulation reliability must be understood.

The first steps in the R&D will be to procure and evaluate various ferrite materials and to study the conducting leads. Beyond that we plan to build a short 0.5 m long model to investigate all of the issues in the magnet design, fabrication and performance. About \$65K is required to procure the ferrites and the conductor; \$20K of this is being provided through the US-Japan agreement. Significant engineering effort will also be required to move the AC dipole design to the CDR stage and to develop the 0.5 m long model. A detailed plan for a conceptual design of the AC-dipole has been developed [81]. Based on this schedule, we would request the already existing group be allowed to proceed at 25% effort, about 1 FTE engineer spread among several people during FY2009.

5.2.4 Accelerator

Mu2e does not require major upgrades to the accelerator complex because of the Boomerang scheme. Since the Accumulator and Debuncher will be used in a new way (8 GeV protons instead of 120 GeV \bar{p}) we must understand what changes will be required. We require a new slow extraction system from the Debuncher, although such systems are not new.

RF Systems The voltage requirements for the momentum stacking system needs to be finalized, including estimates of potential beam loading. Further optimization of the currently understood bunch formation process should be considered,

as well as alternative methods that could help reduce the costs of the required RF systems.

Intensity Limitations:

Estimates need to be performed on the foreseen limitations on intensity in the Accumulator and Debuncher rings when reconfigured for *Mu2e* operation. Many small aperture and large impedance devices can be removed from the stochastic cooling systems, but new RF cavities with higher-than-present voltages, operating with higher-than-present beam currents will be installed. The Accumulator will run at 2–3 times its highest intensity to date, though not stored for hours as at present. The beam will be stored in the Debuncher for approximately 667 ms compared to the typical ~ 2100 ms of present operation, but will contain about 40,000 times more particles than during antiproton production.

Resonant Extraction:

A more developed plan for resonant extraction and its modeling, including the effects of beam momentum spread and realistic apertures, is required. Shielding around the extraction septum area to locally maintain losses of 1-5% needs to be designed. The extraction inefficiency needs to be better estimated. The requirements and expectations for the slow spill feedback circuit need to be developed.

Extraction Line:

A full design of the extracted beam line is required, including appropriate matching into and out of the extinction channel and adequate resolution of physical constraints between the ring and experimental hall.

Extinction Channel:

The extinction channel needs to be taken from conceptual layout to an engineering design, along with appropriate specifications for the required instrumentation for measuring and monitoring the level of achieved extinction.

Radiation Safety:

A careful analysis of the necessary safe guards for running high intensity beams in the antiproton enclosures needs to be performed. Beam loss rates several orders of magnitude greater than present are expected to be encountered. Passive, active, and perhaps electronic safety measures will need to be designed and costed.

Instrumentation:

At present the antiproton rings primarily contain debunched beam whereas the *Mu2e* experiment will form bunched beam in both the Accumulator and

Debuncher rings. An analysis of the present instrumentation and possible modifications or upgrades necessary to monitor bunched beam will need to be performed.

5.2.5 Cosmic Ray Shield

Significant R&D will be required to achieve the design level of 10^{-4} for overall cosmic ray rejection for a large detector. The efficiency of a multi-layer, solid scintillator system must be evaluated and optimized. Among the parameters to be studied and optimized are the fluor type and concentration in the scintillator, the fluor concentration in the WLS fibers, the optimum configuration and density of WLS fibers, the optimum means of embedding the fibers in the scintillator (holes or grooves) and the photodetector. MINOS and others have done considerable work in this area, but many of the parameters listed above optimize at different values for different lengths and geometries because of competing processes. For example, higher fluor concentration in the WLS fiber results in more light but also increases the absorption of light making its way along the fiber to a photodetector. A configuration optimized to *Mu2e*'s precise geometry is required.

We also wish to obtain some of the novel Geiger-mode APDs that have become available from several manufacturers and compare their performance to that of the photomultipliers in the baseline design.

5.2.6 Calorimeter

The collaboration intends to obtain ten PWO-II crystals and ten LSO crystals to measure their performance, with bench tests and 100 MeV test beam electrons. The tests will be made with cooled and room temperature crystals. We also need to obtain large area avalanche photodiodes (APDs) from Hamamatsu and Radiation Monitoring Devices (RMD). They will be tested on the bench and also used in the test beam. Preamplifier/shapers will be fabricated for the APDs. The design of the cooling system and calorimeter mounting scheme needs to proceed in order that realistic energy resolution test can be made with small-scale mock-ups of both, particularly with the LSO calorimeter, where we intend to push the performance to the limit. Hence part of our request is for this design work. Finally, a light pulser system for calibration needs to be designed.

5.2.7 Spectrometer Calibration System

One of the keys to the success of the *Mu2e* experiment is to understand the resolution function of the spectrometer near the endpoint of the decay-in-orbit spectrum. A concept exists to develop an in-situ calibration technique using an electron gun that would gauge the precision, absolute energy scale, and level of systematic uncertainties

associated with the spectrometer while the experiment is live and recording physics data (see Chapter 3). Some time will be required to develop a concept followed by a prototyping stage. The immediate request is for undergraduate student salaries and computing resources.

5.3 R&D Cost Estimate

Table 5.1 provides a rough cost estimate for the *Mu2e* R&D tasks. Costs required to reach final designs are not included. The items in Table 5.1 include the cost of consultants and non-scientific labor at collaborating institutions. Fermilab engineering and technical help would be in addition to this total.

Table 5.1: Requests for Mu2e R&D through CD-3 and for FY2009. The solenoid system design amount is taken from the MECO WBS and inflated for four years at 3.5%/yr.

Item	Total R&D	FY2009 R&D
Completion of Solenoid System Design	\$6,900K	\$250K
Completion of EDIA for Conventional Facilities(FESS)	\$2,225K	\$150K
Tracker Prototypes and Conceptual Design	\$250K	\$100K
Extinction R&D and AC-Dipole	\$100K	\$45K
Calorimeter R&D and Beam Test	\$252K	\$175K
Prototype Veto Module and Conceptual Design	\$200K	\$50K
Electron Gun Calibration System R&D	\$200K	\$50K
TOTAL	\$10,127K	\$820K

Part II

Mu2e in Detail

Proton Beam

6.1 Overview

THE MECO experiment was originally designed around a total data set of 4×10^{20} protons, collected in one to two years. When the planned upgrades to Fermilab’s Proton Source [82] are implemented, there should be an excess proton capacity of this order even during the NO ν A era [83]. The challenge is that there is no practical way to produce the desired beam structure directly from the Booster.

We propose to solve this problem by reusing the 8 GeV Debuncher and Accumulator storage rings, which are both housed in the pBar beam enclosure. Currently, antiprotons are transported from the production target into the outer Debuncher ring, where they are phase-rotated and stochastically cooled. They are then transferred into the inner Accumulator ring, where they are “momentum stacked” onto the core and further cooled for use in Tevatron collisions.

Our proposal is illustrated in Figure 6.1. Booster proton batches would be transported through the Recycler and injected directly into the Accumulator, where several batches would be momentum stacked. These would then be transferred into the Debuncher ring and rebunched into a single short bunch. Finally, the beam would be resonantly extracted, such that this single bunch would become a bunch train. The $1.7 \mu\text{s}$ period of the Debuncher would produce a structure almost ideally suited to the *Mu2e* experiment. An important specification for this beam will be the “extinction factor”, or suppression of out-of-bucket beam, as this is a limiting background for the experiment. This will be discussed in more detail shortly.

6.2 Delivering Protons to the Accumulator

Initially, it was believed that injecting protons into the Accumulator during NO ν A running would require a new transport line to be built from the Booster to the Debuncher/Accumulator enclosure. However, an ingenious boomerang scheme [84] has been proposed which would allow Booster beam to be transported to the Accumulator with no civil construction and indeed minimal beam line modifications of any




Figure 6.1: The relevant parts of the accelerator complex are shown. At right is a detail of the antiproton ring, showing the Debuncher (outer ring) and Accumulator (inner ring). The complete path taken by the protons in the proposed boomerang scheme is shown at left.

kind. This will greatly reduce the burden on the accelerator complex of supporting this experiment.

A key part of the NO ν A accelerator upgrades will be a modification to the MI-8 beam line to allow beam to be directly injected into the Recycler [85]. Booster batches will be slip stacked in the Recycler prior to being loaded into the Main Injector. This will eliminate the time currently spent loading the Main Injector and increase the total proton rate to the NuMI line. In the NO ν A plan, 12 Booster batches will be stacked into the Recycler during every 1.333 second Main Injector cycle. The NO ν A timeline is shown in Figure 6.2. Because this period corresponds to twenty 15 Hz Booster cycles, there are potentially up to eight extra batches available. The baseline design of the *Mu2e* experiment would use six of these, due to longitudinal emittance constraints in the Accumulator/Debuncher.

In order to take advantage of this unused part of the timeline, we would have to add a simple extraction region to the Recycler to direct beam into the existing P150 line, as described in [84]. Protons from the Booster would then make only a partial circuit of the Recycler, after which they would be transported to the Accumulator in the same way we currently transport small numbers of “reverse protons” from the Main Injector. Because the protons do not go all the way around the Recycler, extraction could be done with a simple switched magnet, rather than a kicker. Figure 6.3 shows potential locations for the extraction from the Recycler to the existing P1 line.

With this timeline, up to eight Booster batches can be delivered to the Accumulator every 1.33 second Main Injector cycle. If we conservatively assume Booster

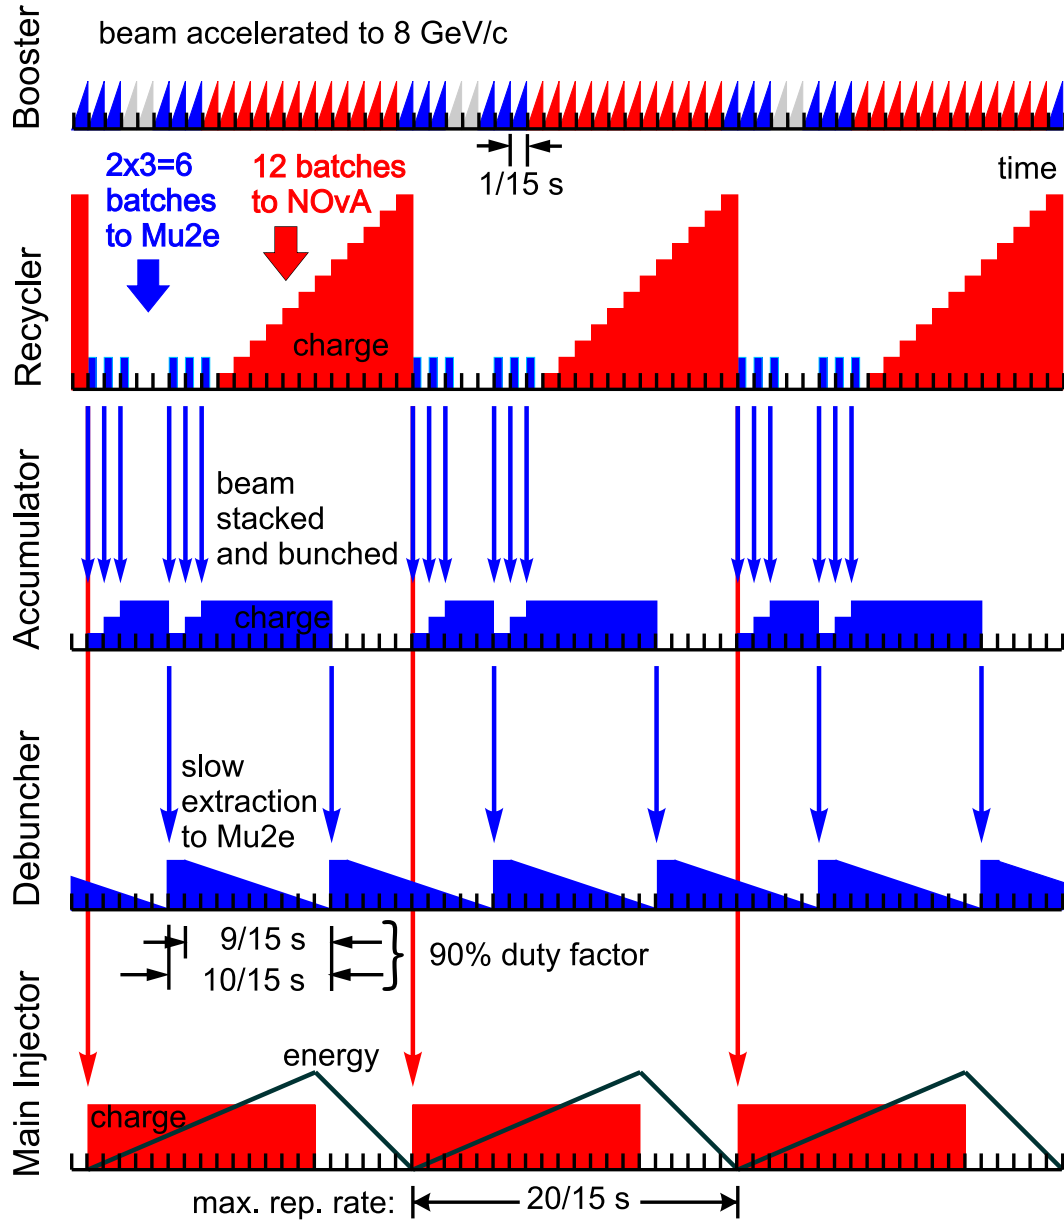


Figure 6.2: A schematic illustration of the timeline for 15 Hz Booster batches in the NOvA era. NOvA proton batches are shown in red, Mu2e in blue. Twelve Booster batches are stacked in the Recycler and then transferred all at once to the Main Injector, eliminating the loading time and increasing protons to the NuMI line. Six of the eight unused Booster batches available while the Main Injector is ramping are sent to the Antiproton Accumulator, three at a time, where they are stacked and bunched and then sent to the Debuncher Ring.




Figure 6.3: Potential locations for transfer lines from the Recycler Ring to the P1 line.

batches of 4×10^{12} protons, this could provide as many as 4.8×10^{20} protons per year to this effort assuming that the total Booster flux could be increased enough to accommodate this. In practice, longitudinal emittance in the Accumulator/Debuncher will likely limit extraction to no more than 6 batches at a time, or 3.6×10^{20} per year, and we will assume this in our baseline planning.

6.3 Momentum Stacking and Rebunching

Momentum stacking in the Accumulator is a straightforward modification of what is done now with antiprotons, and the momentum stacking of protons was discussed in detail for the proposed SuperNuMI (SNUMI) upgrades [86]. Figure 6.4 illustrates the scheme. Protons are injected at an energy that corresponds to the outer orbit in the Accumulator and then captured and decelerated near the core orbit. A large slip factor insures that protons may be stacked very close to the circulating beam without the longitudinal “white space” usually associated with slip stacking and barrier bucket techniques.

We are investigating several schemes to arrive at a single short bunch in the Debuncher. Our baseline proposal is a hybrid scheme, in which the stacked beam is bunched by a 4 kV $h=1$ RF system in the Accumulator, then transferred to the Debuncher, where it undergoes a 90 degree phase rotation by a 40 kV $h=1$ RF system, followed by a capture by a 250 kV $h=4$ RF system. The resulting single bunch has a 38 ns rms length and an energy spread of ± 200 MeV. A simulation of this rebunching is shown in Figure [87]. This scheme has the advantage that the $h=1$ capture in the Accumulator automatically generates a beam free gap to allow for the rise time of the transfer kicker.

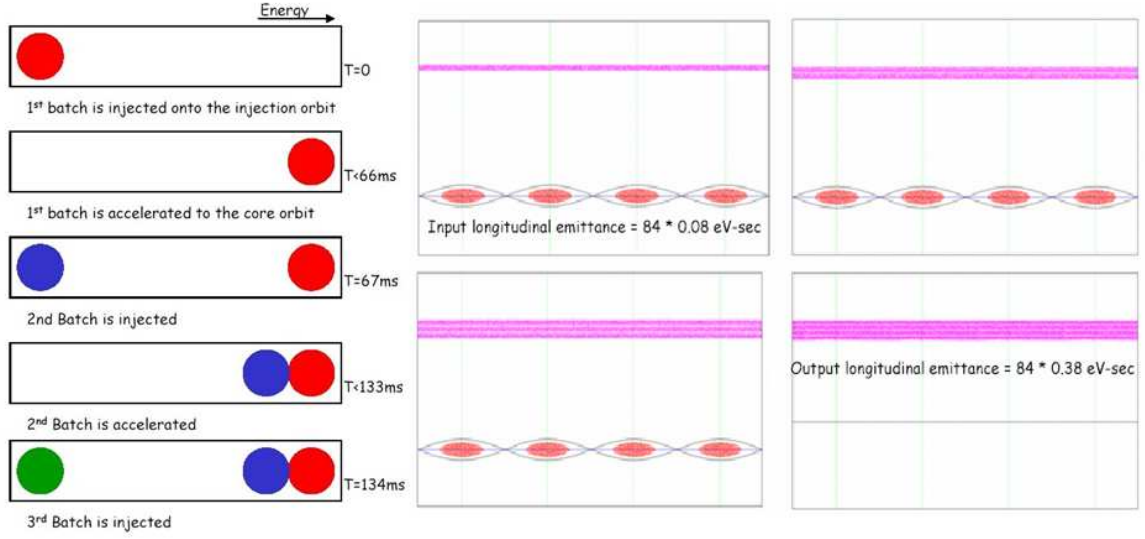


Figure 6.4: The left figure shows a conceptual schematic of momentum stacking. On the right is shown a simulation of the capture and momentum stacking of four Booster batches.

6.4 Resonant Extraction

Resonant extraction is a well established technique to extract beam slowly from a synchrotron. It involves moving the tune of a circulating beam close to harmonic resonance, such that beam becomes unstable and migrates to high amplitude. Generally, the high amplitude particles are intercepted by an electrostatic septum, in which the field is produced by a very thin wire plane, followed by a Lambertson magnet approximately 90° later in betatron phase. In practice, two types of resonant extraction have been widely used:

- Half integer resonance, in which the tune is moved near $\nu = m/2$, where m is an odd integer. The resonance is driven by a set of properly phased quadrupoles. Octupoles are then excited to produce an amplitude dependent separatrix.
- Third integer resonance, in which the tune is moved near $\nu = m/3$ (ν not integer). The resonance is driven by properly phased sextupoles. The separatrix is controlled through tune variation and sextupole strength.

In principle, either (or both) could be used in the Debuncher. Historically, Fermilab has chosen half-integer extraction for a variety of reasons; however, we will choose third integer because the existing working point of the Debuncher is close to a third integer resonance, and because there is wide experience with third integer resonant extraction worldwide. Also, interesting techniques are currently being developed to increase the efficiency of third integer extraction, which we might hope to exploit [88].

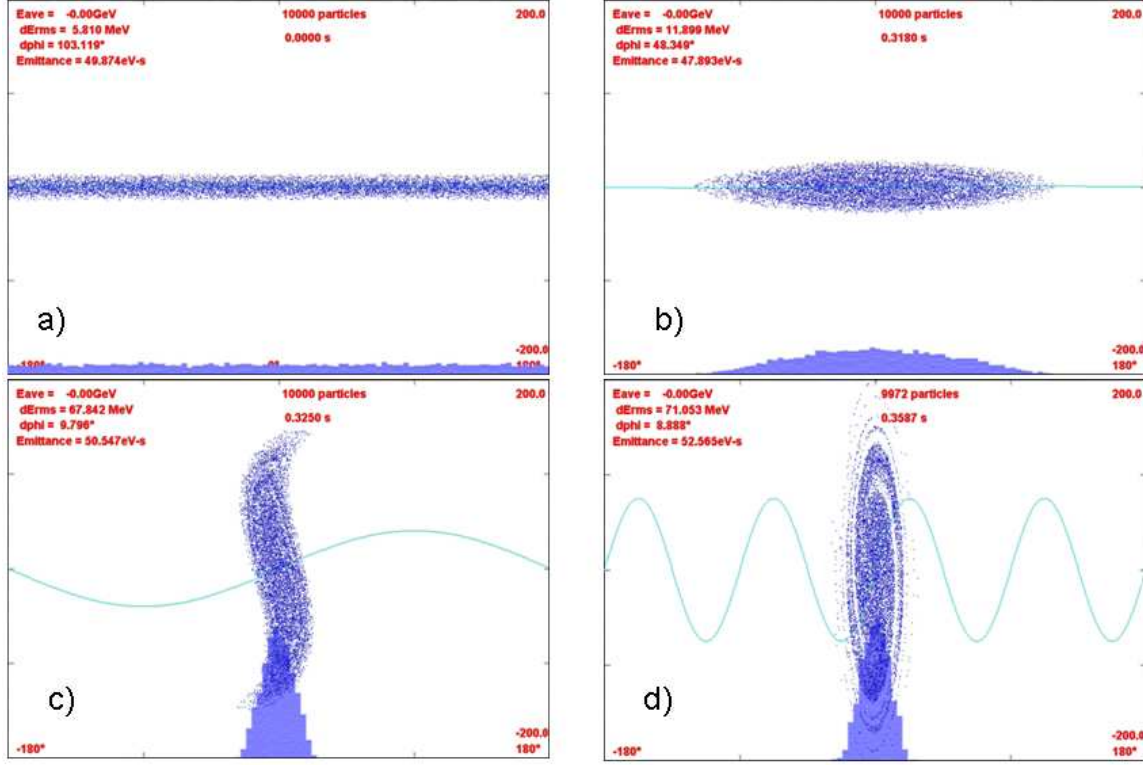


Figure 6.5: The hybrid rebunching scheme. Figures a) and b) show the $h=1$ capture in the Accumulator. After this step, particles are transferred to the Debuncher, where c) shows the $h=1$ phase rotation. The final $h=4$ capture is shown in d).

The inefficiency of any resonant extraction scheme is proportional to w/d , where w is the thickness of the septum plane and d is the width of the septum gap, so it is advantageous to choose a gap as large as possible while staying within the acceptance of the machine. Figure 6.6 shows the proposed location for the extraction septum and Lambertson. An electrostatic septum would be located between the Q103 and Q102 quadrupoles. If we assume the same specifications as a single septum tank from the Main injector — 80 kV over 1 cm by 3 m long — we get 2.5 cm of deflection at the downstream end of quadrupole Q101, the proposed location of the extraction Lambertson. A magnetic field of 0.8 Tesla is required to clear the downstream Q602 quadrupole. This could be accomplished by a 1 m version of a Main Injector Lambertson (± 5 inch extraction channel) followed by a 2 m C-magnet.

The existing tune working point of the Debuncher is $\nu_x/\nu_y = 9.764/9.785$, making the $\nu_x = 29/3$ a logical resonance to exploit for slow extraction. Table 6.1 summarizes the specifications for the slow extraction based on an extraction septum and Lambertson with similar specifications to those used in the Main Injector. Figure 6.7 shows a preliminary OptiM simulation of this resonance, assuming the driving sextupoles are located just inside of the 07 quadrupoles in each straight section. The position of the

Proton Beam




Figure 6.6: Proposed location for resonant extraction system. The electrostatic septum would be placed between quads Q103 and Q102 and the Lambertson would be placed between quads Q101 and Q602.




Figure 6.7: Preliminary simulation of a third order resonance in Debuncher. The position of the extraction septum is superimposed. Resonance parameters are shown in Table 6.1.

extraction septum is superimposed.

6.5 Proton Extinction

The two principal sources of background in *Mu2e* are muon decays-in-orbit (DIO) and prompt processes in which the detected putative conversion electron is produced

Mu2e Proposal

Table 6.1: The approximate parameters of the third integer resonant extraction, with the septum located between the Q101 and Q602 quadrupoles.

Kinetic Energy (GeV)	8
Working tune (ν_x/ν_y)	9.769/9.783
Resonance (ν_x)	29/3
Normalized acceptance (x/y π -mm-mr)	285/240
Normalized beam emittance (π -mm-mr)	20
β_x at electrostatic septum (m)	15
β_x at Lambertson (m)	22
β_x at harmonic sextupoles (m)	14
Septum Position (mm/ σ)	11/4.8
Septum gap/step size (mm)	10
Sextupole Drive Strength (T-m/m ²)	473
Initial Tuneshift ($\delta\nu$)	.048
Septum field (MV/m)	8
Septum length (m)	3

by a beam particle arriving at the stopping target. The first category of background can be reduced by improving the electron energy resolution. The second category is reduced by delivering the proton beam in short bunches separated by $\approx \tau_\mu$ in Al (864 ns), ensuring near perfect proton extinction between bunches, and restricting the search for the conversion of stopped muons to the inter-bunch period.

Table 6.2: The most significant beam backgrounds, for 3.6×10^{20} protons on the primary target and an extinction factor of 10^{-9} . There would be 4 signal events if $R_{\mu e} \approx 10^{-16}$.

Radiative π Capture	0.06
Scattered Electrons	0.04
μ Decay in Flight	0.06
π Decay in Flight	<0.001
Total	0.16

The beam-induced processes simulating conversion are: radiative pion capture that can yield photons with energy of up to 140 MeV; scattered beam electrons; and the decay in-flight of beam muons or pions in the region of the target. In order to reduce these backgrounds, there must be significant proton beam extinction during

the 1 μsec intervals in which the experiment is live. Extensive GEANT simulations of these background processes were performed for MECO [89]. Table 6.2 shows the numbers of events produced by each of these backgrounds for 3.6×10^{20} protons incident on the production target, assuming an inter-bunch proton extinction of 10^{-9} . This corresponds to about half the total background under an expected signal of 4 events for $R_{\mu e} = 10^{-16}$.




Figure 6.8: The AC dipole scheme for beam extinction. A matched pair of dipoles is synchronized with extracted beam bunches such that only in-time beam can pass through a collimator or series of collimators.

Developing the extinction scheme will be an important part of producing our final proposal. The level of extinction will be challenging both to produce and to verify. Ensuring extinction will likely involve several steps. A key component will be a pair of AC dipoles, as illustrated in Figure 6.8. These will be synchronized with the beam such that in-time beam will pass through a narrow channel in a collimator designed to absorb all out-of-time protons. The optimization of this channel is discussed elsewhere [90]. The cost of the AC dipole system is minimized by a beam line design where the β function at their location is large in the bend plane and small in the non-bend plane [91]; however, this makes it challenging to achieve the required phase advance in an extinction channel of reasonable length. The most straightforward design involves six quadrupoles, arranged in mirror symmetric fashion about the collimator. Figure 6.9 shows such a design, corresponding to approximately the shortest layout which can satisfy both the β and phase advance constraints. Studies are currently under way to determine if this scheme can achieve the necessary extinction.

We are also investigating other methods for achieving the necessary level of extinction, which may simplify the magnet design.

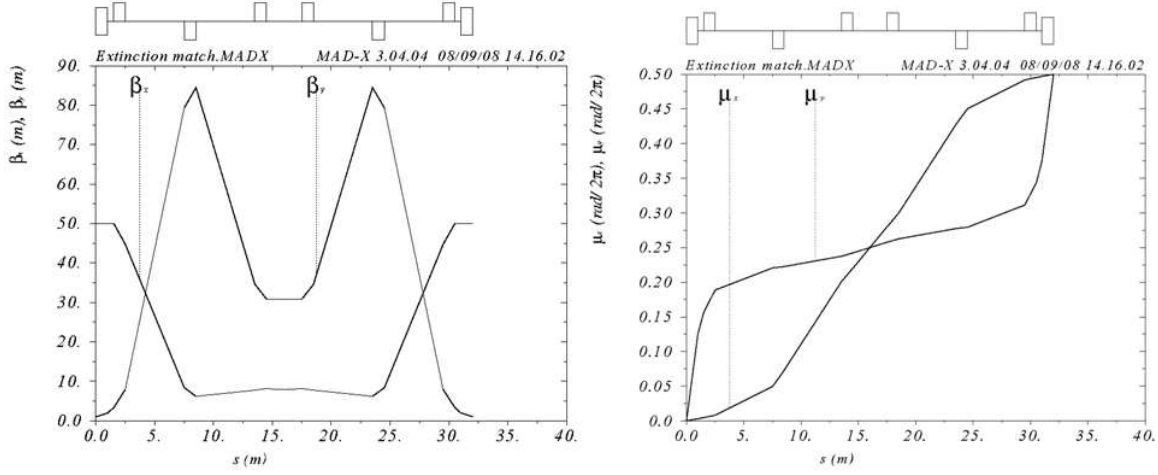


Figure 6.9: Preliminary extinction channel design.

6.6 Total Proton Delivery

In order to supply protons to this experiment, the Proton Source (the Fermilab Linac and Booster) will have to supply protons beyond the needs of the NO ν A program. This will require two things:

- The Booster will have to run beyond the 10.5 Hz average repetition rate required by the NO ν A experiment.
- The efficiency of the Booster must be increased so that the same total beam loss in the tunnel is maintained or reduced as more protons are accelerated.

The rate issues were addressed in a study done recently [92] and it was found that a fairly modest refurbishment program can bring the Booster to continuous 15 Hz operation. It is planned to implement these improvements adiabatically over the next few years out of the Accelerator Division's operating budget, and they should be in place well before the *Mu2e* experiment begins taking data.

As for the Booster efficiency, the Proton Source has already demonstrated a total throughput more or less sufficient for the NO ν A era. At that point, protons to the NO ν A experiment will be limited by the Main Injector capacity, so any improvements to the Proton Source would translate directly to excess protons available for an 8 GeV program, such as *Mu2e*. The requirements to increase the efficiency are summarized in [93]. These should be accomplished through projects that are currently being implemented; namely, an ambitious new corrector system to reduce closed orbit distortions, and a new chopper to increase the efficiency for the creation of the extraction “notch” in the circulating Booster beam.

We have assumed that the Booster can stack and extract six batches to the Accumulator/Debuncher each 1.33 second Main Injector cycle and still maintain acceptable longitudinal emittance. This would correspond to roughly 3.6×10^{20} protons per year. We will develop a realistic proton delivery timetable as part of our proposal, but it is reasonable to estimate that we can deliver the 3.6×10^{20} protons needed by the experiment in something on the order of two years.

6.7 Radiation Safety

Radiation safety is a critical issue for this project. The Accumulator/Debuncher enclosure was designed for an extremely low flux of antiprotons and so is not heavily shielded. The earth shielding is roughly 13ft in most outside areas and is 10ft or less in the service buildings. Tunnel roof loading limits preclude increasing this overburden, which is more than 10ft short of the passive shielding requirements (“Cossairt Criteria”) [94] for the proposed intensity. Unlike the Booster, it is practical to fence in the entire area and declare it “Limited Occupancy”. This will ameliorate the situation somewhat, but it is still certain we will need some sort of safety system involving interlocked detectors (“chipmunks”) of the sort that protect the Booster. It is probable that the service buildings will be locked out entirely while the beam is running at high intensity. Because of the large amount of documentation needed for the shielding assessment, we will begin working on it as soon as possible.

Production Target

7.1 Introduction

THE average number of muons reaching the stopping target per incident proton depends on many factors. Critical to achieving a high stopped muon yield, we must collide protons on a target with as large an atomic number and density as possible. Yield, as well as cooling, depends on target geometry, position and orientation. More generally, muon yield is sensitive to solenoid design and material anywhere in the clear bore of the magnet. Hence, the target, cooling system, and mechanical support must be compact and introduce little additional mass beyond that of the target, thus minimizing pion reabsorption. The target mounting must be sufficiently insensitive to vibrations to maintain good positioning. Furthermore, temperature must be controlled to avoid thermal stress levels that may lead to mechanical failure or geometrical distortion. Still other factors contribute to the stopped muon yield, which are discussed elsewhere in this document; e.g., proton beam energy, production and transport solenoid fields, and the design of the collimators used to select muons of the appropriate momentum and charge.

Production target design has undergone significant evolution. A radiation-cooled target design that is mechanically stable, with acceptable stopped muon yield has been illusive. A water-cooled system with drastically lower operating temperatures has been shown to be feasible. This system results in less than about five percent reduction in stopped muon yield compared to a similar radiation cooled unit. The theoretical and experimental research surrounding target design is discussed.

Negative muon production is accomplished by directing a high-intensity 8.0 GeV proton beam onto a small dense metal target, nominally the size and shape of a pencil, located within a solenoid with an axially graded magnetic field. Approximate target position is indicated by the black star in Fig. 7.1. The primary proton beam strikes the target end-on to produce pions that decay into muons. The proton beam is pointed in the direction of increasing field intensity so that charged secondaries spiral along in a magnetic bottle towards the “closed end”, reflecting many charged particles back towards the TS entrance, thus collecting pions over a large solid angle. A similar system was discussed by Djilkibaev, Lobashev, & collaborators [50, 95], and their ideas were later adopted by proponents of the muon collider [96].




Figure 7.1: Production Solenoid; the grey structure is the iron return yoke. The interface to the TS can be seen at left. The heat and radiation shield, resting against the inner cryostat wall (orange), is visible in this cut-away picture. The shield is composed mostly of copper (purple) and tungsten (red) in the regions of greatest energy deposition.

In particular, pions produced in the target with transverse momentum below ~ 180 MeV/c travel in helical trajectories within the clear bore of the solenoid and decay to muons. Invariance of the quantity p_t^2/B and energy conservation imply that charged particles moving towards the closed end of the “bottle” are reflected downstream along the muon beamline if

$$\sin(\theta) > \sqrt{B_{\text{target}}/B_{\text{max}}},$$

where B_{max} and B_{target} are the values of the axial component of the magnetic field at the target and the upstream end, respectively, and θ is the angle of the pion with respect to the solenoid axis at the target. For the production region magnetic field values, the loss cone, where particles are not reflected, has a half angle of about 30° giving a solid angle acceptance for pion capture $\sim 93\%$. The transport solenoid filters the beam of unwanted particles and passes what remains to the stopping target.

To set the scale, a simple cylindrical target made of a high Z material such as

Production Target




Figure 7.2: Production target installed in the PS. The water-cooled design is shown here, viewed at the upstream end of the muon beamline. The entrance of the proton beam is the opening in the heat and radiation shield seen left of center and nearer the bottom of the figure. Target & service pipes are fully installed. The target assembly is made rigid with cross braces between the larger diameter service input and return pipes. The arrangement is supported by hanging in a dovetail slot cut out of the heat shield.

tungsten, 16 cm long and 4.0 mm in radius weighs about 155 g and receives an instantaneous on-peak power of 10 kW from the 50 kW proton beam expected for the first phase of this experiment. Production of some 10^{18} stopped muons in the planned running time requires 1.8×10^{13} protons/s on the production target. The target-heating is dictated by the macroscopic time structure of the beam: the accelerator cycle time and duty factor.

7.2 Muon Production

The basic principle of the production region is illustrated in a GEANT simulation, Fig. 7.3, showing a view of the production and beginning of the transport solenoids. Design of the production solenoid calls for a graded magnetic field with maximum intensity of the axial component to be 5 T at the upstream end, decreasing linearly to 2.5 T. The target is a long, pencil-sized cylinder and has its axis tilted away from that of the solenoid to intercept the proton beam. The direction of this beam is opposite that of the muon beam in order to reduce the fluence of low energy photons and neutrons into the muon channel and eliminate complications associated with

the interactions of the diffuse exiting proton beam and the beginning of the muon transport. A collimator of radius 15 cm, visible in the figure, defines the entrance to the transport solenoid.




Figure 7.3: GEANT simulation of Production Solenoid; cross-section, viewed from above. Note that coordinates are in centimeters. The proton beam (red line) enters from the right, strikes the target at approximately $(-650 \text{ cm}, 400 \text{ cm})$, and exits through a thin window at the lower left.

Tungsten is a suitable target material for initial studies comparing various cooling schemes; it has a high density, 19.3 g/cm^3 , good refractory properties, and a large pion production cross-section. Tungsten has the highest melting point, $\sim 3683 \text{ K}$, and thermal conductivity of all pure-metal refractories. In addition, this metal has good mechanical stability at high temperature, with yield strength about 5800 psi at 2000 K, and modest thermal expansion coefficient $6.01 \times 10^{-6} \text{ K}$ at 3000 K.

The calculated values of particle fluxes in the beam are based on GEANT simulations of proton interactions in a tungsten target. GEANT plus GHEISHA has been the primary combination of codes employed in our simulations. In order to reduce the uncertainty in the muon yield (and hence the sensitivity of the experiment) due to the uncertainty in the hadronic model of low energy hadron production, we have

Production Target

scaled the results from these simulations by a factor determined from comparison with measured π^- production cross sections in proton tantalum interactions. The effect of this scaling is to reduce yields by a factor of ~ 2 with respect to the GEANT + GHEISHA prediction. The backgrounds that depend directly on pion production rates have also been scaled in the same way.

The data to which we compare our simulation are from interactions of 10 GeV/c protons with tantalum, which is adjacent to tungsten in the periodic table. Measurements [97] include the invariant cross sections for π^- production as a function of pion kinetic energy T and production angle θ measured in the reaction $p^+\text{Ta} \rightarrow \pi^- + X$ over the full angular production range and for $p_\pi > 80$ MeV/c. One mm thick Ta plates with spacing of 93 mm were placed in a 2 m propane bubble chamber that was operated under a magnetic field of 1.5 T. Pion trajectories were confidently identified with minimum momentum of 80 MeV/c ($T = 21$ MeV). The measured average π^- multiplicity at 10 GeV/c is 1.51 ± 0.03 . The experimental π^- inclusive differential cross section measurements, together with a phenomenological fit, are shown in Fig. 7.4. The dependence of the invariant cross sections on the π^- kinetic energy is well approximated by an exponential function: $f = C \exp(-T/T_0)$. The total pion production cross section for Ta at 10 GeV/c found by integrating this formula with fitted values of C and T_0 is 2.36 barn. With a nuclear inelastic cross section for Ta of 1.56 barn, there is good agreement with the measured pion multiplicity of 1.51.

To compare the muon flux simulation with the experimental data, a Ta proton target ($\rho = 16.6$ g/cm³) with length 19.34 cm (1.67 nuclear interaction lengths) and radius 0.4 cm was studied. The muon flux was simulated using GEANT + GHEISHA. A proton beam with Gaussian shape and $\sigma_x = \sigma_y = 0.2$ cm was caused to impinge on the target. The proton interaction point was chosen using GEANT, and then π^- were produced at that point according to the measured production cross sections. The π^- were then tracked using the GEANT and the resulting μ^- yield calculated. The ratio of this μ^- yield based on measured production cross sections to that based on GHEISHA is 0.54. We scale all results that depend on pion yield at 8 GeV/c down by a factor of 2 to account for this difference and the energy dependence of the production cross section, the latter taken from a GEANT calculation.

As previously mentioned, and here in more detail, the yield of muons depends on the target shape, the proton energy, the value of the field in the production and transport solenoids, the clear bore of the production solenoid, and the size of the collimators. The yield was optimized [98, 99] with respect to variations in these parameters. It is relatively insensitive to small variations in target position and to the target length in the range of 12–20 cm. The yield decreases by about a factor of two with target radius variations between 3 mm to 9 mm due to absorption as the pion exits the target and as it passes through it again while moving in a helical trajectory in the production solenoid. We currently use a target radius of 3 mm. This has some implications on target heating, as discussed below. In addition to being necessary because of the incoming proton beam angle, the target tilt also helps




Figure 7.4: Negative pion inclusive differential cross section in different angle intervals for 10 GeV/c protons incident on tantalum [97]. The lines are the result of a fit to an exponential form: $f = C \exp -T/T_0$.

Production Target




Figure 7.5: Distribution of positive (top) and negative (bottom) muon creation points in the Production Solenoid. Most positive muons are generated in the walls of the solenoid, surface muons, whereas negatives are largely from negative pion decay in flight.

reduce scatter of pions following a helical trajectory. For a 5 T maximum field in the production solenoid and a 15 cm radius collimator, the stopped muon yield decreases by only $\sim 3\%$ in going from a 30 cm to 20 cm radius clear bore. This region may thus be available should more shielding be necessary. To reduce pion loss, the target support structure should have the lowest mass possible.

7.3 Stopped Muon Yield

A study of how stopped muon yield depends on radius reveals that the optimum target radius is 3.0 mm as shown in Fig. 7.6. Studies of stopped muon yield versus target length indicate no appreciable gain in yield is achieved beyond 16 cm. Further modifications of target geometry will result from a closer study of the relationship between target pressure drop and flow rate.




Figure 7.6: Stopped muon yield per primary proton versus target radius for a water-cooled target. Yield is reported relative to that for radiation cooling; from Figure 5 of [100].

Table 7.1: Optimization of stopped muon yield as a function of coolant containment shell thickness and gap size.

Water thickness (mm)	Ti wall thickness (mm)	μ^- stops per proton	Acceptance loss (%)
0.00	0.00	0.0050	0.0
0.20	0.15	0.0049	2.7
0.25	0.15	0.0048	4.1
0.30	0.15	0.0048	4.5
0.40	0.15	0.0047	5.8
0.50	0.15	0.0047	6.3
0.25	0.20	0.0048	4.5
0.25	0.30	0.0047	6.7
0.25	0.40	0.0047	6.0
0.25	0.50	0.0047	5.4
0.50	0.30	0.0041	17.8
0.50	0.50	0.0048	4.6
2.35	0.76	0.0037	27.0

7.4 Target Heating

The use of a heavy target in a very intense proton beam requires careful consideration of target heating. The power deposited in the target is not very sensitive to the hadronic code used, and is determined from a GEANT + GHEISHA simulation. The calculated average energy loss per 8 GeV/c proton is equal to ~ 0.7 GeV/proton and ~ 0.8 GeV/proton for target lengths 16 cm and 20 cm, respectively, equivalent to a peak power of 9.4 kW and 10.2 kW. The longitudinal distribution of the average energy loss per primary proton is shown in Fig. 7.7.

Table 7.1 was used to guide design calculations for cooling. Based on later studies comparing the effect of a water-cooling system on stopped muon yield to that of radiation-cooled, see Fig. 7.6 and [100], target radius was optimized to 3.0 mm. Subsequent to this, energy deposition was re-evaluated, leading to a total instantaneous power of 7550 Watts; Fig. 7.7 shows the results of these calculations in more detail.

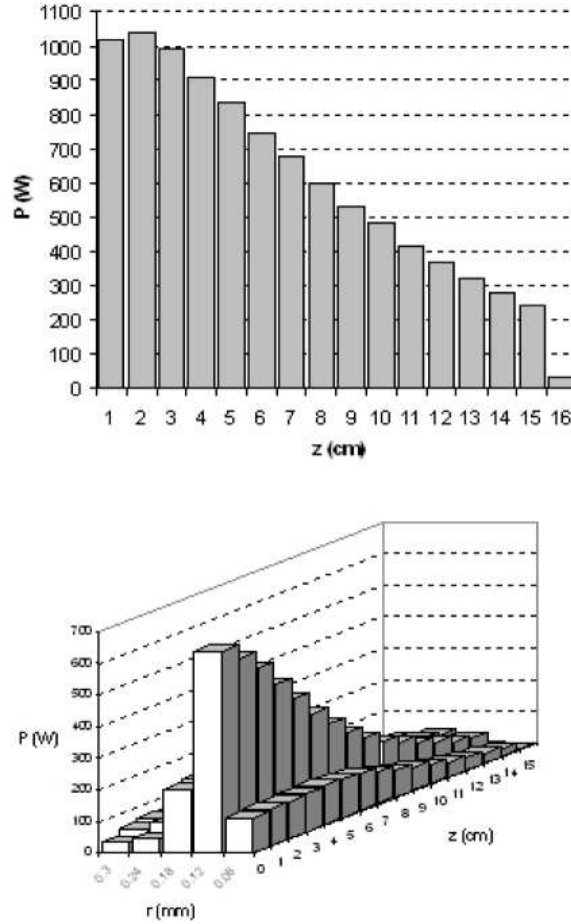


Figure 7.7: Pion production target power distributions with 3.0 mm radius and 16 cm length cylinder. These are the distributions used in cooling calculations, scaled to 9500 Watts total instantaneous power; the actual GEANT value is 7550 Watts.

7.5 Target Cooling

Two cooling strategies have received significant study: *radiation* and *forced convection* with water as the coolant. Radiation cooling is attractive because it provides a high muon yield, is simple to operate and virtually maintenance-free. On the other hand, high operating temperatures lead to evaporation of some of the target material; requiring strategies to control contamination of the rest of the muon beam line and detector area, and require good control of thermal stresses. Convective cooling makes drastically lower operating temperatures possible, extending the possibility of achieving even greater muon beam intensities. An essential requirement for such a system is to maintain a non-boiling condition, since gas formation would impede flow through small water channels. This type of cooling provides more latitude in selecting target materials with higher densities and thermal conductivities, thus raising the possibility of reduced thermal stresses and an increased useful lifetime of the target. However, this cooling system requires greater maintenance, and significantly more system design consideration. Protocols for coolant storage and disposal will be required.

A water-cooled target has been designed using a turbulent water stream in an annular coolant channel surrounding a cylindrical target; see Fig. 7.8. The coolant circulates in a closed loop, passing through a recuperative heat exchanger. A conceptual design of the full layout is shown in Fig. 7.10. In this section, we will discuss the results of design studies and prototype testing.

The current design, see Fig. 7.8, calls for a gold or platinum rod, 16.0 cm long, with a 3.0 mm radius, and cooling channel gap size 0.3 mm. The coolant containment shell is made from a high-strength titanium alloy; it consists of five pieces: a cylindrical shell, two endcaps, and inlet & outlet pipes. The wall thickness of the cylindrical shell is 0.51 mm (0.020") and the pipes have a 3.175 mm (0.125") outer diameter, wall thickness 0.55 mm (0.021"), and 25 cm long. The target rod is held in position within the shell by three rounded feet at each end, positioned 120 degrees apart, each 0.3 mm tall and about 1.0 cm from the ends of the cooling channel.

Assuming an inlet water temperature 20 C and a volumetric flow rate of one gallon per minute, our calculations, shown in Fig. 7.10, demonstrate that the maximum target surface temperature (at the outlet) can be held at 71 C. Target surface temperature is given by the curve with $r = 3.0$ mm. The top curve, labeled $T_{critical}$, gives the boiling point assuming 25 cm inlet and outlet pipes. Flow and cooling tests, employing induction heating [101–103], have demonstrated that such a target system is feasible. Further tests are needed for full system development for the first phase of the experiment. Significantly higher heat loads, factors of 2-8, will require target system redesign but it is deemed feasible. Figure 7.10 is a conceptual design of the cooling system. The loop is designed to be a constant flow cooling system. There are three monitors for temperature; T_R (reservoir), T_{in} (inlet), and T_{out} (outlet), two for pressure, P_{in} , and P_{out} , and one volumetric flow sensor FM. The single-speed

Production Target




Figure 7.8: Cross-sectional view of current target cooling design. In our design, the beam strikes a gold target end-on from the left. The target shell, end caps, and inlet & outlet pipes are made of titanium. The target has a slight taper at the inlet which helps reduce the operating pressure; the coolant channel then narrows to 0.3 mm.




Figure 7.9: Target and coolant temperature at fixed radii as a function of position along the length. Our chief calculation tool has been CFDesign, a heat and mass transfer program designed for solving complex engineering problems. The results shown here are for a worst-case scenario, steady state heating, with power distribution shown in Fig. 7.7 and 9500 Watts total instantaneous power. Flow rate is 1 gallon per minute.

pump must provide a steady (non-pulsed) flow. There are many highly efficient, reliable, low impedance, compact recuperative heat exchangers commercially available. Coordinated system control is essential. Flow is regulated with two needle valves, NV_1 and NV_2 . The gate valves GV_1 through GV_4 , and DV are necessary for servicing. All valves should be watertight, durable (e.g., stainless steel), with position indicators, and manual & remote control. All valves, the pump, and heat exchanger cold side supply should be mechanically slaved together to ensure they are operated properly; more importantly, the entire system must have electronic control and monitoring interlocked with the proton beam. All sensors and hardware, except the titanium assembly, will be located far from radiation and magnetic fields; thus everything “above” GV_3 and GV_4 , including these valves, will reside away from the PS.

Construction and materials for coolant transport must be selected to maintain control over coolant leakage. Corrosion resistance, strength, radiation shielding must guide material selection. Figure 7.2 shows 0.5-inch titanium alloy service pipes connecting to the inlet and outlet pipes. These larger pipes extend out through the Production Solenoid vacuum bulkhead. A robotic arm will allow remote manipulation of the entire titanium assembly. The water pipes that connect to the target assembly and extend to the pump should be made of corrosion resistant high-strength stainless steel. The target circulation system distance is expected to be approximately 40 feet; thus, the entire system should utilize 80-90 feet of stainless steel pipe. Proper wall thickness should be evaluated when coolant chemistry (along with other factors) is better understood. In any case, during target installation and removal, the two titanium and two steel pipes must be capped until connection is made. All joints, possibly with the exception of the connection just mentioned, should be welded continuously to ensure proper seals.

The amount and types of radioactive materials and chemical compounds that are formed in the coolant affect the size and composition of the reservoir R; the reservoir may also require a level indicator, as a safety precaution. Laboratory safety guidelines will determine the number of allowable Curies at a specified distance from the tank, and thus its design. Shielding may be required. Filtration needs are determined by the same coolant contamination factors mentioned above. A more appropriate filter location may be necessary.

Direct target temperature monitoring is unnecessary and poses no additional risk in operation; it is also advisable due to possible risks to pion reabsorption. At constant flow, water pressure and temperature monitors will serve as sensitive indicators of the state of the target. For example, a gradual accumulation of material in the flow channels would also cause an instant response and similar rise in pressure on P_{in} . Any narrowing of the cooling channel would produce a greater coolant velocity, thus raising the target cooling efficiency. A rise in target surface temperature will result in an instant rise in T_{out} . A complete risk assessment of the cooling system will be required, and concerns addressed before a design is finalized.

Production Target

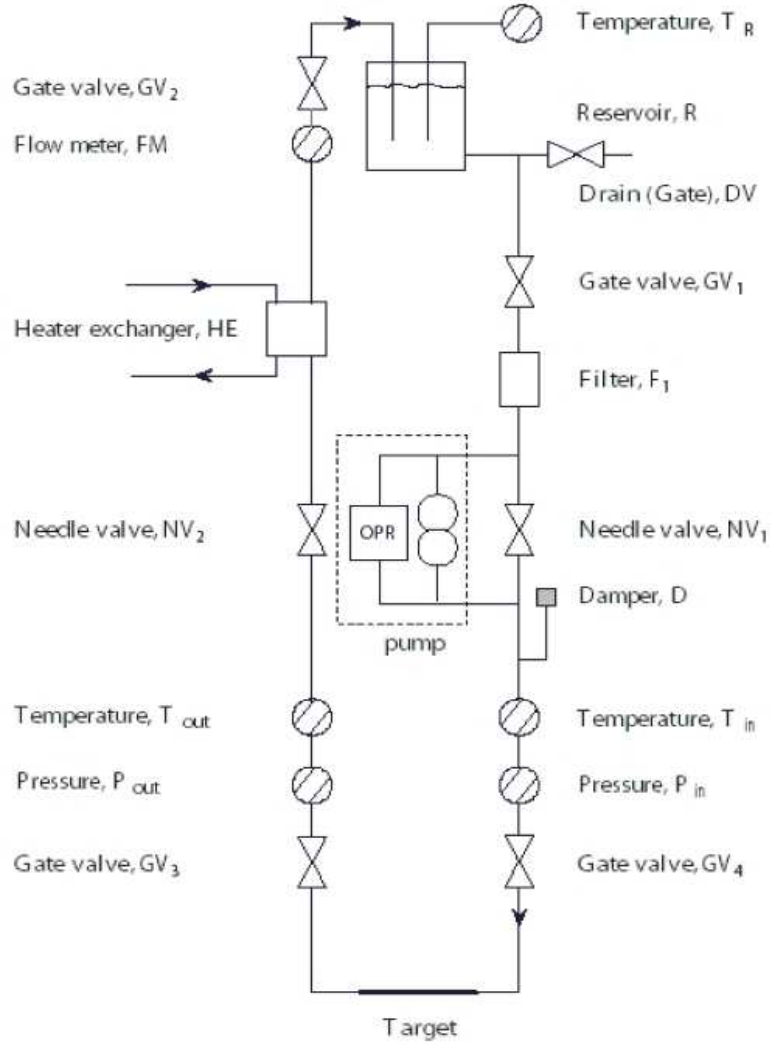


Figure 7.10: Conceptual design of the recirculating water-cooled target system. Note the over-pressure relief valve across the pump.

The reliability of the system shown in Fig. 7.10 has been extensively tested. We have built a mobile target test system, based on this design (with small practical modifications), which we have used to test full-scale target and cooling assembly prototypes, with and without heating. Target heating was done using induction heating with a 0.5 MHz 20-kWatt power supply.

Mechanical stresses in target assembly joints and components are very unlikely to reach dangerous levels. We recommend, based on our experience assembling numerous target prototypes, that joints be electron beam welded, to avoid damaging or distorting the flow channels. Compared to the importance of proper & continuous welding of all joints, individual component failure is remote. We have computed and verified the internal pressures likely to be present in the full-scale coolant recirculating

system; these results are discussed below.

Surface roughness changes as coolant impurities interact with the target surface. Impurities arise from target, coolant, and containment assembly materials undergoing nuclear fragmentation when they are in direct contact with the proton beam. This subject has not been fully explored yet.

The material that collects in the coolant and its concentration determine how it will be handled and stored. Tritium has a half-life of 12.3 years, making it one subject of possible concern. A cursory study of ^3H production using GEANT [104] showed that over the planned running time 4.72×10^{-6} moles of tritium stop in the coolant that lies within the heat shield. Less than 1% of all tritium is produced by beam interactions in the water itself. The primary origin of tritium is due to beam particles interacting with the target. Tritium that stops in the coolant comes, primarily, from beam-induced nuclear fragmentation at the target surface; in fact, other elements in the cooling water also originate this way. When tritium comes in contact with oxygen it will form tritiated water. Tritium can also form if suitably slow neutrons are present, by interacting directly with H or with ^3He . Leaving aside for now the neutron spectrum in the Production Solenoid region, the total tritium produced is small enough (based on [104]) that the tritiated water concentration will not reach a maximum. However, a mixture of tritium and water will result in some hydrogen gas production, as well as ^3He (decay product, but not a health hazard). There is also no danger present from hydrogen, since the concentration in air must rise above 4% before risk of explosion could arise. The low-energy electron, < 18.6 keV, from ^3H decay can be easily shielded. We plan to carry a complete study of coolant radiochemistry to understand what hazards exist and how to prevent them.

7.6 Heat and Radiation Shield

The heat and radiation shield protects the superconducting coils of the Production Solenoid from the ionizing radiation generated by the primary proton beam striking the production target within the warm bore of the PS. Details of the expected performance and activation can be found in Ref. [105]. Figure 7.11 shows a cutaway illustration of shield volumes and materials used in the physics simulation (GEANT3) to estimate the effect of variations in material choices and dimensions. In the figure, the muon beam upstream elements of the Transport Solenoid including the first collimator, shown in grey, can be seen in their positions immediately downstream of the heat shield.

The shield must limit the instantaneous local heating in the Production Solenoid coils due to radiation originating in the production target to less than 25 W/g at any point within the coil volume. The total steady-state power dissipation in the cold mass of the PS must be limited to less than 150 W. GEANT simulations [106] indicate that the inner radius of the shield must be greater than or equal to 25 cm to maintain the required muon stopping rate per incident primary beam proton. The

Production Target




Figure 7.11: Cutaway view of the heat and radiation shield within the warm bore of the Production Solenoid. In the figure, copper is blue and tungsten is red, while the stainless steel volumes are in several colors to distinguish regions.

76 metric ton shield must be actively cooled to remove approximately 16 kW of heat generated within its material volume by radiation from the production target. The vacuum boundary for the production region is at the PS inner cryostat wall, which places all elements of the shield within the production region vacuum. By design the shield receives extremely high doses of radiation in order to prevent the PS coils from bearing that burden. Therefore, any materials used in the shield construction must be able to survive large (to be determined) radiation doses without degrading to the point of failure. This includes any vacuum or fluid gasket materials or special low-friction bearing surface materials. The assembled shield shall have no line-of-sight cracks from any point in the production target to any point in the PS cold mass.

With the nominal 4×10^{13} protons/s beam on target, approximately 16 kW of heat is generated within the volumes of the heat shield by the passage of radiation. Nearly half (~ 6.2 kW) is generated within the PHW3 volume, i.e., the tungsten insert at the smallest radius near the production target; see Fig. 7.11. Table 7.2 lists the steady state radiation heat load in each section of the shield. In the table, the volume PCIW represents the inner wall of the PS cryostat as defined in the MIT Conceptual Design

Mu2e Proposal

Report, while the PSTB volume (as defined above) represents the additional material needed for strengthening the wall to support the heat shield.

Table 7.2: Estimated steady state radiation heat loads

Volume	Power (W)
PHC1	1200
PHC3	720
PHC4	100
PHW1	2700
PHW2	1500
PHW3	6200
P3CL	2000
PCLG	300
PCIW	41
Total:	15,741

The inner shell of the PS cryostat is a large deep cylinder that can support the weight of the shield during insertion and operation provided the loads are distributed properly to the shell, and the shell is held round at several points along the length. To achieve this goal a series of annular ring stiffeners must be incorporated into the shield's structure. The exact number of annular plates remains to be determined. The concept includes a radiation damage resistant bearing material (possibly Dicronite) on the outer edge of each ring to facilitate installation.

Muon Beamline

IN this chapter, we discuss the muon beamline; the goal is to produce 10^{11} stopped muons per second in the secondary target located in the Detector Solenoid (DS). This beamline includes the collimators and anti-proton stopping window in the Transport Solenoid (TS), proton and neutron absorbers, beam stop, and vacuum system. The TS system filters the particle flux (see Fig. 8.2) producing a momentum- (< 0.08 GeV/c) and charge-selected muon beam, with good reduction in contamination from e^\pm , μ^\pm , π^\pm , p , and \bar{p} during the detector live-time. The muons have high efficiency for stopping in the muon stopping target, and the electron momentum spectrum from muon decay in flight cuts off well below the electron momentum from μ - e conversion, 105.0 MeV/c in Aluminum and 104.3 MeV/c in Titanium. Muons not stopped in the target are transported to the muon beam-stop. Protons and neutrons originating from muon capture in the stopping target are attenuated by absorbers to minimize detector background rates. Finally, vacuum requirements are discussed.




Figure 8.1: Particle arrival time at the muon stopping target.




Figure 8.2: Particle momentum at entrance to the Detector Solenoid.

8.1 Collimators

The design of the transport solenoid is discussed at length in Chapter 9. The purpose of the collimators is to exploit the slow drift, perpendicular to the plane of the S-shaped TS magnet, in opposite directions by positive and negative charges. The vertical displacement midway through the S-shape TS magnet is:

$$D[m] = -\frac{Q}{e} \frac{\pi}{0.6B[T]} \frac{P_L^2 + 0.5P_T^2}{P_L[GeV/c]}, \quad (8.1)$$

where e is the charge of the electron, and P_L (P_T) is the component of the momentum along (transverse) the magnetic field. Besides this slow drift, the particles execute fast gyrations with radius:

$$a[m] = \frac{P_T[GeV/c]}{0.3B[T]}. \quad (8.2)$$

As can be seen from the above equations, the collimators filter the beam favoring low-momentum particles. The anti-protons are highly non-relativistic, and stop in the thin anti-proton stopping window. Almost all the electrons are relativistic and largely pass before the start of the detector live-time, which begins about 700 ns after the end of the proton bunch hits the target. The pions largely decay before the start of the detector live-time. All positively charged particles and high momentum particles are highly attenuated by the collimators.

Cylindrical collimators are placed in the first and last straight sections of the TS. The first collimator is made of copper. The last is made of polyethylene (CH_2) loaded with either Boron or Lithium to protect the detectors from neutron backgrounds. The important dimensions are given in Table 8.1. The mid-TS collimator is designed to accept only negative muons. Thin annular foils with inner radius 8.5 cm are inserted in collimator 3 to stop low energy muons which would miss the muon stopping target. Shielding in the S-shaped solenoid to protect the TS coils from radiation produced at the proton target. The first two collimators with the TS shielding are shown in Fig. 8.3.

Table 8.1: Parameters of three TS collimators. Position of the center is in the standard coordinate system with global origin at the center of the muon beamline.

Collimator	Center			Inner	Length [m]
	position [cm]			radius [cm]	
1	390.4,	0.0,	-345.4	15–17	1.0
2	0.0,	−9.0,	0.0	15	0.8
3	−390.4,	0.0,	343.0	12.8	1.0




Figure 8.3: The first and second collimators (H and G) and the shielding to protect the TS coils (S).

Collimator sizes were optimized to remove or heavily suppress electrons above 100 MeV, a potential source of background. GEANT simulation studies using the collimators described above in a full simulation of the experiment with 10^7 primary protons on the production target completely eliminated 100 MeV electrons at the exit to the TS. Furthermore, by design of the graded magnetic field, 100 MeV electrons originating in the production solenoid will have insufficient P_T to hit the detectors.

8.2 Absorbers

The absorbers in the Detector Solenoid area are designed to moderate protons and neutrons. Figure 8.4 shows a schematic view of the DS magnet, revealing the boron or lithium loaded polyethylene proton absorber and neutron absorbers inside the DS vacuum. The largest potential contribution to the tracking detector rate is from protons from muon capture in the stopping target. Without shielding, the average rate in individual tracking detector elements is above 10^6 Hz. However, the average rate from protons is reduced to 10^5 Hz with the proton shield [107]. The interface between the neutron absorbing shielding material and the DS is as follows:

- Individual pieces of the shielding material will rest against the walls of the DS warm bore. These pieces have an axially varying weight load as high as 23.6 kg/cm at the upstream end to 6.34 kg/cm at the downstream end of the magnet.
- The shielding material will bolt to the threaded bars welded to the warm bore walls.




Figure 8.4: Schematic drawing of the Proton and Neutron Shields in the Detector Solenoid. DAB1 represents the proton shield. DAB2-4 represent different sections of the neutron shield.

Neutrons produced in the muon stopping target are a potential source of background in the cosmic ray veto counters that surround the Detector Solenoid. The MARS [108] Monte-Carlo program was used to simulate these neutrons. Although these studies are continuing, acceptable fluxes were obtained with boron loaded polyethylene. Lithium loaded polyethylene was also studied, but it is more expensive and did not significantly reduce the flux in the detectors.

8.3 Muon Beam Stop

The muon beam stop (MBS) is designed to absorb the energy of beam particles, which consist mainly of e^- and μ^- (see Table 8.2) that reach the end of the detector solenoid, while minimizing the background to the detectors from the muon decays and captures in the beam stop, especially during the detector live-time, which begins about 700 ns after the proton bunch hits the target [109]. Of the beam particles that arrive at the stopping target region, 50% of muons and 84% of electrons continue towards the dump. Bremsstrahlung photons produced by electrons, and the products of muon decay and capture in the beam stop can hit the detectors if the beam stop is placed too close. Near the rear of the detector solenoid the axial magnetic field intensity drops along the beam stop direction; this is a critical feature of the beam

stop. The field gradient reflects most charged particles produced in the stop away from the detectors.

Table 8.2: Particles flux per beam proton at the Detector Solenoid entrance from FLUKA simulations. Late particles have an arrival time after 600 ns.

All μ^-	All μ^+	All e^-	Late e^-	All e^+	Late e^+	All π^-
0.018	8.2×10^{-6}	0.23	4.6×10^{-4}	0.032	7.4×10^{-5}	2.7×10^{-6}

During the beam stop optimization we satisfied the following requirements: the beam stop should be far from the detectors and the solid angle from the beam stop to the detectors should be minimized to reduce the flux of secondary particles leaving the beam stop in the direction of the tracker and calorimeter. The inner radius of the beam stop cylinder walls can not be too small, otherwise the muons and decay in flight electrons start to hit the walls earlier in the region close to the calorimeter. This defines a requirement that the first part of the beam stop should have the inner radius only somewhat smaller than the minimal distance from the calorimeter to the axis. The second part of the beam stop should have a small radius, but at the same time allow all electrons and muons to propagate to the downstream face of the beam stop. We have chosen the radius of the second part of the beam stop $R = 32$ cm and the position of the downstream face of the beam stop $z = 1700$ cm, in the global coordinate system. One can see in Fig. 8.5 and Fig. 8.7 that in this case all the late electrons (arrival time after 600 ns) and muons hit the walls of the beam stop in the region $z > 1600$ cm. In this region the field is less than 0.05 T (5% of the 1.0 T magnetic field in the tracker region). Almost all the beam particles hit the downstream face of the beam stop.

To make a choice of the material to be used on the inner surface of the beam stop, the following considerations were taken into account. There are some benefits to high atomic number on the inner surface of the beam stop, as almost all muons are quickly captured before the detector live-time. The muon lifetime in the Lead, for example, is only 80 ns. Unfortunately, the neutrons from the muon capture cause unacceptable radiation damage of the APDs in the calorimeter. The calculated neutron flux in the calorimeter region from the muon capture in the beam stop is shown in Table 8.3. The neutron flux in the calorimeter from the CH₂ beam stop is less than the neutron flux from muon capture in the aluminum stopping target, 3×10^{10} neutrons/cm² over the life of the experiment, whereas the Lead beam stop gives considerably more neutron flux in the calorimeter. The need to limit the APD radiation damage from neutrons led to our choice of CH₂ for the inner surface of the beam stop.

Finally the beam stop has a 2 cm thick stainless steel cylinder which covers the whole region between the calorimeter and the front face of vacuum closure cap (see Fig. 8.7). From the inner surface it is covered by CH₂, with 2 cm thickness on the first 75 cm along the z axis and 9 cm thickness on the remaining 268 cm. To decrease




Figure 8.5: Trajectories of beam electrons with late arrival time (after 600 ns). In this plot, the electron trajectories start at the end of the TS. Red indicates electrons, blue gammas.




Figure 8.6: Trajectories of beam muons (green) and electrons (red) from muon decays. Blue indicates gammas, dashed black lines indicates neutrons. CH_2 is indicated by yellow lines.

the flux of bremsstrahlung photons at the vacuum pump locations, the beam stop stainless steel cylinder is covered from outside by a 2 cm thick Lead layer, which is 184 cm in length. The cosmic ray shield inside the steel cylinder is covered by a 10 cm thick CH_2 layer with a 10 cm diameter hole in the center for the muon stopping target monitor pipe.

The chosen design of the beam stop not only largely suppresses the contribution to the rates from secondary interactions in the beam stop region, but also leaves enough

Table 8.3: Peak neutron flux in the calorimeter from muon capture in the beam stop during the life of the experiment (10^7 s).

Inner surface material	Peak neutron flux [n/cm ²]
Pb	4×10^{11}
CH ₂	6×10^9




Figure 8.7: Beam stop design showing CH₂ (yellow), stainless steel (red), Lead (blue), and the calorimeter (black cross-hatched).

room for cables in the vacuum volume outside of the beam stop.

8.4 Vacuum System

Vacuum is required in the detector solenoid mainly to limit backgrounds from muons stopping on gas atoms, followed by either muon decay or capture. The DIO endpoint (also the $\mu - e$ conversion electron energy) is given in Fig. 8.8 for different elements [56]. As discussed earlier, the DIO process is a steeply falling background with an endpoint energy given when the two neutrinos have zero energy. The muon mass is $105.66 \text{ MeV}/c^2$. Nuclear recoil reduces the energy for light nuclei and the Coulomb binding energy reduces it for heavy nuclei. The end point for aluminum is 104.97 MeV . Oxygen, present in the residual gas from small air leaks, and fluorine, present in the residual gas from tracker gas CF_4 leaks, have the highest DIO endpoint, 105.1 MeV . The difference between these endpoints and the aluminum endpoint, 0.14 MeV , is not

significant compared to the detector resolution. Therefore, for the aluminum target, the gas does not introduce a “physics” background, and only contributes to the overall detector rates. These issues still have to be studied for a Titanium target, including low Z impurities in the target itself. Our estimate shows a detector solenoid vacuum pressure of 10^{-3} Torr should be adequate for the aluminum target [110].




Figure 8.8: Muon decay-in-orbit endpoint energy vs atomic number.

The ability to withstand high voltage discharge in a gas is a function of the pressure. The minimum voltage difference without gaseous discharge for Nitrogen is only 275 V. This occurs when the distance d and pressure p satisfies $pd = 0.75$ Torr-cm. This is called the Townsend limit, and occurs when the energy acquired between gaseous collisions is greater than the ionization potential, and the number of collisions is sufficient to initiate a gaseous discharge. At lower voltages, a spark will not occur, but a glow discharge will develop. For our geometry, a gas pressure of 10^{-3} Torr is adequate for the detector high voltage [111].

Scattering of beam particles into the detector acceptance in 10^{-3} Torr vacuum was evaluated and found to contribute negligibly to the detector rates. The multiple scattering and energy loss of the conversion electrons contributes negligibly to the signal resolution at 10^{-3} Torr.

The Detector Solenoid requires a closure at its downstream end to:

- Close off the vacuum space of the magnet warm bore.

Muon Beamline

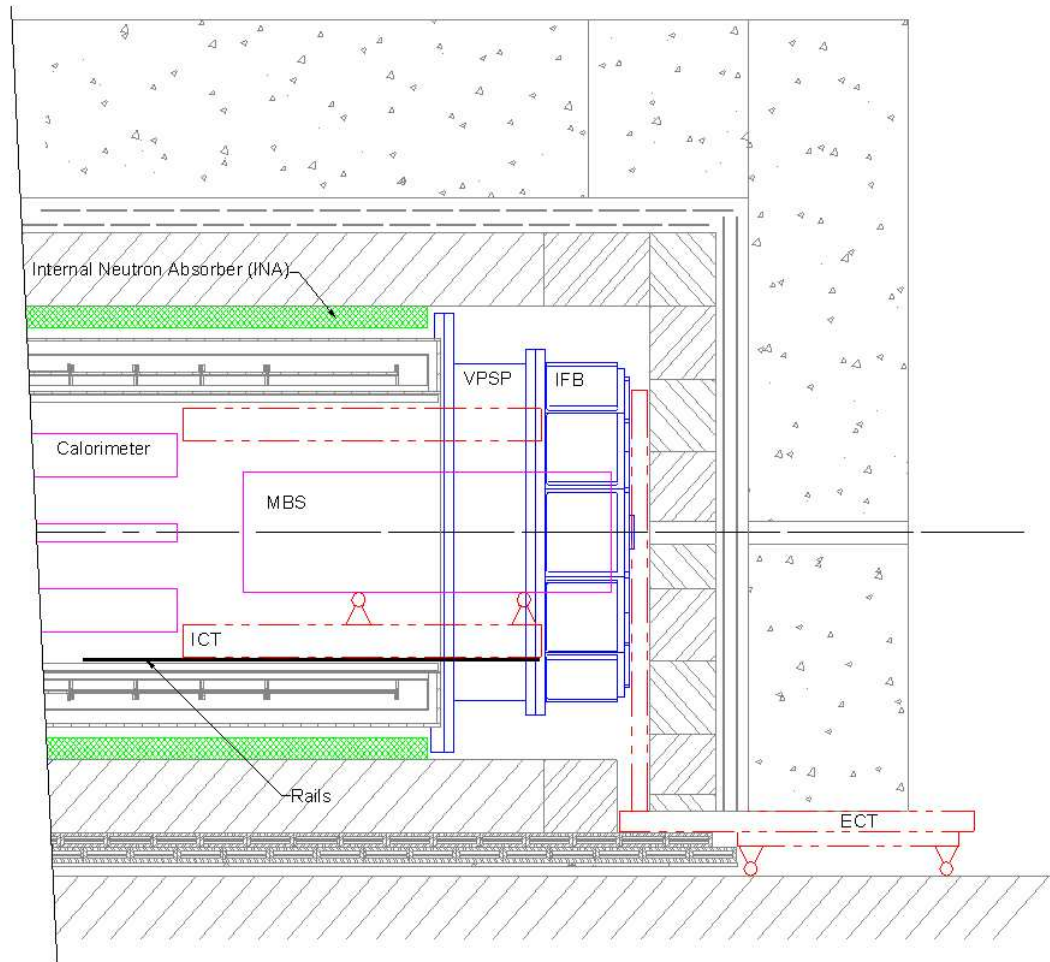


Figure 8.9: DS Vacuum Closure Elements. Note the Internal Neutron Absorber (INA), in green, that lies between the outer cryostat wall and the inner wall of the iron magnetic return yoke. It adds further protection to the active cosmic ray shield from slow neutrons from the DS. (The muon beam stop in this drawing is from an earlier design.)

- Provide a connection and mounting place for the pumps used to generate and maintain the warm bore vacuum.
- Provide instrument, power, gas and cooling feed-throughs to bring these lines through the vacuum barrier.
- Provide for the orderly routing of cables.
- Provide partial mounting and support for the Muon Beam Stop (MBS).

As the design of the vacuum closure evolved, it became apparent that the best way to satisfy the design requirements was to do so with separate elements which worked together. The vacuum closure is made up of the following four elements:




Figure 8.10: showing the three cryo vacuum pumps, one turbo pump, and four spare ports.

1. Vacuum Pump Spool Piece (VPSP)
2. Instrumentation Feed-thru Bulkhead (IFB)
3. Internal Cable Tray & MBS support (ICT)
4. External Cable Tray Cart (ECT)

This is illustrated in Fig. 8.9. The vacuum pumps mounted on the VPSP are shown in Fig. 8.10. The three CF-10 cryopumps [111] have a total effective pumping speed of 6×10^3 l/s for air, 4×10^3 l/s for CF_4 , and 5×10^3 l/s for C_4H_{10} . The latter two are suggested tracker gases. With three cryopumps, we can isolate one for maintenance or regeneration while maintaining the DS vacuum with the other two cryopumps. The pump-down of the 2.5×10^4 l DS volume begins with a SDV-320 screw dry vacuum pump located outside the shielding. When the pressure drops below about 0.5 Torr, the gate valve to the hybrid SM141001 molecular-drag/turbo-molecular pump is opened. When the pressure drops below about 0.01 Torr, the gate valves to the cryopumps are opened. The turbo-pump must be shut down before the magnet is energized, as the fringe magnetic fields are above the pump specifications. The DS pressure vs time during this initial pump-down phase is shown in Fig. 8.11. The magnetic field specification for the cryopump is $B < 0.12$ T. The fringe field at the location of the cryopump is 0.11 T. Studies showed that with simple magnetic shielding of three 5 cm thick iron plates, the fringe field at the cryopump could be reduced to 0.5 T with an acceptable influence on the DS main magnetic field [111]. The physics/engineering optimization will be performed when funds are available.

The main radiation issue with cryo pumps is the Teflon coating on the cryopump displacers. Teflon degrades after a radiation dose of only 100 Gy. The vendor test data of the proposed rad hardened cryopumps showed no degradation in performance after 5×10^5 Gy. The calculated dose at the cryo pump location with the new MBS design [109] was 1000 Gy over the lifetime of the experiment (10^7 s), giving many orders of magnitude safety factor.




Figure 8.11: DS pressure during the initial pump-down. The cryopump gate valves are opened when the pressure is below 0.01 Torr.




Figure 8.12: Calculation of the DS pressure vs time after pump-down begins for the out-gassing vacuum load only.

The main gas loads are due to out-gassing of the shielding and beam stop materials, and leaks and gas diffusion from the tracker gas system. A calculation of the DS pressure vs time after pump-down begins for the out-gassing vacuum load [111], without bake-out, is shown in Fig. 8.12. Our design goal is 10^{-4} Torr pressure after one day of pump down, so the present DS out-gassing load is acceptable without bakeout. This will be re-evaluated as the engineering design progresses. The straw leak rate extrapolating from the leak rates measured from single straws [112] is less than 0.1 Torr l/s. With two cryopumps, this gives a steady state DS vacuum pressure of 3×10^{-5} Torr. Although the physics requirement is 10^{-3} Torr, we want to measure our detector rates, etc. as a function of gas pressure down to 10^{-4} Torr to verify our calculations, so the system design is adequate even if the tracker system gas load is three times greater than our extrapolation using the single straw leak rate measurements. Our vacuum system design will be re-evaluated after the tracker system gas load has been measured.

The thin anti-proton stopping window, Fig. 8.13, also serves to separate the DS and PS vacuum spaces. This is necessary to prevent neutral radioactive atoms from diffusing into the detector region. The PS vacuum system has not been designed yet; however, it should be simpler due to the much smaller vacuum volume without detectors within it.

The pump-down procedure is first the PS-DS bypass valve will be opened and both the PS and DS will be pumped down. Once the pressure is well below 1 psia (50 Torr), the bypass valve will be closed. The vacuum control system will open the bypass valve if the pressure differential is greater 1 psia. After the first run,




Figure 8.13: Schematic of the anti-proton stopping window.

the window will be removed and tested to destruction to verify our radiation dose degradation calculations, etc.

Figure 8.14 shows the vacuum instrumentation feed-through bulkhead (IFB) just downstream of the vacuum pumps. The IFB provides the DS vacuum closure and will be withdrawn together with the detector for maintenance between runs. The IFB provides the downstream termination of the internal cable tray. This design features modular, removable instrumentation feed-through panels. There are twelve facet covers, as seen in Fig 8.15, each with net area of 1935 cm^2 , and twelve wedge covers, as seen in Figure 1.20, each with an area of 1255 cm^2 . The calorimeter feed-through pin counts are given in Table 8.4. The last column shows the required IFB area. The calorimeter requires the largest number of vacuum feed-throughs, and uses about one half the IFB available area. The tracker feed-through requirements are shown in Table 8.5. The magnetic measurement system has not been specified yet.

The conceptual design, Fig. 8.16, for the PS vacuum closure and pump system has yet to be fully developed. However, preliminary calculations show that a smaller screw dry vacuum pump, a hybrid SM141001, and two CF-10 cryopumps would be adequate.

Muon Beamline

Table 8.4: Calorimeter IFB cabling.

Crystal count	1800	Area (cm ²)
Signal pins	3600	1162-3175
HV pins	3600	7200
Preamp power	216	300
Gating and calibration	144	4363
Gas lines	8	800
Cooling	8	800
Thermosensors	576	397

Table 8.5: Tracker IFB cabling.

Data, clock, trigger, sync	4×18 optical channels
Slow controls	6×18 copper line pairs
Power	2 copper line pairs
HV	18 copper line pairs
Gas	18 lines




Figure 8.14: Vacuum closure two-piece configuration showing the VPSP with vacuum pumps and the IFB. The VPSP will remain attached to the solenoid while the IFB will be withdrawn during internal detector servicing.




Figure 8.15: Perspective drawing of the IFB showing the facet covers and the wedge covers.

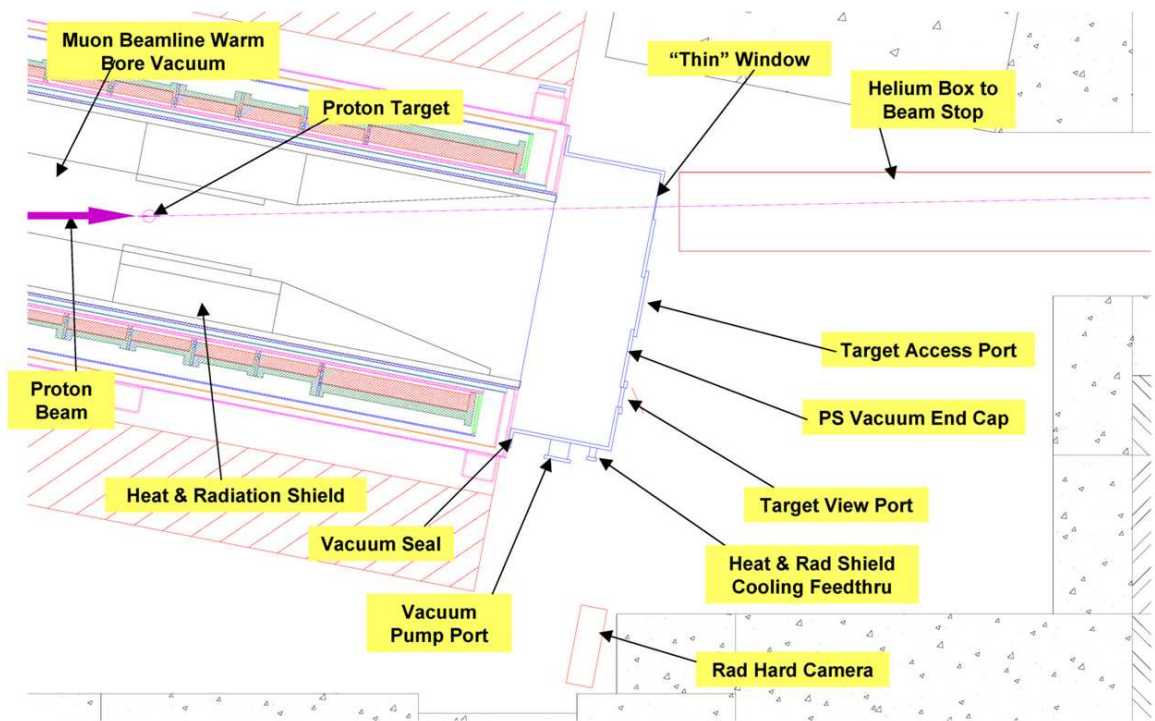


Figure 8.16: Conceptual design of production solenoid vacuum closure and proton beam exit window with helium box leading to the proton beam stop.

Production, Transport, and Detector Solenoids

THE solenoid magnet design for the *Mu2e* experiment is based on the extensive design and engineering work performed as part of the proposed MECO experiment [113]. The work included defining magnetic field specifications, coil sizing and placement calculations, cryogenic and quench analysis, and mechanical engineering of the support structure [114–121] leading up to a thoroughly documented magnet conceptual design report [122]. Relatively minor changes were made to the magnet design between the publication of the conceptual design report in June 2002 and the baseline design presented at the RSVP review in April 2005 [123]. In this chapter an overview of the magnet design is presented.

The magnet system consists of three superconducting solenoidal magnets with iron shielding and iron endcaps. The main units are Production Solenoid (PS), the Transport Solenoid (TS), and the Detector Solenoid (DS) as shown in Fig. 9.1. The entire magnet system consists of the coils, iron shielding, mechanical support structures, cryostats, power supplies and quench protection, and vacuum systems. A view of the assembled magnet is shown in Fig. 9.2

A high energy proton beam is directed at a solid target near the center of the PS. The protons enter through an opening between the PS and TS coils as shown in Fig. 9.3. The axial magnetic field of the 4 m long 1.5 m diameter PS is graded with the downstream (of the protons) end at 5.0 Tesla and the upstream end near the TS at 2.5 Tesla. As explained in Chapter 7 the pions and muons produced from the target are contained and directed toward the TS by the axially graded magnetic field of the PS.

A fraction of the muons resulting from pion decay are captured in the PS and directed toward a 0.5 m bore diameter Transport Solenoid (TS), which is S-shaped with an axial length of about 13 m. Since the TS is a toroid curving in the horizontal plane, particles of opposite charge drift in opposite directions vertically, with magnitudes which depend on their longitudinal and transverse momentum components. Collimators at the center of the TS (just after the first 90 deg bend) are used for particle charge and momentum selection. The second half of the TS transports the remaining particles to the DS and centers the muon beam at the entrance of the DS.




Figure 9.1: Sketch of the Mu2e magnetic elements showing the layout of the coils and the iron shielding.




Figure 9.2: View of the assembled magnet system. (Taken from Ref. [122].)

The magnetic field in the TS decreases from 2.5 Tesla at the PS end down to 2.0 Tesla at the DS end.

The muons entering the DS are brought to rest in a stopping target inside the 10 m long by 1.9 m bore Detector Solenoid (DS). The magnetic field on the axis of the system is carefully designed to decrease from 2.0 Tesla at the entrance of the detector solenoid to 1.0 Tesla before the start of the detector. In the region of the detector

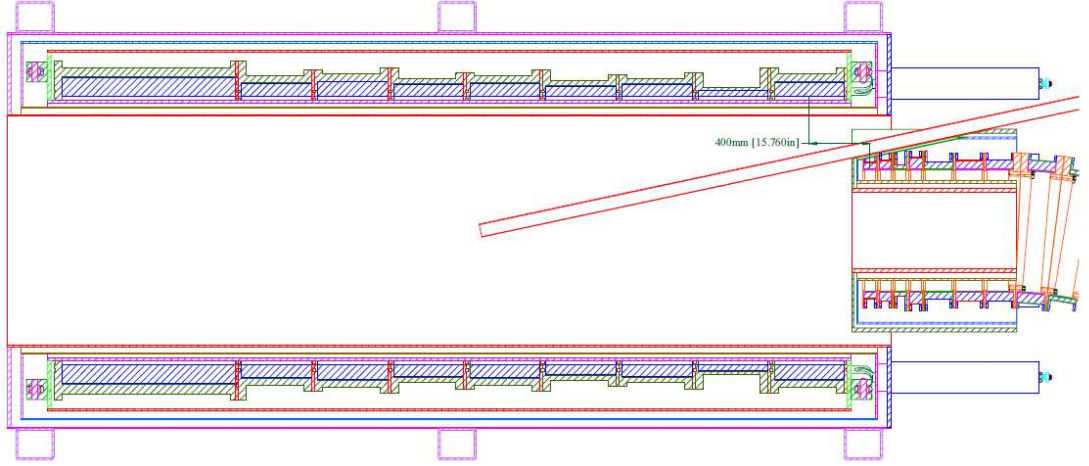


Figure 9.3: Proton beam interface. Proton beam enters the Mu2e apparatus between the coils of the PS and the TS.

the magnetic field is designed to remain constant at 1.0 Tesla with a high degree of uniformity to provide the required resolution in the tracker. The muon stopping target is placed in the region of decreasing magnetic field to improve the acceptance of decay electrons as explained in Chapter 7.

A massive iron shield surrounds the DS to absorb cosmic ray particles that are a source of background in the experiment. There is also a massive shield around the PS to provide radiation shielding of the environment from the production target. Endcaps are added to the ends of the PS and DS shielding to reduce the PS and DS fringe field within the volume of the TS.

The solenoid magnetics are designed to capture muons and transport them to the stopping target and prevent particles from being trapped in local field minima. These requirements are formed into a set of specifications on the magnetic field quality and limits on the magnetic field gradient as defined in [115] and listed in Table 9.1. The specifications is a set of rules specifying fields in a volume limited by a surface formed by a circle of a variable radius, R_{fs} , normal to the S-shaped axis of the magnet system. The global axial coordinate, z , is measured along the axis of the magnet system with $z = 0$ at the upstream end of the PS where the field is 5 Tesla. Field specifications apply limitations on deviations of the field from the field spec, as well as on the values of field derivatives, dB/ds . Here B is the axial component of the magnetic field, and s is measured along lines, equidistant from the axis of the magnet system in any fixed azimuthal axial plane.

The steel shield [124] is included in the magnetic analysis. The DS shielding is a 13.312 m long square box with a 0.5 m wall thickness, coaxial with the DS coil, and having an inner square aperture 3.1 m on a side. An endcap is placed on the DS shielding which is 0.3 m thick with a 0.65 m inner diameter octagon cut out of the center. The PS shielding is a 6.3 m long square box with a 0.75 m wall thickness,

Table 9.1: Field specification characteristics.

Coil	L (m)	Z_{min} (m)	Z_{max} (m)	R_{fs} (m)	$ \Delta B/B $ max	where	dB/ds (T/m)
PS0	4.00	0.00	4.00	0.30	0.050	$r = 0$	< 0
TS1	1.00	4.00	5.00	0.15	0.005	$r = 0$	< -0.02
TS2	4.60*	5.00	9.60	0.15	0.010	$r < R_{fs}$	na
TS3	1.95	9.60	11.55	0.19	na	na	< -0.02
TS4	4.60*	11.55	16.15	0.15	0.010	$r < R_{fs}$	na
TS5	1.00	16.15	17.15	0.15	0.005	$r = 0$	< -0.02
DS1	4.00	17.15	21.15	0.30-0.70	na	na	var $< 5\%$
DS2	4.00	21.15	25.15	0.70	0.002	$r < R_{fs}$	na
DS3	2.00	25.15	27.15	0.70	0.010	$r < R_{fs}$	na

coaxial with the PS coil, and having an inner square aperture of 3.2 m on a side. The steel is assumed to be 1010 carbon steel.

The coil design for the solenoids assume that the surplus SSC cable will be used in the construction. An iterative process is used to determine the coil parameters including the number of coils, placement of coils, number of turns in the coils and the current. After iterating, a set of 96 coil sizes and placements was determined [125] which met the field specs even when the effects of the iron shielding was included in 3D magnetic modeling codes such as Opera-3D. In some cases the practicality of the engineering design led to a slight relaxation of the field specifications but these did not affect the performance of the muon beamline.

The superconducting coil is wound as a set of 96 coil units, each wound on its own mandrels. A 1 mm thick ground insulation of fiberglass is wrapped over the mandrel OD, and a 1 mm thick ground insulation of G11CR is placed inside each of the mandrel side flanges. Coil winding starts, and ends at the OD of the winding. In order to achieve it winding begins at the ID of the winding, and goes in two directions. First, a single pancake along one of the sidewalls of the mandrel is wound. This pancake is isolated from the remainder of the winding by an insulating plane of 1 mm of G11CR. The other end of the conductor is used for layer winding the rest of the coil between the pancake winding, and the other sidewall of the mandrel. A 1 mm thick layer of insulation is introduced between the pancake, and the rest of the winding for voltage isolation.

The magnet system is contained in four cryostats: PS, TSu, TSd, and DS. A scale drawing of the magnet system is included as Fig. 9.4, and a labeled, 3-D rendering of the system is shown in Fig. 9.2. The two TS cryostats are divided by a thin window [126] which serves to separate the upstream and downstream bore vacuum systems. The upstream TS cryostat (TSu) contains the series-connected TS1 and TS3u (TS3 upstream) coil sets, together labeled TS13u, located in the straight

Production, Transport, and Detector Solenoids

sections and the TS2 coil set located in the 90 deg bend. Similarly, the downstream TS cryostat (TSd) contains the series-connected TS3d (TS3 downstream) and TS5 straight-section coil sets, together labeled TS3d5, and the TS4 coil set in the second 90 deg bend. The PS coils are subjected to a nuclear heat load that could approach 190 W at 4.5 K, depending on the final proton beam target shield design. The PS coils, therefore, are cooled with a combination of natural-convection and/or forced flow helium that is in direct contact with the ID and OD of the windings. The TS and DS coils are conduction cooled using thermal heat drains connected to helium tubing. The cryogenic system has been designed to facilitate the cooling of any cryostat while any of the others are either warm or cold.

The magnet system will operate with six power supplies, one for each of the six coil sets: PS, TS13u, TS2, TS3d5, TS4, and DS. The power supplies are designed to ramp the coils to full current in one hour. In this conceptual design, the supplies are also required to provide a reverse voltage capability, which enables a slow discharge of each of the six coil sets in under one hour. The parameters of the power supplies are listed in Table 9.2.

Table 9.2: Coil Set Inductance, Operating Current and Minimum Charging Rate

Coil Set	Inductance (H)	Operating current (A)	Minimum charging rate (A/s)	Minimum charging voltage (V)	Maximum time to zero current (A/s)
PS	17.73	3500	0.972	20.0	3100
TS13u	2.045	1500	0.417	1.0	650
TS2	0.523	4000	1.11	1.0	270
TS3d5	1.19	1500	0.417	1.0	400
TS4	0.386	4000	1.11	1.0	200
DS	3.78	4000	1.11	5.0	1450

Structural designs have been completed for all magnet systems and the reader is referred to the extensive documentation in references [122, 125] for more detail. The PS coils employ aluminum shells to contain the hoop and axial loads of the winding. The TS and DS coils are self-supporting to their Lorentz loads, but have internal mandrels to facilitate gravity and overall axial support. Axial and gravity loads are carried in all cryostats by spherical rod-end rods that serve as cold to warm supports and also accommodate the motions associated with cooldown. The PS and DS coils are designed to contract axially toward the TS, and the two TS cryostats are designed such that their coils contract axially toward the central Kapton window. Final adjustment of the rod nuts when the coils are cold, and following magnetic surveying during installation, will ensure that the coils are properly positioned for

Mu2e Proposal

operation. The magnet system has been designed structurally to enable any or all coil systems to be powered at once.




Figure 1.1 – The layout of the MECO Magnet system

Detector

THE *Mu2e* detector is shown in Fig. 10.1. It consists of the muon stopping target, the straw-tube tracker, and the electromagnetic calorimeter. The stopping target sits in the graded part of the Detector Solenoid field; the tracker and calorimeter are in the constant 1 T field.

The detector must have no elements in the region of the intense transported beam (which consists of 50 billion muons per second alone), a large acceptance for the ~ 100 MeV conversion electrons and small acceptance for the decay-in-orbit electrons coming from the stopping target. The dominant background are electrons coming from the high-energy tail of the decay-in-orbit distribution. Since the decay-in-orbit endpoint energy corresponds to the conversion energy this background can only be defeated by good energy resolution. To achieve our sensitivity goals that resolution must be better than a FWHM of 1 MeV. To minimize muon captures the Detector Solenoid is evacuated to less than 10^{-3} Torr. The detector must also have a minimal amount of material since multiple scattering is the limiting factor in the energy resolution.




Figure 10.1: The *Mu2e* detector in the Detector Solenoid. The muon beam stop is not shown.

In what follows we discuss the muon stopping target, the muon capture monitor, the tracking detector, and the electromagnetic calorimeter.

10.1 Muon Stopping Target

The baseline aluminum target has a mass of 159 g in 17×0.02 cm thick disks. The disks are arranged parallel to each other in 5 cm intervals, centered on the Solenoid Magnet axis and with each face perpendicular to it. Their radii range from 8.30 cm at the upstream end to 6.53 cm at the downstream end. The target is placed in the graded portion of the Detector Solenoid magnetic field, with the upstream disk at 1.57 T and the downstream disk at 1.30 T.

The stopping target is configured to optimize several competing factors:

- The target must be sufficiently thick in the direction of the muon beam trajectories to stop a large fraction of the incoming muons.
- The target must present the minimum possible path length to hypothetical conversion electrons (at least those which would fall within the acceptance of the detector) as they exit the target and as they spiral in the magnetic field. The width in the electron's energy loss distribution, due to variations in the amount of material traversed and due to energy straggling, are major contributors to the energy resolution of the conversion electrons. This implies that the target should be made as thin as possible.
- The radii of the target (e.g. extent of the target away from the solenoid axis) should be minimized to reduce the number of DIO electrons which can reach out to the inside radius of the tracker, and present background.
- The thickness of the target should be minimized in order to help control backgrounds from bremsstrahlung caused by beam electrons traversing the target, or delta rays produced in the target by energetic cosmic ray muons, or other cosmic ray interactions, etc.

Care must be taken in the design of the target support wires to be thin enough to minimize the number of stopped muons in them. If muons stop in the wires at a significant distance from the solenoid axis, then low energy DIO electrons could fall into the acceptance of the tracker and cause unwanted low energy background. An initial design will give an extra 10% background rate of DIO electrons, which is acceptable. In addition, since the DIO endpoint energy, which is the same as the conversion energy, decreases with increasing Z , as shown in Fig. 10.3, a high- Z target cannot have low- Z supports.

Target configurations similar to aluminum would also be used for titanium and other elements which are under consideration as target material.

The target material selection is similarly based on competing factors:




Figure 10.2: The stopped muon lifetime vs Z .




Figure 10.3: The muon decay-in-orbit endpoint energy vs Z .

- The target material must be chemically stable and form-able into thin, self-supporting sheets.
- The conversion energy for high Z nuclei is considerably below that for low Z targets (Fig. 8.8). When using a high Z target (recall that the DIO endpoint energy is the same as the conversion electron energy) care must be taken that there are no low- Z materials in the target or near it, otherwise DIO electrons from the low Z material can have energies as large as or larger than the conversion electron in the high- Z material and present a severe background.
- The target must be chosen so that the muon radiative capture endpoint energy is below the conversion energy, otherwise photons from reactions of the form $\mu^- + (A, Z) \rightarrow (A, Z - 1) + X + \gamma$ can fall into the conversion electron energy window, presenting a background. (Here X represents pions and other secondary particles.) The nucleus (A, Z) should be chosen so that the rest energies of the possible combinations of $(A, Z-1)+X$ are all a couple of MeV above that of the original nucleus, (A, Z) . This is satisfied for most target materials, including aluminum (2.6 MeV), titanium (600 keV for one 10% abundance isotope, several MeV for others), and gold (700 keV).
- The muon conversion rate increases with Z , reaching a maximum at Se and Sb of about 2.5 times that of Al (Fig. 10.4). For high- Z targets there is also a dependence of the conversion rate on the nature of the interaction (scalar, vector, dipole).
- The fraction of muon captures inside the measurement window should be maximized.

The capture rate is proportional to Z^4 for low Z , while the normal decay rate of a bound muon is about the same as for a free muon. For aluminum, about 60% of mouns

are captured on the nucleus, and the remaining 40% decay in orbit; the lifetime of muonic aluminum, 864 ns, is well suited to the time structure of the beam at Fermilab. The spacing between proton pulses is 1700 ns and the measurement period is between 800 ns and 1700 ns after injection. About half of the nuclear captures occur in the measurement time window.

In the case of a titanium target, the number of captures in the measurement window, with a muon lifetime of 330 ns and a capture fraction of about 85%, is a bit less than for aluminum; however the decay fraction is less and therefore there is less DIO background.




Figure 10.4: The muon conversion rate vs. Z , normalized to the rate in aluminum [27].




Figure 10.5: The muon capture rate and decay-in-orbit rate vs Z .

For a high- Z target like gold, the muon lifetime is 88 ns and almost all muons capture on the nucleus. The number of captures in the measurement window is reduced by about a factor of five to ten relative to aluminum and titanium. In transit to the target, muons have a lifetime of $\gamma\tau$, where $\tau = 2200$ ns. The average transit time is about 300 ns, therefore losses due to decay of in-flight muons are small. Once the muons are captured in atomic orbit, the lifetime is shorter, and in the case of gold far shorter; only those muons which stop in gold just before the beginning of the measurement period will result in a significant contribution to the capture rate during the measurement period. Therefore data rates on a target like gold will depend critically on the rate of late-arriving muons (see Fig. 8.1). Any studies to improve the muon beam would include trying to find ways to stop as many muons as possible just before the measurement period begins, in order to improve rates on high Z targets.

We conclude that aluminum and titanium are viable targets. The conversion rate in titanium is about a factor of 1.7 that of aluminum. Measurements on much higher Z targets are possible, but at a considerable loss of data rate. Since the linear scaling of the conversion ratio with Z begins to fail at medium values of Z , and thereafter starts to decrease, gold has no advantage over aluminum in that respect. However, gold is more sensitive to the details of the interaction as shown in Fig. 10.5.

Table 10.1: Transition energies for muonic atoms: Aluminum and Titanium

Transition	Al (keV)	Ti (keV)
$3d \rightarrow 2p$	66	189
$2p \rightarrow 1s$	356	1021
$3p \rightarrow 1s$	423	1210
$4p \rightarrow 1s$	446	1277

10.2 Muon Capture Monitor

10.2.1 Purpose and Method

Given the complexity of generating and collecting low energy negative muons in the production solenoid, and transporting them via the transport solenoid to the target foils in the detector solenoid, it is evident that some means of confirming the rate and integral number of negative muons which stop on the target foils is crucial. It is equally evident that such a device will prove indispensable in the initial process of tuning conditions for the proton beam and the solenoids.

An effective and reliable Muon Stopping Target Monitor can be established by observing the prompt production of muonic x-rays which signal the formation of muonic atoms in the target foils. This objective can be achieved if it is possible to conveniently locate a germanium detector where it can view, without serious deadtime, photons coming from the target foils. Such x-rays are unambiguously characteristic of a muonic atom's atomic number Z . In addition, the highest yield x-ray is the $2p \rightarrow 1s$ radiative transition which confirms the arrival of a muon in the initial state needed for μ -e conversion. Other observable x-rays, having substantial yields and signaling arrival in the $1s$ state, are the $3p \rightarrow 1s$, and the $4p \rightarrow 1s$. Typically the $3d \rightarrow 2p$ transition which populates the $2p$ state also appears in the energy spectrum. The study of such exotic atom x-rays has a long and productive history which closely parallels the development of semiconductor spectrometers, the Si(Li) for detecting low photon energies and the intrinsic Ge for medium-to-high energies. Members of this Collaboration have had extensive experience in these fields.

A germanium detector crystal of significant size should be used to collect the energy spectrum of the muonic x-rays whether the target foils be aluminum or titanium. Table 10.1 lists approximate energies for x-rays from Al and Ti. Recording the rate of these full energy photons at a FWHM resolution of 2.2 keV unambiguously identifies and monitors negative muons arriving in the $1s$ atomic level of the stopping target material.

In principle, this detector is also sensitive to the possible generation of pionic or antiprotonic x-rays if these negative particles were to reach the target foils. However,

observable x-ray energies from these hadronic atom transitions are cut off when the radiative rate from a given upper atomic level is overtaken by the competing strong nuclear absorption rate. This coupled with the short lifetime of 26 ns for pions should rule out a pionic target-atom signal. On the other hand a beam contamination of antiprotons stopping on the aluminum foils might have a measurable yield for the circular $4f \rightarrow 3d$ x-ray, 202 keV in Al or 588 keV in Ti, and possibly the $3d \rightarrow 2p$, 586 keV in Al or 1679 keV in Ti.

10.2.2 Location for the Germanium Detector

Three requirements determine the best location for the Ge detector to view the muon target:

1. The detector should only view the target, if possible. Hence the first requirement is for good collimation ahead of the detector.
2. Because of the extraordinarily high x-ray rate, about 10^{11} Hz, the detector must be far from the source, along a low-attenuation path, and
3. The detector must be lie beyond the DS magnetic field where it can be serviced periodically with cryogenic liquid and annealed to repair neutron damage.

Figure 10.6 shows an optimal layout for the Ge detector which satisfies these requirements. The photon spectrometer is placed along the axis of the detector solenoid, at the downstream end of the movable concrete shield wall. From there it views all foils head-on, with the front foil 15 m and the downstream foil 14 m away. Collimation is conveniently provided by bore holes in the 0.5 m steel wall and the 1.0 m concrete wall. A sectioned vacuum pipe runs through these walls starting at the back face of the detector solenoid. The pipe is not coupled to the detector solenoid so that the section through the steel wall can travel with the wall whenever it is necessary to gain access to the detector solenoid.

Transmission of 356 keV x-rays passing through all 17 aluminum foils is 90%. At the back end of the DS vessel a 5 cm dia. stainless window of 0.2 cm thickness and thin windows on the vacuum pipe transmit 85% at this photon energy. For muons stopping in the steel window the muonic x-ray energy is $E(2p \rightarrow 1s) = 1426$ keV. Beyond the DS vessel a vacuum pipe transports photons through the stainless steel wall and the heavy shielding block wall to the photon spectrometer endcap. Copper ring collimators placed within the vacuum pipe limit photons to the central 3 cm dia. portion of the germanium crystal, defining a target center-to-crystal fractional solid angle of 3.1×10^{-7} . At a muon stopping rate of 1×10^{11} Hz the germanium detector will process the K-series x-ray events ($2p \rightarrow 1s$, $3p \rightarrow 1s$, $4p \rightarrow 1s$) at 31 kHz. A 45 cm^3 germanium crystal with 3 cm depth is expected to capture full-energy events for the 356 keV x-ray of muonic aluminum with an efficiency of about 50%. For the 1021 keV x-ray of muonic titanium this efficiency drops to 33%. The combination of

Detector




Figure 10.6: Layout of Muon Beam Monitor

high, full-energy event efficiency and excellent peak resolution (2.2 keV) assures that the muonic atom formation process is well determined.

10.2.3 Calibration

The spectrometer system can be calibrated in the standard way, which typically involves one or more calibrated radioisotopes. It happens in our case that a single source of Europium-152 (12.7 year half-life) spans the energy range of muonic x-rays for either an aluminum or a titanium target. A 10 μCi strength source placed just outside the detector endcap can accomplish this. Typically calibration data is collected simultaneously with the experimental data if the muonic lines are not overlapped by calibration lines. The well established gamma energies and experimental yields for ^{152}Eu are illustrated in the semilog plot of Figure 10.7.




Figure 10.7: Pulse-height spectrum of ^{152}Eu . Energies in keV.

To obtain an absolute measurement of the muonic atom formation rate it is necessary to make an absolute calibration of the total efficiency (detector energy efficiency \times solid angle) for gammas leaving the target position and being detected by the germanium detector. For this a special measurement of the total efficiency can be made anywhere once the Ge detector system is available. A stronger ^{152}Eu source ($100 \mu\text{Ci}$) can be used in a full scale mockup of target foils, Ge detector, windows, collimators, and vacuum or helium environment. Locating the source at various target foil positions, spectra can be taken over long time-interval runs to map out the total efficiency.

10.2.4 Selection of Germanium Spectrometer System

A number of vendors supply complete spectrometer systems. Three prominent firms are Princeton Gamma Tech (PGT), Perkin-Elmer (Ortec) and Canberra.

Among our early concerns about viewing in this location the target with such a system were the following:

1. The event rate in the detector for muonic x-rays originating in the target foils is quite high when compared to previous muonic x-ray experiments :

With today's high rate preamplifiers of the transistor-reset variety this is no longer a problem. These are rated as capable of maintaining a FWHM resolution

Detector

of 2.2 keV while handling a count rate of 800,000 MeV per sec. This should readily satisfy our requirements.

2. The detector rate for the interaction of fast muon-capture neutrons (> 0.5 MeV) is high enough to induce damage in the detector's germanium crystal :

This stems from the fact that fast neutrons can cause nuclear dislocations in the crystal lattice thereby building up trapping centers for the signal charge carriers. At a level of 10^9 dislocations per cm^2 in n-type germanium a deterioration of peak shape resolution will begin to show in terms of a low energy tail. The nuclear interaction length in germanium is 88.3 g/cm^2 . Therefore in a 3 cm thick crystal 17% of the neutrons can interact. Given that half of the muons capture on aluminum, producing fast neutrons, this leads to a rate for inelastic neutron events in the Ge crystal of 2.6 kHz. In a 10^7 sec run there will be a need to rehabilitate the crystal every 1-to-1.5 months. Should damage become evident it is possible to anneal the crystal off-line using the vendor's neutron radiation repair kit. For n-type germanium the charge carriers are electrons. When trapped in a lattice defect these can be released quickly with less serious reduction in the total collected charge than for p-type germanium. Hence we would select an n-type intrinsic germanium detector with a fast-reset preamplifier for best high-rate performance. The need to periodically anneal the detector for damage could take this system down for possibly a shift or two. For this reason we believe there is a need for a backup detector in the Target Monitoring system.

3. The transport solenoid delivers a flux of electrons to the detector solenoid which is about 500 times greater than that for negative muons in the initial 150 ns following a proton pulse. In the interval between 150-to-750 ns, when most of the muons arrive, this ratio falls toward 1.0. These electrons have energies ranging up to 70 MeV. Bremsstrahlung photons are produced in the forward direction by 16% of the electrons passing through the target foils. These photons, therefore, add to the muonic x-ray flux arriving at the germanium detector :

Beam electrons are to be swept from the vacuum pipe by applying a transverse field of 0.1 T across the pipe. This would provide a deflection sufficient to remove the electrons from the pipe within a field length of less than one meter. The best location for this sweeper would be between the stainless steel wall and the back concrete block. Bremsstrahlung photons would be dealt with by inhibiting the preamplifier during the 0-to-150 ns interval.

10.3 The Tracking Detector

For the tracking apparatus for the *Mu2e* experiment, we plan to use the design developed by the MECO collaboration, which is described in the MECO draft TDR. The

description below is adapted from that document.

10.3.1 Physics Requirements

One potential limiting background in the *Mu2e* experiment is muon decay in orbit (DIO). We note that the endpoint of the DIO electron energy is the energy of the electrons emitted in coherent muon conversion. Both the absolute normalization and energy spectrum have been calculated [127, 128] and these calculations agree to a precision of about 25%. The small discrepancy in the calculations can be traced to different approximations for nuclear recoil effects and relativistic corrections to the muon wave function. Thus the level of background is related directly to the precision of the electron energy measurement, so that, in principle, the background can be made arbitrarily small by improving resolution. Since the DIO spectrum falls as the fifth power of the difference between the endpoint energy and the energy of the DIO electron, the background level is very sensitive to resolution. To reduce background, the central part of the resolution function must be minimized and all high energy tails suppressed.

To find the constraints on the width of the resolution function and its high energy tails, we define the detector response function as $f(x)$, where x is the difference of detector-reconstructed energy and the actual energy of the conversion electron: $x = E_{det} - E_0$. Generally, $\int f(x)dx < 1$, due to finite acceptance. If one defines the signal region as $x > \Delta$, then one can define an acceptance function $A(\Delta)$ as

$$A(\Delta) = \int_{\Delta}^{\infty} f(x)dx.$$

Assuming the normalized background takes the form [128]

$$\frac{dn}{dE} = -Cx^5,$$

with $x < 0$, one can define a noise function $N(\Delta)$ as

$$\begin{aligned} N(\Delta) &= \int_0^{\infty} Cx^5 dx \int_{x+\Delta}^{\infty} f(y)dy, \\ &= \int_{\Delta}^{\infty} f(y)dy \int_0^{y-\Delta} Cx^5 dx, \\ &= \int_{\Delta}^{\infty} \frac{C}{6} (y - \Delta)^6 f(y)dy. \end{aligned} \tag{10.1}$$

One can further define a noise/signal ratio R as

$$R(\Delta) \equiv \frac{N(\Delta)}{R_{\mu e} A(\Delta)}, \quad (10.2)$$

where $R_{\mu e}$ is the μ - e conversion branching ratio, taken to be 10^{-16} in the following studies. Estimates of C are inferred from the references [127, 128], giving $C \approx 0.6 \times 10^{-16} \text{MeV}^{-6}$.

Figure 10.8 shows the acceptance as a function of σ using a Gaussian detector resolution function, with Background/Signal=0.05. Compared with the acceptance for a perfect detector with $\sigma = 0$, the acceptance is 90% at $\sigma = 300$ keV, and 70% at $\sigma = 400$ keV. The acceptance drops quickly for $\sigma > 400$ keV. Thus, we require $\sigma < 400$ keV, or FWHM < 1 MeV. Constraints on the magnitude of a high energy tail can be qualitatively inferred from Equation 10.2. To illustrate, if we require the extra background contribution from the high energy tail to be less than 2% of the acceptance, then the magnitude of any high energy tail above ΔE should be less than $0.2 \text{ MeV}^{-6} \times \Delta E^6$; e.g., a high energy tail above 5 MeV should be less than 10^{-5} .




Figure 10.8: Acceptance as a function of σ , using a Gaussian detector resolution function; assuming $R_{\mu e} = 10^{-16}$ and Background/Signal=0.05. The acceptance is normalized to 1 for a perfect detector, $\sigma = 0$.

10.3.2 Tracking Detector Overview

The tracking detector is located in a uniform 1 T magnetic field. The goal of this detector is to measure with good efficiency the parameters of the helical trajectory

of electrons. The uncertainty in helix parameter measurements is dominated by multiple scattering in the tracker. A second source of error in the determination of the energy comes from pattern recognition errors. This latter source does not reduce the acceptance significantly, but is a potential background because it might generate high energy tails in the resolution function. Spatial resolution of detector hits along the helical trajectory does not contribute significantly to energy resolution. In addition, inferring the energy of an electron depends on knowing the energy loss in the stopping target and in any material between target and detector, e.g., proton absorbers. Energy loss has two effects. One is to broaden the central part of the resolution function and introduce a small mean energy loss. The second is to introduce a low energy tail. This latter effect is essentially equivalent to a loss of acceptance and does not introduce backgrounds.

A “good geometry” spectrometer should measure the radius of the projected circular pattern of the electron with minimal error. The helical path is sampled at three positions, with the first and last positions separated by approximately 180° . A minimum of three position measurements is required, but more points are needed to overconstrain the fit and reduce backgrounds. This is particularly important because accidental hits in the detectors can be combined with true signals to yield a trajectory that reconstructs with energy in the μ - e conversion region, about 105 MeV.

Given the need to minimize material, straw tubes in vacuum are the technology of choice for tracking. Large systems of straw tubes are common in high energy physics experiments. A large system of straw tubes in vacuum has operated successfully at the COSY-TOF facility[129]. The fact that the electrons spiral $\sim 180^\circ$ in vacuum in the center of the Detector Solenoid is a special feature that minimizes the contribution of multiple scattering to the resolution.

These general design considerations lead to a detector geometry referred to as the Longitudinal Tracker (L-tracker), and is the baseline tracker to which other designs have been compared. The L-tracker consists of an octagonal array of eight detector planes placed symmetrically around the Detector Solenoid axis, plus eight more planes (also referred to as vanes) projecting radially outward from each vertex of the octagon. Each plane is approximately 30 cm wide, 260 cm long, and has 3 layers of straw tubes in a close-packed arrangement. A hit position in the radial and azimuthal direction is determined by the straw position and the drift time on the anode (wire). The hit position in the axial direction is determined by the centroid of the imaged charge from the anode wire, as collected on cathode strips etched on two thin, mylar sheets sandwiching the straw planes and vanes. In some studies reported here, the octagonal array has been modeled as a cylinder.

When all straw tube wires are parallel to the DS axis, helical trajectories return to the same r - ϕ point (but different z) after each orbit. Pattern recognition studies have shown that to reduce backgrounds to an acceptable level more than one turn along the helix must be measured. To eliminate multiple “hits” on the same detector element, the planes and vanes of the L-tracker are rotated by a small angle (typically

15 mrad) about an axis perpendicular to the planes. This tilt also prevents low energy electrons generated by gamma rays from depositing all of their energy in a single straw by following a magnetic field line. The detector length is 2.6 m, chosen on the basis of extensive studies that have been done for detectors with lengths between 2.4 m and 2.9 m. For a 2.4 m detector, 39% of conversion electrons emitted with $p_t > 91$ MeV/c have at least 6 hits in the tracking detector; a 2.9 m detector guarantees that two full helix turns are measured for the same class of events. This is discussed further in Section 10.3.3.

The minimum radial distance to the octagonal planes is 38 cm in order to make the rates from DIO electrons small compared to those from photons and protons. A single turn of a typical conversion electron trajectory crosses the octagon twice and either one or two vanes; we refer to these as 3- and 4-hit turns. Figure 10.9 views the tracker along the DS axis with three trajectories superposed. The transverse mo-




Figure 10.9: Cross section of the tracking detector, the stopping target, and trajectories for electrons created in the target with transverse momentum of 55, 91, and 105 MeV/c. The trajectories are positioned to show the minimum allowed detector radius that keeps rates from muon DIO manageable.

menta of these trajectories (referenced to the stopping target position) are 55 MeV/c (the momentum exceeded by only 0.3% of decay in orbit electrons), 91 MeV/c (the transverse momentum of a conversion electron emitted at 60° to the beam axis), and 105 MeV/c. Note that the figure shows the maximum physical stopping target radius, 8.30 cm; while a circle inscribed inside the octagon is 38 cm in radius.

10.3.3 Tracking Detector Performance

Detector Rates

High rates in the detector may limit the experimental sensitivity in several ways. First, detector occupancy might be so high that the efficiency in collecting valid position information is reduced. This problem may be addressed by making finely segmented detectors with short integration times. Second, noise (accidental) hits may be combined with those from a lower energy trajectory to create a trajectory consistent with that of a conversion electron (105 MeV). This is a problem common to many high rate experiments that look for rare events, and it may be suppressed by reducing time resolution, by having sufficient redundancy in particle position measurements, and by discriminating against signals from particles other than the electron of interest. Third, there may be contributions to the trigger rate due to pile-up of lower energy signals in the trigger detector. This possibility may be reduced by segmentation and geometric design of the trigger detector.

The detectors are active only (in principle) between muon pulses. At this time few, if any, protons hit the production target, so particle flux in the beam channel is very low. However during a 200 ns period following the proton pulse, there will be a very high flux of charged particles passing through the muon beam channel and the detector solenoid. In this time interval the detector rates will be so high that straw tube gains may have to be reduced by a factor 10-100. We also note that the magnetic field is designed so that there are no magnetic traps, to ensure that there are no late arriving particles, or ones making multiple passes through the detector.

During the active time window, detector rates arise from four main sources.

- Muon DIO electrons potentially produce high detector rates, but the detector is designed so that the majority of these are restricted by the magnetic field to radii less than 38 cm.
- Roughly 60% of stopped muons are captured on nuclei. These captures produce photons, neutrons, and protons from nuclear de-excitation. Approximately 2 γ 's, 2 neutrons, and 0.1 protons are emitted on average for each capture.
- Beam electrons emit bremsstrahlung radiation as they traverse the stopping target. These photons may Compton scatter and pair produce in the detectors. The muon beam line is designed to reduce the flux of low-energy electrons that reach the stopping target in the measurement time window; however, beam electrons are the dominant contribution to detector rates during the time immediately following the proton pulse at the production target.
- Albedo from the muon beam stop can reach the detector, but beam stop design reduces these rates to a low level.

Detector

The products of each of these sources also strike the beam stop, contributing to the detector rates. Muons stopping in the dump represent less than half the total flux of beam muons, furthermore, the probability that a muon decaying in the dump causes a hit in one of the detectors is significantly less than that of a muon in the stopping target.

Rates have been calculated using full GEANT simulations of the interior of the detector solenoid. In this simulation, first, the spatial distribution of stopped muons was calculated using GEANT, then the distribution of the source of particles that potentially cause detector hits was chosen. Energy spectra of particles emitted from nuclei following muon capture and electrons from muon DIO have been taken from the literature, as discussed later; these particles were then generated isotropically with the appropriate energy distribution and tracked through the magnetic field. All physical processes in the materials of the detector solenoid, and muon beam dump were included. Some of these rates depend on the amount of material in the tracking detector, and this has been modeled in some detail, including the structure supporting the straws, cabling, etc.

Detector design is driven by the need to be insensitive to the majority of the approximately 10^{11} muon decays per second; this is an advantage of $\mu^- N \rightarrow e^- N$ experiments with respect to $\mu \rightarrow e\gamma$, since the signal electrons have twice the maximum energy of electrons from μ^- decay in vacuum. For muon DIO, the spectrum extends to 105 MeV as is shown in Figure 10.10. To simulate detector rates from this source, electrons were generated with this spectrum in accordance with the previously determined stopping distribution in the target, and the hit rate in the tracker was calculated. The electron rate is dominated by muon DIO, protons from muon capture in the stopping target, electrons traversing the detector, and from bremsstrahlung photons that pair produce or Compton scatter in the tracking detector (often after first scattering somewhere else in the detector solenoid). The total contribution is significant only in the octagonal planes of the detector; a full breakdown of the contributions to the tracker rate are shown in Table 10.2.

There is an unavoidable flux of γ 's, protons and neutrons produced by muon capture on nuclei. Every μ^- capture results in the production of excited nuclear states, radioactive nuclei, and/or neutron emission with the possibility of subsequent neutron induced nuclear gamma rays. This results in photons originating from various places in the detector solenoid, some fraction of which are not associated with the beam gate. Almost all of these photons are less than a few MeV (the binding energies of the most probable excited nuclei after μ^- capture are less than 4 MeV). The simulation assumes a flat energy spectrum out to 10 MeV at a rate of 1.8 γ 's per μ^- capture.

Protons are also emitted during the μ^- capture process. The proton spectrum, which has energies predominately below 15 MeV, was taken from an experiment [130] using muons stopping in emulsion. The shape of the spectrum is almost Gaussian, centered at ~ 7.5 MeV proton energy, with a ~ 5.5 MeV width and a high energy




Figure 10.10: Plot of the differential and integral electron energy distributions for μ decay in orbit. Circles are the integral of the distribution for energies above the abscissa value, and crosses are the differential distribution.

tail extending to above 50 MeV. The normalization is somewhat uncertain and depends strongly on nuclear size. The best available data on the normalization is from Budyashov et al. [131]; other experiments [132–135] also report measurements on different nuclei. The assumed proton spectrum, which is conservatively chosen to be the largest reported flux of 0.15 protons per μ^- capture, is shown in Figure 10.11. The protons have relatively high momentum, but low kinetic energy and are easily absorbed.

The largest contribution to the rate is from protons; the total instantaneous flux of protons exiting the stopping target is $\sim 0.5 \times 10^{10}$. Without shielding, the average rate in individual tracking detector elements would be well above 0.3 MHz. However, the protons can be attenuated significantly by a set of absorbers. One is a carbon fiber conical shell of thickness 1.0-3.0 mm surrounding the stopping target, with inner radius, 45.0 cm, at the DS entrance, and 69.0 cm three meters downstream. The geometry of this upstream absorber is chosen so that it is not hit by 105 MeV electrons originating in the stopping target. A second polyethylene absorber consists of a cylindrical shell of thickness 0.5 mm and with radius slightly smaller than the inner radius of the tracker. It extends from just downstream of the stopping target to the beginning of the tracker. The effect of proton absorption in the stopping target and the absorbers is shown in Figure 10.11. The lowest momentum protons are fully absorbed, but the remaining protons typically have a mean ionization rate $\sim 10 \times$ minimum ionizing.




Figure 10.11: The distribution of the kinetic energy of protons originating from μ^- capture on Al, taken from a fit to the data of reference [130]. The shaded histogram gives the distribution of protons that cause hits in the tracking detector.

Rates due to photons were investigated in a similar calculation and reported in Table 10.2. Many of the photons interact in the tracker after scattering in other material, but tracker hits are caused by low momentum electrons from Compton scattering or pair production, and these typically make multiple passes through a given straw within a very short time.

Neutrons are produced during the μ^- capture process. A neutron spectrum, typical for our target, can be created from experimental data [136, 137]. Neutrons below 10 MeV are produced with a thermal distribution, and there is an exponential tail above 10 MeV. Detector rates have been calculated assuming two neutrons are emitted per μ^- capture.

The rate of neutron hits is sensitive to the detailed geometry of the detector solenoid, but the present simulation suffers from the fact that neutrons are tracked only down to 10 keV, at which point they deposit their remaining energy locally. The neutron flux will be attenuated through the use of appropriate neutron absorbers, for example in the region upstream around the stopping target and outside the conical proton absorber. In any event, neutron induced signals in the tracker often do not have three straws hit in a cluster and thus can be removed from inclusion in the

Table 10.2: Peak detector rates R_{wire} at the beginning of the measurement interval 700-1600 ns. The average number of hits from one background particle in the entire detector is $\langle N_{event} \rangle$. The subscript t denotes backgrounds originating at the stopping target and the subscript bd denotes backgrounds originating from the beam dump.

#	Background type	Rate (GHz)	Probability to hit detector	$\langle N_{event} \rangle$	R_{wire} (kHz)
1	e_t	70	32.0×10^{-5}	0.5	20
2	n_t	61	14.2×10^{-5}	1.0	15
3	γ_t	61	24.8×10^{-5}	1.5	42
4	p_t	5	362.0×10^{-5}	1.8	66
5	$eDIO_t < 55$	24	9.8×10^{-5}	0.45	1.7
6	$eDIO_t > 55$	0.06	127.0×10^{-5}	7.0	0.6
7	n_{bd}	14	7.1×10^{-5}	1.5	0.6
8	γ_{bd}	14	8.3×10^{-5}	1.5	2.0
9	$eDIO_{bd} < 55$	60	8.9×10^{-5}	0.31	2.0
10	$eDIO_{bd} > 55$	0.16	18.2×10^{-5}	0.5	0.02
11	$eDIF$	0.0007	1	11.1	10.7
12	Total				144

pattern recognition codes.

Late arriving beam electrons also cause additional detector hits. The hits come from bremsstrahlung in the stopping target, with the γ 's subsequently Compton scattering or producing pairs in the tracker. This contribution, denoted e_t , is given in the first line of Table 10.2. We note that late arrivals may be reduced by improved design of the muon beam.

The total rate per detector element is ~ 150 kHz. These rates are lower than those in the straw chambers of similar construction, used in BNL E781. During a 30 ns gate with a typical drift velocity of $100 \mu\text{m}/\text{ns}$, the average occupancy will be under 1%, and many of the signals induced by these particles may be distinguished from those induced by conversion electrons. Some examples are:

1. Hits by protons have large pulse height compared to conversion electrons.
2. Hits by electrons from Compton scattering and pair production will also typically have high pulse height since these make multiple turns through a single straw.
3. Hits can be distinguished by their time structure which will not be consistent with that of a spiraling 105 MeV electron.

4. Hits due to neutrons will have high energy deposition and will typically not activate all three layers of a straw tube plane.

Response To Signal Events

In this section, we discuss the tracking detector performance for true signal events in a 2.4 m long detector. This study used a full GEANT simulation of the target and detectors [57, 138]. It incorporated the full Moliere scattering, Landau fluctuations in the energy loss, and Gaussian measurement errors. The errors in position measurements were taken to be $\sigma_x = 0.2$ mm, $\sigma_y = 0.2$ mm and $\sigma_z = 1.5$ mm. The simulation used electrons originating from the target disks, exiting isotropically and appropriately distributed according to the relative muon stopping probabilities in the disks.

The study reported here does not use a full model of the tracking detector including drift times in the straw cells. Positions of the electron trajectory-straw tube plane crossings were recorded, referred to as the *cluster* position, and then randomized according to the resolutions given above. Straw chamber inefficiencies are also not incorporated in this study; however, by allowing up to two of the six or more clusters on each trajectory to have one of the three straw signals missing, a single cell efficiency of 97.5% would result in a 3% loss of events due to straw chamber inefficiency.

In addition to generating simulated cluster positions from the DIO or conversion electrons, clusters induced by the high ambient flux of protons, neutrons and photons emitted following muon capture were superimposed on events. This was done by using a GEANT simulation to produce events in which these background particles produced hits in the tracking detectors. A large sample of these events were produced and stored, and randomly sampled at the appropriate rate. In our studies, the average number of noise clusters is eight, as determined from the estimated flux of background particles in a 15 ns time window around the electron time.

The analysis could be further improved by taking advantage of correlations of the pulse shapes or timing between the straws and pads. A conservative approach has thus been taken in the pattern recognition analysis. Since individual straw signals are not simulated, many of the tools that could be used to reject noise clusters have either not been incorporated into the analysis or have been incorporated only in an approximate and rather conservative way. Other background rejection techniques are discussed in the following section on backgrounds from pattern recognition errors.

About 60% of all conversion electrons hit the detector. Figure 10.12 shows a typical event in the simulation. The number of times the electron helical trajectory turns within the detector region is determined by its pitch angle $\theta_p \equiv \arctan(p_t/p_l)$, and detector length. Figure 10.13 shows the distribution in the number of clusters in the tracking detector for conversion electrons that hit the detector. A small tail extends beyond 14 clusters from electrons that lose significant energy in the tracking




Figure 10.12: Simulated conversion electron trajectory.

detector and thus make many turns in it. Figure 10.13 also shows the distribution in pitch angle for the same electrons at the upstream end of the detector. The




Figure 10.13: The left frame shows the cluster number distribution for conversion electrons within the tracking detector acceptance. The right frame is the pitch angle distribution for the same set of events.

requirement is that signal events have a measured value of θ_p in the range $45^\circ < \theta_p < 60^\circ$. The lower limit on θ_p eliminates electrons originating in regions with magnetic field of 2 T or larger, e.g., in the transport solenoid or in the final collimator, and also minimizes backgrounds from beam electrons that scatter in the target. The upper

limit helps to reduce backgrounds from cosmic rays or events produced in the proton absorber. This limit is equivalent to a requirement that the electron propagates upstream to the target without being reflected in the increasing magnetic field.

The pattern recognition procedure is to first select clusters that form circles in the xy (transverse) projection and then match them with z (axial) clusters to look for good helical tracks. All possible combinations of clusters are tried, including combining all xy clusters with all z clusters in a given vane or segment of the octagon. The momentum for the helix track is then determined by a fitting algorithm that uses a likelihood method to be described below. The fitter returns the most probable momentum and a corresponding likelihood at this momentum. Since the fitter is very CPU intensive, preliminary selection criteria are used to obtain potentially good circles in the transverse plane and full helices in three dimensions so that the number of possible trajectories analyzed by the fitter is minimized. These criteria are chosen to reduce computing time while producing a negligible reduction in acceptance for signal events.

We now briefly describe the fitter. It works on the principle of the maximum likelihood method, and determines the most likely momentum of a particle that made the helical trajectory. The strategy notes that the trajectory mainly deviates from a helix due to multiple scattering in the detector elements, but individual segments between adjacent hits are helical. The fitter exploits this to determine the trajectory of each segment between detector crossings as a function of the electron momentum p_e and then calculates a likelihood value $L(p_e)$ for the full trajectory as a function of p_e . This likelihood value is simply the product of the scattering probability at each detector position.

$$L(P_e) = f_2(\theta_2)f_3(\theta_3)...f_{n-1}(\theta_{n-1}) \quad (10.3)$$

Here n is the total number of hits and $f_i(\theta)d\Omega$ is the probability that the particle scattered into the solid angle $d\Omega$ at θ_i in the detector element where the i th hit was recorded. The parameter $f(\theta)$ takes a Gaussian form for small angles and has Moliere tails for large scattering. It is also possible to incorporate energy loss and the detector spatial resolutions in Equation 10.3. This is discussed in detail in reference [138].

The most probable value of p_e is that which maximizes the likelihood. To estimate the error on the value of p_e from the fitter, the distribution in the likelihood vs. p_e in the region of the peak is fitted with a Gaussian form. The σ of this fit is denoted σ_{p_e} and it gives a good estimate of the uncertainty in the measured value of p_e . The parameter σ_{p_e} and the maximum likelihood value are found to be powerful discriminants against events with badly fit trajectories. The above algorithm is derived assuming the hits which are used are those actually made by the particle track (i.e. no pattern recognition errors). The same algorithm is found to work well even with noise, after applying additional selection criteria discussed below.

The following are the selection criteria imposed to select well measured particle trajectories following the track fitting.

1. The value of the likelihood is required to be greater than some value.
2. The scattering angle at each detector element is required to be less than 0.08 radian.
3. The fitted uncertainty σ_{pe} is required to be less than 600 keV.
4. The total number of clusters is required to be at least 6 for a 2.4 m detector. This requirement varies with detector length, and it significantly reduces high energy tails in the resolution function, primarily from pattern recognition errors.
5. The fitted trajectory is required to have a cluster at each intersection of the trajectory and a detector element.
6. The projection of the fitted trajectory to the point where it intercepts the electron calorimeter is required to agree with the position the electron entered the calorimeter to within 20 cm.
7. The energy of the GEANT primary electron at the entrance of the electron calorimeter is required to be at least 75 MeV.
8. An event is rejected if a lower momentum track is found with a suitably relaxed set of selection criteria. This significantly reduces background from pattern recognition errors with essentially no loss of acceptance for the signal events.

The intrinsic energy resolution (excluding the effects of energy loss in the target, but including spatial resolution in the tracking detector and the effect of noise) is found to be $\sigma_{RMS} = 110$ keV. The effect of energy straggling in the target causes the resolution function to deviate from a Gaussian shape at low energies but does not introduce a high energy tail into the resolution function. The FWHM of the response function is 900 keV. We summarize the efficiencies of the critical selection criteria in Table 10.3. Further suppression of some backgrounds can be obtained by restricting

Table 10.3: Selection criteria used in the electron momentum measurement.

Selection criterion	Efficiency
At least 6 hits in tracking detector	0.44
Detected energy above ~ 103.6 MeV	0.62
Required pitch angle at the detector	0.88
Requirements on fitting quality	0.83
Position match in electron calorimeter	0.97
Overall acceptance	0.193

the upper limit on the electron energy. For example, restricting the electron energy to be $103.6 \text{ MeV} < E_e < 105.1 \text{ MeV}$ introduces negligible acceptance loss.

Backgrounds Induced by Pattern Recognition Errors

We next turn to a discussion of backgrounds due to very high energy tails in the resolution function of the spectrometer. These have been shown to be primarily due to an analysis combining hits from a low energy electron and random accidentals (noise clusters). We refer to this as *pattern recognition errors*.

While the preceding discussion of backgrounds explicitly includes the possibility of background due to pattern recognition errors, it is limited due to the finite statistics of the calculation, consisting of 10^7 fully simulated events, including noise. The statistical level of this simulation is sufficient to calculate the background arising from DIO electrons with energy above 100 MeV, of which there are $\sim 10^5$ for the total experiment run time. For DIO events below 100 MeV there are many more events, thus additional calculations are required in order to estimate the expected background level. These calculations are described in detail in reference [57]. Background could, in principle, arise from a variety of sources: multiple scattering, large tails in position resolution, etc. DIO electrons in this energy range primarily produce potential background by pattern recognition errors [57], but at a level well below the sensitivity of the experiment. Those calculations are briefly described here. In this discussion, a background event is defined to be a reconstructed track with momentum above 103.6 MeV. The pattern recognition and fitting code imposes an upper limit of 120 MeV on the reconstructed energy; hence the background is in the interval $103.6 \text{ MeV} < E_e < 120 \text{ MeV}$.

Two independent studies were made to determine the number of mis-reconstructed events. Both used GEANT simulations of the detector, but they used different pattern recognition and background rejection strategies. The calculations proceeded by determining the cluster positions of the DIO electrons and superimposing, on average, an additional 24 noise clusters. This is larger than the expected noise rate and allows an efficient determination of the most probable topology of background events.

The more powerful of the two pattern recognition strategies applied the same reconstruction procedure to the DIO events, described in the previous section. Results for two different detector lengths and several energy intervals for DIO electron events are given in Table 10.4. A minimum of five clusters was required to form a helix in the 2.0 m detector, and six for the 2.4 m detector.

It was possible to generate sufficient statistics for DIO electrons with $E_e > 95 \text{ MeV}$ to show by direct simulation that the background rate is negligible. For lower energy, it is not possible, given available computing power, to generate the full sensitivity of the experiment. Hence, a study of a sample of events with somewhat relaxed selection criteria is used to infer the properties of events that are likely to cause background. The basic procedure is to relax selection criteria and study the number of noise hits necessary to make background events. It is found that the background events typically use a large number of noise hits; the high energy part of the resolution function is dominated by events with many noise hits. To study a sample of events with higher probability to produce background, events are generated with higher

Table 10.4: Pattern recognition results for two different detector lengths and several production windows. A background event is an electron reconstructing with momentum exceeding 103.6 MeV/c and satisfying the selection criteria listed in the preceding section.

Detector length	DIO energy window	Events thrown	Background events
2.4 m	95-100 MeV	5×10^7	0
2.4 m	80-95 MeV	1×10^8	0
2.4 m	75-80 MeV	4×10^7	0
2.0 m	95-100 MeV	4×10^7	2

than expected noise rates, and the events are weighted appropriately. The specific procedure used to calculate the expected level of background is to weight each event by suppression factors corresponding to the random probability that the particular event with a certain number of noise hits would occur if the simulation were done with the expected noise rate rather than an inflated noise rate.

There are several suppression factors to apply. A mean noise rate three times higher than expected is used, so there is an event suppression factor of $f_t = (1/3)^{n_t}$, where n_t is the total number of xy and z noise clusters used in the fitted trajectory. Furthermore, the detector plane-helix crossing angle in the transverse plane is determined to high precision; this angle can be inferred locally using the full straw-tube drift time information with a resolution of about 50 mrad. These angles can be compared with the local helix crossing angle to reject clusters if they are not consistent. A conservative (high efficiency) cut on the agreement is estimated to result in a suppression factor of $f_\theta = (1/8)^{n_{xy}}$, where n_{xy} is the number of xy noise clusters. This is equivalent to the statement that only 1 noise cluster in 8 would have local cluster information consistent with being due to a particle with a well known trajectory inducing the hits. An event suppression factor is taken as the product $f_t \times f_\theta$, and this is then averaged over the background event sample; this average factor is denoted f_s .

For electrons between 95–100 MeV, f_s is estimated to be 0.005 for a larger sample of backgrounds satisfying relaxed cuts. Since there is no background found after the final cuts, P_b , the probability of producing a background per DIO electron is calculated as if the sample corresponded to one background event: $P_b = 0.005/5 \times 10^7 = 10^{-10}$. Multiplying this probability by the total number of DIO electrons in this energy window, 0.0005 background events are expected.

Similarly, for electrons between 80–95 MeV, f_s is calculated to be 0.00006 for a relaxed sample of background. For this sample, $P_b = 0.00006/10^8 = 6 \times 10^{-13}$, and the background from DIO electrons between 80–95 MeV is estimated to be $6 \times 10^{-13} \times 9 \times 10^9 = 0.005$. For electrons between 75–80 MeV, no background was found even when the fitting criteria were considerably relaxed. Since P_b is smaller for lower energy

Detector

electrons, the value of P_b for 80–95 MeV electrons is used and multiply by the total DIO electron number, as this overestimates the expected background. By taking into account the triggering efficiency, 0.6 for an 80 MeV electron, the background from DIO electrons in the energy range 75–80 MeV is found to be $6 \times 10^{-13} \times 0.6 \times 5 \times 10^{10} = 0.02$ events.

In this study, only events with four hits per helical turn were accepted. Other studies have shown that the background level for reconstructed events with three hits per helical turn have approximately the same background level. Summing all backgrounds and accounting for the exponential dependence of the detector rate during the 700 ns detection time, the total expected background is 0.067 events. Doing the same analysis, but assuming a noise rate two times higher than expected results in an expected background of 1 event.

These background events are consistent with being uniformly distributed in the energy interval of interest, $103.6 \text{ MeV} < E_e < 120 \text{ MeV}$. Hence, the number of events expected in the signal window size, $103.6 \text{ MeV} < E_e < 105.1 \text{ MeV}$ is about a factor 10 less, or 0.006 events at nominal noise rates.

The above calculation is conservative in a number of ways. First, we have used essentially no energy matching in the calorimeter and only very loose position matching requirements. Second, additional rejection against noise hits can use pulse height information. Protons are heavily ionizing and such hits can be rejected with high probability. Photons cause hits by Compton scattering or pair production; in both cases the electrons make helical orbits within a drift tube and can also be distinguished by their high apparent ionization. Third, no correlation is assumed between xy and z clusters; however, both pulse height and time matching can be used to reject incorrect pairings. Finally, events with a low momentum track can be rejected. Currently only low momentum tracks with 4 hits per helical turn are found; rejecting low momentum tracks with 3 hits per turn will further reduce backgrounds.

A more straightforward strategy, employing helical roads, was also used in an independent simulation and analysis of background from DIO events. This algorithm reconstructed events with both 3 and 4 hits per turn and used similar noise rates. It also used energy matching and tighter space matching in the electron calorimeter, but did not use local track angle information. Although it achieved background rejection a few times worse than the strategy described above, it did give another set of direct simulation results on the sensitivity of the background rate to factors like the noise rates, local track angle information, and detector length. The dependence of the expected background level on the noise rate is similar to that of the other analysis, as is the dependence on the energy of the muon DIO electron.

We conclude that if accidental rates are as expected, the muon DIO background from event mis-reconstruction would be < 1 event with a 2 m detector and is significantly under one event for 2.4 m. This is not a limiting background, however. Making a longer detector provides significant further rejection and insurance against detector accidental rates that are higher than calculated. Some additional background

rejection tools not currently used, provide further rejection possibilities. Based on these studies, we propose a tracking detector length in the range 2.4–2.9 m. The actual length will be chosen by balancing construction constraints against the desire for redundancy in event reconstruction.

10.3.4 Mechanical Construction

Straw and Pad Designs

The baseline tracking detector is constructed of 5 mm diameter straw tubes in the geometry described in Section 10.3.2. Both the octagon plane and vanes are composed of detector planes formed by three layers of nearly axial, close-packed straws [139]. The straws have 25 μm thick carbon loaded Kapton walls, so that the axial coordinate of a hit can be measured by reading the induced charge on capacitively coupled foils placed on both sides of the octagonal plane and vane detectors [140]. These foils are made of 25 μm Kapton having 5 mm pitch copper strips etched perpendicular to the wire direction. The number of channels is given in Table 10.5.

Table 10.5: Number of detector channels

Detector Element	Component	Number		
		per Element	Total	Readout
Octagon	Straw Tubes	180	1,440	TDC and ADC
	Strips	1,040	8,320	TDC and ADC
Vanes	Straw Tubes	180	1,440	TDC and ADC
	Strips	1,040	8,320	TDC and ADC

Straws of length 2.6 m require one intermediate support for the anode wire. The straws and wires will be attached to manifolds which provide gas and electrical connections. The manifolds reduce the acceptance of the apparatus since electrons passing near them must be rejected from the sample.

The position resolution constraints are well within what has been demonstrated in a number of experiments. The simulation assumes a Gaussian position resolution of 200 μm for the drift coordinate and 1.5 mm for the axial coordinate z , for the mean coordinate of a cluster of hits (3–4 drift coordinates and two z coordinates). Typical drift resolution in detectors of this type is 160 μm for each drift coordinate [141] and $\sim 10\%$ of the pad width for capacitive strip readout. Operation in a magnetic field will somewhat degrade the performance of the system, but this is not expected to present significant difficulties. In any event, simulations have shown that tails in the position resolution of a few percent which extend out to a straw diameter do not adversely affect the momentum resolution.

Straws similar to these, although somewhat shorter, have been successfully used in BNL Experiment E871 [142]. Low density straw tube systems of 2.7 m, 2.6 m, and 2.4 m have been successfully built and tested by several groups [139], and readout of induction pads through resistive straw cathodes has been demonstrated [140].

Deformation of Straw Tubes

Deformation of the straws when loaded by gas pressure and wire tension was investigated. If treated as a cylinder, the internal gas pressure results in an outward force of ~ 2.0 N per straw, and this exceeds the expected wire tension of ~ 0.5 N. At issue is the extent to which the straw deforms due to this loading. It is noted that long straws are generally placed under an outward tension of about 1.8 N in order to align them before wiring. However stretching of the straws after installation into their planes could cause out of plane deflections resulting in variation in efficiency over the straw length. The fractional stretch of a straw was tested by increasing the pressure in a sealed straw with one end fixed and the other free. The typical fractional change in length is 0.04% for one atmosphere overpressure. This would not over-tension the wires, but will likely require a mounting system in which one end of the tracking detector can move axially as the interior to exterior pressure changes. Results for the resistive straws may be different.

Straw Operation in Vacuum

We tested the leak rate of straws under vacuum. Our tracking simulations have assumed we would use straws similar to those used in E871, which are made of two layers of Kapton, each 0.0005 in thick, spiral wound with a half strip overlap. The inner layer has ~ 1000 Angstroms of copper deposited on the interior. We tested the leak rate of both the bulk straw material and the end fittings by measuring the rate of rise of pressure in an evacuated tube containing sample straws with one atmosphere pressure. The rise of the chamber pressure was measured as a function of time after the pump valve was closed. The rise decreased with initial pumping time, indicating it was due to out-gassing of the exterior of the straw. The residual rise after 5 days of pumping corresponded to a leak rate of $\sim 2 \times 10^{-8}$ l/min/m for the bulk straw, and a leak rate of $\sim 3 \times 10^{-9}$ l/min per end. These leak rates, when scaled to the full spectrometer, are well within modest pumping rates. Furthermore, straw tubes have been operated in vacuum in a previous experiment [143, 144]. We have also constructed prototype low-mass gas and electrical manifolds in which a fraction of a “vane” module has been assembled.

10.3.5 Pad Readout

A significant feature of the straw chamber spectrometer is the use of pad readout to determine the coordinate along the straw. This technique has been previously

demonstrated [140]. Straws are constructed of carbon loaded Kapton, which have a resistivity between $500 \text{ k}\Omega$ and $1 \text{ M}\Omega$ per square. A resistive cathode allows an electrical signal to be induced on copper pads deposited on a thin film of Kapton placed outside and perpendicular to the straw assembly. One pad layer will be on each side of a three-layer straw plane as shown in Figure 10.14. Using strips 5 mm wide and interpolating the amplitude of the charge deposited on these pads, we anticipate a position resolution of the charge centroid on the anode wire of $\leq 1 \text{ mm}$. A drawing of the pads is given in Figure 10.15. In the detector simulations, a resolution $\sigma = 1.5 \text{ mm}$ was used, which is significantly larger than what is routinely achieved. Small prototypes were built by both the Houston[78] and Osaka groups for MECO.




Figure 10.14: Detail of the *L*-tracker straws and pads. There are three layers of hexagonal close packed straws. The outer layers are resistive, while the middle straws are conductive. Pads running perpendicular to the straws pick up induced signals to determine the second coordinate.




Figure 10.15: Sketch of a typical set of eight pads with the cables.

Choice of Straw Resistivity

To determine the axial hit position, the collected anode charge must be imaged and read from cathode strips placed perpendicular to the straws. The straws must be

Detector

constructed of resistive material with reasonably thin walls ($\sim 25 \mu\text{m}$). Resistive Kapton can be supplied in thicknesses $\geq 19 \mu\text{m}$ and with standard resistivities $0.5\text{--}1.0 \times 10^6 \Omega/\text{square}$.

The expected signal from a cathode pad was studied by using an equivalent circuit model as shown in Figure 10.16. The model has 640 nodes. The center node, m , is

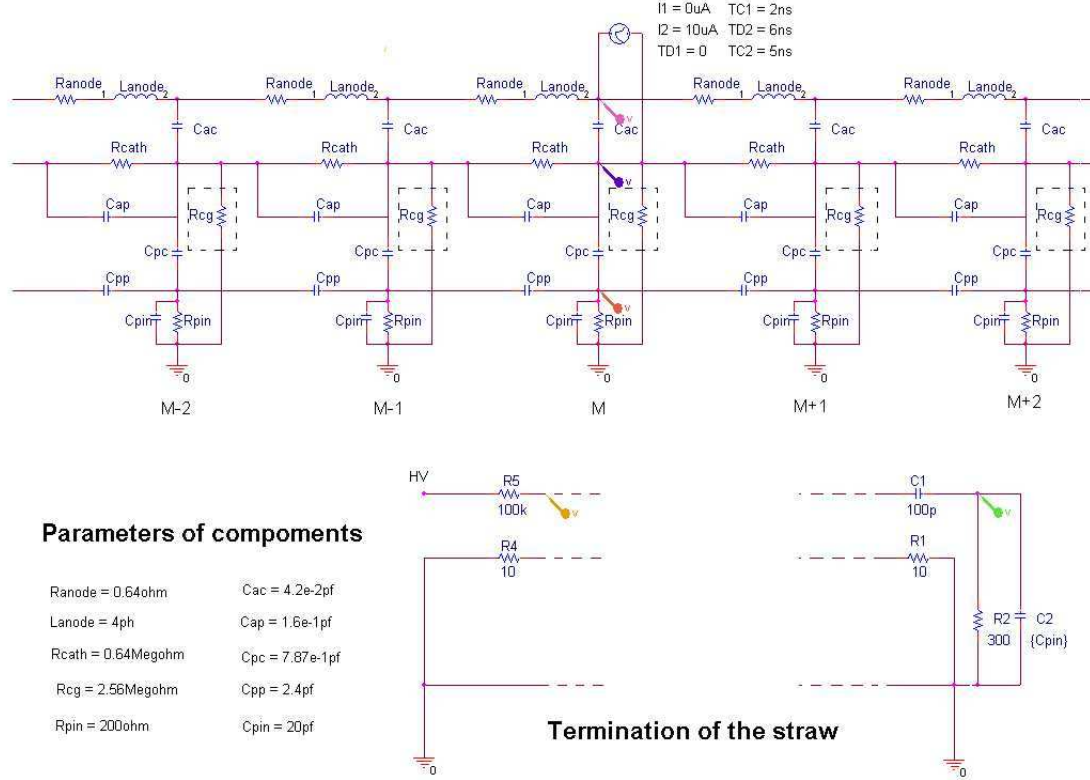


Figure 10.16: Equivalent circuit to study the effect of cathode resistivity on the strip signal.

injected with a current signal having 5 ns rise and 15 ns fall times. The strip signal is shown as a function of the cathode resistivity (a $640 \text{ k}\Omega$ value of R_{cath} is equivalent to $1 \text{ M}\Omega$ per square). The integrated charge on the strips, m , $m+1$, etc. is consistent with a static image model [145–147], indicating that ratio of charge on the $m \pm 1$ strips to the central charge is approximately 20%. The induced signal for various values of R_{cath} is shown in Figure 10.17. A cathode resistivity between $0.5\text{--}1.0 \text{ M}\Omega$ per square provides a sufficiently transparent foil for the straws. A summary of this work has been published [148].

The recharge time following the particle flash associated with the beam microstructure also limits the straw resistance. This has been studied using the circuit model




Figure 10.17: Induced signal on the strips as a function of the cathode resistivity. A 640 k Ω value of R_{cath} is equivalent to 1 M Ω per square.

above. The semi-circular region of the straw opposite to the induction strips must be grounded to reduce the total cathode resistance. If this is not done the straws cannot recover from the beam flash.

Another option is to use 30 μ thick straws made of carbon-loaded PEEK (poly-ether-ether-ketone). Since these tubes cannot be manufactured more than 1 m long, three tubes would be glued together to obtain the required 2.6 length.

10.3.6 Drift Gas

Electron Drift Properties in the Magnetic Field

Because of the expected rates, and particularly the required recovery time of the detector after the beam-flash, we intend to use a fast drift gas such as 80% CF₄ with 20% isobutane [149]. In addition, we plan to limit the gas gain to approximately 5×10^4 . Although this will also limit the spatial resolution of hits, our requirements of $\sigma = 200 \mu\text{m}$ are easily obtained. magnetic fields up to 2 T in CF₄/isobutane have been studied [150] in the laboratory. At 1 T, the Lorenz angle is $\approx 45^\circ$ at 1 keV/cm and $\approx 20^\circ$ at 4 keV/cm. The drift velocity in the drift direction varies from 70 $\mu\text{m}/\text{ns}$ at 1 keV/cm to 120 $\mu\text{m}/\text{ns}$ at 4 keV/cm. The magnetic field in the M μ 2e detector region is constant along magnet axis, 1 T, but the E field varies as a function of radial distance r . We expect the drift velocity along the radial direction to be 50 $\mu\text{m}/\text{ns}$ at 1 keV/cm and 110 $\mu\text{m}/\text{ns}$ at 4 keV/cm.

10.3.7 An Alternate Design

There are design and construction challenges in the longitudinal tracker.

1. The assembly of long straws is very flexible, and it may be difficult to stiffen them without introducing either significant complications or unwanted extra material.
2. The mechanical properties of a system of long straws made of different materials under the stress of vacuum may cause problem for stable and reliable operation.
3. Obtaining material with the correct resistivity for the outer straws is difficult.
4. The large number of inductive readouts are especially sensitive to noise and grounding issues.

As a result, another design, called the transverse tracker, has been suggested. This device is modular, easy to build, and uses only conductive straws without any inductive readout. It features many single layers of straws transverse to the magnetic field.

The issue is that inductive readout design has useful advantages for redundant pattern recognition, including

1. Precise pad-wire timing, simplifying the matching of pairs of coordinates into points.
2. Pulse-height correlations will also matching pairs of coordinates.
3. The three adjacent layers of straws helps reduce the left-right ambiguities inherent in straw tubes.
4. The trajectories will hit nearby straws.

The open issue is whether inferior pattern-recognition capability of the transverse tracker is still sufficient to identify the 105 MeV electrons with ample resolution and no high-energy tails. A preliminary investigation has suggested that the answer is yes [151], and we await a more realistic Monte Carlo simulation to convince us that the transverse tracker hardware will do the job.

10.4 Electromagnetic Calorimeter

The electromagnetic calorimeter is used to form the experimental trigger, to provide a reference time of about 1 ns and space point resolution of about 1 cm for track reconstruction, and a redundant measurement of the electron energy. The calorimeter must operate in a high-rate environment with a radiation dose of about 160 Gy. Its response must be sufficiently fast that when the detector is turned on 700 ns after the start of the prompt flash, little residual light remains. The calorimeter must be able to operate with no untoward effects in the 1 T magnetic field of the detector solenoid. It should have little acceptance for the enormous rate of decay-in-orbit elections coming

from the stopping target, and large acceptance for ~ 100 MeV electrons coming from the stopping target. Its energy resolution must be good enough to enable the trigger to operate at a rate of $\lesssim 1$ kHz.




Figure 10.18: Isometric view of the Mu2e detector showing the straw tracker and the electromagnetic calorimeter, along with the track of a conversion electron.

In what follows we describe a calorimeter design that meets the above specifications.

Table 10.6: Calorimeter parameters.

Number of vanes	4
Vane size ($\phi \times r \times z$)	$3 \times 30 \times 135 \text{ cm}^3$
Inner / outer radius	39.0 cm / 69.0 cm
Crystal type	PbWO_4
Crystal size	$3.0 \times 3.0 \times 12.0 \text{ cm}^3$
Crystal length	12.0 cm (13.5 X_0)
APD per crystal	2
APD size	$1.3 \times 1.3 \text{ cm}^2$
Crystals per vane	$10 \times 45 = 450$
Total crystals	$4 \times 450 = 1800$
Total APDs	3600
Operating temperature	-24°




Figure 10.19: Calorimeter front view. Electrons spiral in a clockwise direction.

10.4.1 Calorimeter Description

The calorimeter is situated at the rear of the Detector Solenoid, just behind the magnetic spectrometer, and in front of the muon beam dump. It consists of 1800 PbWO_4 crystals arranged in four vanes, as shown in Fig. 10.18 and Fig. 10.19. The PbWO_4 crystals are $30 \times 30 \times 120 \text{ mm}^3$ in size, or $13.5 X_0$ long. (The electrons impacting the calorimeter travel in helical orbits with incident angles ranging from 40° – 65° ; the mean angle being 55° . Hence, the minimum effective material depth is $14.9 X_0$ and the average is $16.5 X_0$.) Each crystal is read out by two avalanche photodiodes. Each vane has 450 crystals arranged in a 10×45 array ($r \times z$). The innermost face of the calorimeter is at a radius of 39.0 cm, roughly that of the innermost straws; the outer radius is at 69.0 cm. Each vane is 135 cm long. Like the straw-tube tracker, the calorimeter has been placed such that the bulk of the decay-in-orbit electrons — those with energies below 55 MeV — have trajectories that lie inside its inner radius and hence are not detected, as shown in Fig. 10.4.1 which gives the geometric acceptance as a function of electron energy. The acceptance for conversion electrons coming from the stopping target is essentially the same as that of the straw-tube tracker: all emanated at angles above $\pm 60^\circ$ from the solenoid axis are accepted. To shield the calorimeter from the intense flash of soft gammas at the beginning of each

microbunch, a 1 cm deep lead shield will be placed in front of the first column of crystals in each vane (at the front z face of the calorimeter).




Figure 10.20: Incident angle of electrons at the calorimeter. An angle of 90° in perpendicular to the module faces.




Figure 10.21: Calorimeter geometrical acceptance. The acceptance for conversion electrons is similar to the straw-tube tracker at $\sim 50\%$.

10.4.2 Crystal Choice

The electromagnetic calorimeter employs lead-tungstate (PbWO_4) crystals read out by avalanche photodiodes (APDs), providing an energy resolution of $\sigma \sim 7 \text{ MeV}$ (@ 100 MeV) and a position resolution of $\sigma \sim 1.0 \text{ cm}$ (in r and z). We have chosen PbWO_4 for the calorimeter because: it is fast, it has a short radiation length, and it will survive the anticipated radiation dose. The PbWO_4 crystal parameters are given in Table 10.7.

In the last fifteen years considerable research on the manufacture and performance of PbWO_4 has been done. A PbWO_4 calorimeter has been built and operated for the PrimEx experiment [152]; much larger PbWO_4 calorimeters have been built for the CMS [153] and ALICE experiments [154, 155]; and the PANDA experiment is poised to begin fabrication of a PbWO_4 calorimeter designed to detect gammas with energies similar to that of the conversion electrons of *Mu2e*.

New scintillating crystals with much better light yields have been developed, such as $\text{LSO}(\text{Ce})$, which emits at 402 ns with a 40 ns decay time. In tests done by R.-Y. Zhu at Caltech a light yield 200 times that of PbWO_4 was measured, providing a $\sim 10\%$ FWHM energy measurement of 0.51 MeV gammas from ^{22}Na [156]! Radiation damage studies are promising. At the present time these crystals are too expensive, however there is hope that the cost will be reduced in the future. The use of such

Detector

crystals for the *Mu2e* calorimeter would significantly improve the energy resolution, which although not needed for triggering, would decrease backgrounds by providing a redundant measurement of the electron energy with comparable resolution to the straw tracker. It would also alleviate the need for a crystal cooling system.

Table 10.7: Calorimeter crystal parameters [157].

Crystal	PbWO ₄ (Y)
Density	8.28 g/cm ³
Radiation length	0.89 cm
Molière radius	2.0 cm
Interaction length	20.7 cm
dE/dx	13.0 MeV/cm
Refractive index ^a	2.20
Hygroscopicity	None
Luminescence	420 nm
Decay time	6 ns
Light yield ^b	1.2%
$d(\text{LY})/dT$	−1 to −3%/°C

^aAt λ_{max} .

^bNaI(Tl) = 100%.

The relatively low light yield of PbWO₄ — ~20 photoelectrons (phe) per MeV for small samples and about half that for large crystals — requires cooling the crystals. This increases their light yield: in going from 20°C to −23°C gains of a factor of three have been obtained [158]. Unfortunately, recent studies show that the radiation hardness decreases for cooled crystals [158].

Much research has been done to increase the light yield of PbWO₄ by adding additional dopants, such as molybdenum and lanthanum. Unfortunately, although the light yield can be increased, it is usually in conjunction with an additional slow component with a time scale greater than 1 μ s.

Members of the PANDA experiment, in collaboration with the Bogoroditsk Technical Chemical Plant, have carried out an extensive R&D program to produce higher light-yield PbWO₄ crystals with no additional slow component to their response. [157]. The crystals, which they have dubbed PWO-II, are grown from high-purity raw material and doped with lanthanum and yttrium in concentrations of $\lesssim 20$ ppm. They have succeeded in producing large crystals with a light yield of 20 phe/MeV at 18°C, double that of similarly sized crystals produced by Bogoroditsk for the CMS endcap calorimeter [157].




Figure 10.22: Energy resolution for incident photons measured by a 3×3 array of $20 \times 20 \times 200 \text{ mm}^3$ PWO-II crystals. (From Ref. [159]).

Figure 10.22 shows the energy resolution of a PWO-II array from a test beam study by members of the PANDA collaboration [159]. At 100 MeV they measure a resolution of $\sigma = 5.5 \text{ MeV}$ with the crystals cooled to 10°C , and $\sigma = 3.91 \text{ MeV}$ when cooled to -25°C . If these resolutions can be obtained with mass produced crystals then the use of PWO-II without cooling would be considered, alleviating the need for a complex cooling system, although the APDs would still be cooled to reduce their noise, and increase their gain. Figure 10.23 shows the time response of the scintillation light from PWO-II crystals for several different operating temperatures. Like normal PbWO_4 , the PWO-II crystals respond quite quickly to the initial excitation: over 95% of the total light is emitted within the first 50 ns.




Figure 10.23: Time integral of the light from PWO-II crystals. (From Ref. [160]).




Figure 10.24: PbWO_4 crystal, 2 APDs, and a preamplifier fabricated by NYU for MECO prototype tests.

Currently there is sufficient capacity worldwide to produce the 1800 crystals needed by *Mu2e* in a year's time, with large capacity plants in Shanghai, China (SICCAS [161]), Bogoroditsk, Russia (BTCP [162]), and Murmansk, Russia [163], as well as other sites [164, 165]. To date only BTCP has the capability to manufacture PWO-II crystals. Because of this our baseline design assumes the use of cooled “normal” PbWO_4 crystals.

10.4.3 Calorimeter Readout

Because of the high magnetic field of the Detector Solenoid (1 T) large-area avalanche photodiodes (APDs) will be used for the photodetector.* APDs have several other favorable properties besides their ability to operate in high magnetic fields: they are compact, $\sim 200\mu\text{m}$ thick; have high quantum efficiencies, 70–80% at wavelengths where the PbWO_4 outputs most of its light; and are relatively inexpensive. Large area APDs, $\sim 1.0 \times 1.0 \text{ cm}^2$, have recently become available and are manufactured by Radiation Monitoring Devices (RMD) [166], Advanced Photonics [167], and Hamamatsu [168].

APDs suffer from several problems that, for example, photomultipliers do not have: they have low gain, ~ 100 ; the gain has a large temperature dependence, $\sim 2\%/^\circ\text{C}$; and they are particularly sensitive to charged particles. Their low gain dictates the use of very low noise amplifiers, and operation at low temperatures. Gain

*We are keeping tabs on progress on novel photodetectors that are currently under development, such as single-photon avalanche photodiodes that have recently been introduced by several companies, although not in sizes large enough for *Mu2e*.

stability requires strict control of the temperature and voltage of the APD. Because of their anomalously large response to charged particles, such as charged particle punch-through from the electromagnetic shower (the so-called nuclear counter effect), two APDs will be used to read out each crystal, which allows the anomalous responses to be identified and eliminated both at the trigger level and in the off-line analysis. It also doubles the light yield, improving the energy resolution by $\sqrt{2}$.

The APDs will be glued to the back of the PbWO_4 crystals. They will be cooled, along with the PbWO_4 crystals, to -24°C , decreasing their dark current by a factor of 50. Cooling the APDs also increases their gain: CMS has measured an increase of $2.2\%/^\circ\text{C}$ at a gain of 50 [169]. Their bias voltage will be about 300–400 V.

Several $1.3 \times 1.3 \text{ cm}^2$ APDs manufactured by RMD were procured and tested by the New York University (NYU) group for MECO. They were glued to $3.0 \times 3.0 \times 14.0 \text{ cm}^3$ PbWO_4 crystals cooled to -24°C , and tested with cosmic ray muons [170] (Fig. 10.24). With an APD gain of 200 and APD capacitance of 130 pF, and using a low-noise preamp of their own design, they measured a dark current of 10 nA, giving an equivalent noise of 0.7 MeV. This is on the order of the estimated pile-up noise of 0.9 MeV. The measured number of photoelectrons was 38 phe/MeV with two APDs. From these cosmic-ray tests they estimated an energy resolution of $\sigma = 4.1 \text{ MeV} @ 100 \text{ MeV}$.

The APDs will be read out by charge-sensitive preamplifiers located near the APDs. A preamplifier was designed, built and tested for MECO by the NYU group. A line driver on each preamp card takes analog signals outside the cryostat vacuum volume. A shaper-amplifier with $\text{RC} = 50 \text{ ns}$ further amplifies the signals before they are digitized by the wave-form digitizing system. In order to minimize the footprint of the preamplifier we are considering implementing it as an ASIC.

The strong dependence of APDs on the bias voltage ($\sim 2\text{--}3\%$ per volt) dictates a stable low-noise power supply. The system developed by CAEN for the CMS electromagnetic calorimeter meets these requirements [171].

10.4.4 Calorimeter Cooling

The crystal and APD region of each vane will be enclosed in a copper heat sink cooled to -24°C by welded coils connected to an outside chiller by insulated conduits. The copper box side is thin (1–2 mm) on the crystal face where the electrons enter, and has penetrations for the APD leads on the readout side. There is no steady heat load (APD dissipation is negligible) but the cooling-down times will be long since heat transfer in the vacuum of the Detector Solenoid will be inefficient.

The preamps (with a dissipation $\sim 200 \text{ mW}$ per channel) will be enclosed in their own copper heat sink box, and operate at the same -24°C temperature. Individual preamp ground shields will be “heat-shorted” to the heat sink box. The total dissipation of $\sim 100 \text{ W}$ per vane will be extracted by a separate outside chiller. If the conduction cooling is insufficient to keep components at a reasonable temperature,

a gas enclosure for the preamps could be used, although this would complicate the design.

The gain variation of PbWO_4 is typically -1 to -3% per degree C, depending on the crystal, operating temperature, and manufacturer [172]. The APD temperature variation is 2.5 – 3.5% per degree C. Maintaining crystal temperature to about ± 0.5 °C will be sufficient to prevent large trigger rate fluctuations from this source.

As mentioned above, one advantage of using PWO-II or one of the new high-light-yield crystals is the fact that they would not have to be cooled. In the event that the APDs would still have to be cooled thermoelectric coolers would be used.

10.4.5 Calorimeter Mechanical Support

The 450 crystals in a vane will be supported by a square-cell honeycomb of carbon-epoxy composite laminate (~ 0.2 mm wall thickness) similar to that being used by CMS for its electromagnetic calorimeter [153]. A section of the honeycomb is shown in Fig. 10.25. This dead material dilutes the active mass of the calorimeter by 0.3% . The square honeycomb structure is contained in, and attached to, the metal box that forms the heat sink for cooling the detector. The outer-radius element of the vane-box is a heavy plate that is the attachment for the calorimeter vane.




Figure 10.25: Carbon-epoxy square-cell honeycomb support for the PbWO_4 crystals.




Figure 10.26: Front view of the calorimeter cage, cart, and rails supporting the four calorimeter vanes.

Each of the four vanes of the calorimeter is attached to a rigid rolling cage that is entirely outside the detector region's maximum radius (Fig. 10.26). The cage allows the calorimeter to be rolled in and out for installation and servicing on rails attached to the inner wall of the cryostat.

10.4.6 Calorimeter Radiation Dose

GEANT Monte Carlo simulations have been done to estimate the expected radiation dose to the calorimeter [48]. These simulations indicate an expected dose of about 160 Gy in one year of running, with a dose rate of about 0.03 Gy/hr. The calorimeter irradiation is highly non-uniform since the bulk of the dose comes from soft gammas (~ 0.1 MeV), that deposit their energy in a very thin layer on the surface of the crystals. The front face of each vane is shielded with 1 cm of lead, which reduces the dose by a factor of 25–30. The dose from neutrons, which interact uniformly throughout the entire detector, at ~ 1 Gy is negligible, as is the energy deposited from electrons and protons.

It is known that radiation only affects the transmission, not the scintillation mechanism in PbWO_4 . At these dose rates we expect only about a percent reduction in the light output [173].

Radiation damage from neutrons and protons displacing atoms from sites in the crystalline lattice increases the dark current in APDs. The effect is reduced by cooling the crystals. The neutron flux at the APDs is several orders of magnitude greater than the proton flux. With an estimated total neutron flux of 5×10^{10} n/cm² we do not anticipate a significant increase in the dark current.

10.4.7 Calorimeter Performance

The calorimeter trigger is described in detail in Chapter 12. The trigger is formed by dividing each calorimeter vane into overlapping 5×5 supercells, summing the energy in each supercell, and requiring that the energy be greater than a preset threshold. A GEANT Monte Carlo simulation of the calorimeter trigger was performed, where the energy in each cell was smeared by electronic noise, photostatistics, and pileup. The results are shown in Fig. 10.27, where an energy resolution of $\sigma(E) = 5.4$ MeV at 100 MeV has been assumed. This simple threshold trigger with good acceptance produces a trigger rate within the acceptable bandwidth of the data acquisition system (DAQ). An increase in the energy resolution to $\sigma(E) = 7.9$ MeV doubles the trigger rate, but it remains within the DAQ bandwidth limits.

10.4.8 Calorimeter Calibration

The absolute calorimeter energy response must be continually monitored, and the trigger threshold adjusted, in order to keep the trigger rate at $\lesssim 1$ kHz. To get an idea of the scale by which the calibration must be determined, a change in the calorimeter energy threshold from 80 MeV to 75 MeV increases the trigger rate by a factor of three [174]. In addition, the cell-to-cell gain variations must be equalized at the hardware level on the same scale. Sources of short and long term variation include crystal and APD radiation damage, APD gain drift, temperature variation of the crystals, APDs, and electronics, and APD voltage fluctuations.




Figure 10.27: The calorimeter trigger acceptance and trigger rate as a function of the trigger energy threshold. The blue curve is the trigger acceptance for conversion electrons with accepted trajectories in the tracker. The red curve shows that the trigger rate, dominated by DIO electrons, is acceptable over a wide range of threshold energies. Note that above a 80 MeV threshold there is a rapid fall off in acceptance. Energy in each cell was generated by a GEANT3 simulation and includes electronic noise, pileup, and an energy resolution of 5.4 MeV at 100 MeV.

A sample of PbWO_4 crystals, tested at NYU using a ^{137}Cs 0.662 MeV gamma source, gave an average crystal non-uniformity (maximum to minimum gain difference) of 3.3%. The range for the crystals was 1.8–4.8%. The large extent of the electromagnetic showers in the crystals (compared to that of ^{137}Cs 0.662 MeV gammas) should mitigate the intra-crystal non-uniformity, and this nonuniformity is not expected to present a problem. The largest crystal-to-crystal differences, which will be corrected for by properly setting APD voltages, was 11% in the sample tested.

A calibration system will be designed to measure the absolute gain and the cell-to-cell variations, allowing the system hardware to be periodically tuned for long term variations, and the short time scale fluctuations to be corrected offline. Voltage and temperature monitoring and control, although not part of the calibration system per se, will be implemented to keep fluctuations to within allowable levels.

The absolute gain, uniformity, and temperature dependence of individual crystal APD assemblies will be tested prior to assembly of the calorimeter. Although this is feasible using cosmic rays, a radioactive source or an electron accelerator is preferred. APD voltages will be set to equalize gains at the start of the experiment. The original calibration measurements will be used to set the initial energy calibration constants.

A flasher system will be used to monitor gain changes of the crystals and their associated APDs. Laser, LED, and Xenon flasher systems that have been developed for other similar crystal calorimeters are being evaluated [175, 176]. Stabilities of

better than 0.1% (r.m.s.) have been reported [175], which is better than what we need.

The absolute energy calibration can be provided by measuring the front part of the calorimeter response to electrons above 80 MeV, whose energy is measured to about 0.20 MeV by the tracker, and which arrive at a rate of about 200 Hz. The energy calibration for the back part of the calorimeter, which is less often hit by these electrons, can then be determined from relative calibrations. If necessary, the Detector Solenoid magnetic field will be lowered to provide a higher rate beam of lower energy electrons which will go through the tracker and calorimeter.

There will thus be two almost independent systems to monitor the gain of most of the calorimeter. Relative APD/crystal response changes will be tracked on a short time scale with a flasher system. Absolute gain measurements from incident electrons will be compared with the results from the flasher system to study the systematics involved. The gains of cells near the back of the calorimeter, for which only the flasher system and perhaps cosmic rays are available, will then be corrected as needed. APD voltages will be periodically tuned to keep the hardware trigger threshold constant to about 2 MeV

Individual cell temperatures will be monitored by temperature sensors connected to the back of the crystal (see for example the Rugby Ball used in GRAAL [177]). Once a trigger occurs, the measured gain versus temperature curves will be available for use on line, and will be used off-line to fine tune the calorimeter energy measurement.

Cosmic Ray Shield

COSMIC ray (CR) muons can produce backgrounds from delta ray production in the target and detector region, muon decay-in-flight electrons, and muons mistaken as electrons. These backgrounds are defeated by both passive and active CR shields. The layout of the CR shield is shown in Figure 11.1 and Figure 11.2.

The Cosmic Ray flux under ground falls off with depth, as shown in Figure 11.4. It is likely that the apparatus will be located about 20' below ground, but the experimental area design has not been finalized. Hence, in the following discussion, we conservatively assume an apparatus placed on the surface, as were all previous experiments searching for muon to electron conversion.

The cosmic ray background rate will be monitored between spills and when the beam is off, allowing a direct measure of the level of background. The cosmic ray background rates will be measured as soon as the detector and detector solenoid are in place, allowing modifications to the shield if necessary.

11.1 Cosmic Ray Background Rate Calculation

The calculation is based on measured cosmic fluxes from the literature and a GEANT simulation of the shielding and detector. Muons dominate the flux of particles penetrating any significant amount of shielding. Their energy spectrum at sea level is essentially flat below 1 GeV, and then falls with a power law approximately given by $E^{-2.5}$, with E in GeV. This is shown in Figure 11.3

The angular distribution is approximated by $dN/d\theta \sim e^{-1.43\theta}$. The muon flux is about 60% positive. For decays and incident muons mistaken as electrons, only negative muons contribute. For delta rays and pair production in materials in the target and detector region, both μ^+ and μ^- contribute. The calculation accounts for the following sources:

- Muons penetrating the shielding and decaying in the detector solenoid.
- Muons penetrating the shielding, interacting in the target, detector and other material, and making electrons.

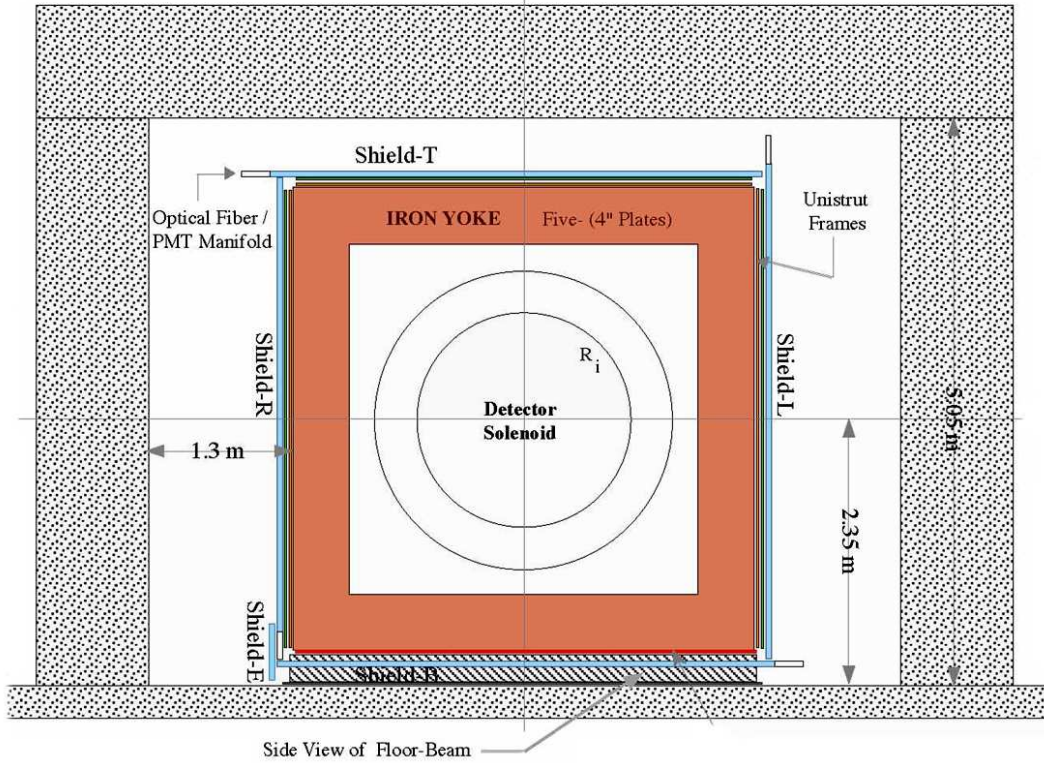


Figure 11.1: End view of Cosmic Ray Shield

- Muons penetrating the shielding, scattering in the target or other material, and the muon being mistaken for an electron.
- Muons interacting in the shielding and producing other particles (photons or hadrons) which then interact in the detector to produce electrons. These events may not deposit energy in a veto counter.

The shield configuration simulated consists of 0.5 m of steel surrounding the detector solenoid, followed by a triple layer of scintillation counter detectors, and 2.0 m of heavy concrete shielding blocks. The effect of the magnetic field in the steel shielding has not been included in the background calculation; it should reduce the particle flux inside the solenoid by curling up low energy muon tracks.

The probability of particles penetrating the shielding was calculated by simulating muons incident on the shielding normal to the surface. This underestimates the attenuation since it underestimates the average path-length. The flux of particles exiting the shielding consists of photons, muons, electrons, positrons and lesser numbers of low energy hadrons. The differential intensities for γ , e^\pm and μ^- fluxes emerging from the nominal shielding are obtained. These fluxes were used as input to the calculation of the probability of producing a 100 MeV electron from cosmic rays.

Cosmic Ray Shield




Figure 11.2: Top view of Cosmic Ray Shield




Figure 11.3: Flux of cosmic ray muons at the earth's surface (1000 g/cm^2) as measured by the CAPRICE94 experiment (Ref. [178]).




Figure 11.4: Cosmic ray flux vs. zenith angle at the surface and depths of $10'$ and $20'$

Essentially all particles penetrating the shielding resulted from processes that deposited energy in the scintillation counter between the concrete and steel. A small flux of photons emerged without depositing energy in the scintillator. They resulted from bremsstrahlung by a μ which then ranged out before passing through the scintillator. The probability of getting a photon with energy exceeding 100 MeV is approximately 2×10^{-6} . This contributes a negligible amount to the potential background from photons caused by processes that did deposit energy in the scintillator, assuming an inefficiency in the scintillator for detecting a penetrating charged particle is 10^{-4} .

To estimate the total background, the penetrating flux of γ , e^\pm , and μ^\pm was caused to impinge on the volume inside the detector solenoid. Particles were generated on the interior of a cylindrical shell (the magnet coil) according to the calculated flux of particles penetrating the shielding, and weighted by the cosmic ray flux as a function of zenith angle. The simulation of the resulting propagation and interactions was done including the effect of the magnetic field. All kinematic properties of all particles which intersected any part of the tracking detector were recorded, and the following selection criteria imposed:

- The particle charge is negative.
- The momentum is in the range $100 \text{ MeV}/c < p < 110 \text{ MeV}/c$.
- The number of clusters in the tracking detector is more than 3.
- The pitch angle is in the range $45^\circ < \theta_p < 62^\circ$.
- The closest distance to the solenoid axis at the stopping target is less than 10 cm.

- The track has fewer than 3 missing hits in the fitted trajectory.

The selection on the pitch angle accounts for the fact that electrons produced in the stopping target have allowed values in this range. The last selection criterion eliminates electrons originating in the middle of the detector, for which the fitted trajectory is predicted to pass through an octagon or vane detector 3 or more times without it having done so. The current event selection criteria are more stringent than these.

A total of 9×10^7 particles were generated, distributed according to the particle type and momentum distributions calculated as described above. Taking the duty cycle of the accelerator to be 50%, the detection time for conversion electrons to be 900 ns each $1.7 \mu\text{s}$, assuming that we veto cosmic ray induced events using the veto scintillation counter with an efficiency of 0.9999, and accounting for the illumination area, this corresponds to approximately 200 times the nominal *Mu2e* running time of 10^7 s.

A total of 24 particles (weighted as described above) satisfying the above selection criteria was found. Most of these events were caused by particles produced downstream of the tracking detector (in the electron trigger detector, for example), moving upstream through the tracking detector, reflecting off the B field, and then moving downstream through the tracking detector a second time. All but 3 (weighted) events have at least 4 clusters in the tracking detector caused by the backward moving particle. This background is suppressed since we will detect and reconstruct the backward moving track, out of time in the tracking detector. Additionally, events produced in the electron trigger detector will have energy deposited at a time and position which can be deduced from the tracking information. We assume that these events will be vetoed with high efficiency. Of the particles with fewer than four clusters produced by the backward moving track, three were muons and can be eliminated by a time of flight requirement between the tracking detector and electron trigger counter. Three others had more than 200 MeV deposited in the electron trigger counter, and can be eliminated by a clean requirement in a small time and position window. One had only 27 MeV deposited in the electron trigger counter. The three remaining events were a μ decay upstream of the tracking detector and two delta rays, one produced in the target and one in the straw detector. From this result the cosmic ray background is predicted to be 1.6 events at 100 times the nominal 2×10^7 s running time, or a background of 0.016 events.

11.2 Beam Induced Rates

We have used MARS to simulate the counting rate in the Cosmic Ray Veto counters arising from neutrons produced in the muon stopping target. Principally this comes from the 2 MeV γ produced by neutron capture on hydrogen.

In the annular region immediately aligned with the stopping target, the rate for $E_{\text{deposited}} > 200$ keV in the inner layer of 1-cm thick scintillator is less than 5×10^6 for 2.5×10^{11} stopping muons. Elsewhere, it is negligible. We believe that that this is acceptable, but the Monte-Carlo studies are continuing. We have *not* simulated the effect of the general neutron background in the experimental area on the Cosmic Ray Veto counters, but as part of our prototyping effort we plan to test a module in an accelerator environment.

11.3 Passive Cosmic Ray Shield

Figures 11.1 and 11.2 show a configuration of concrete vertical blocks and horizontal roof beams. Vertical blocks, $5.28 \times 1 \times 1$ m³, stand back one meter from the inner steel box to afford access to the active scintillator shield. Roof beams, $8.26 \times 1 \times 1$ m³, top off the structure.

Concrete shielding blocks are fabricated from a mixture of special aggregate mineral material and cement, and reinforced with steel rods. Key ingredients in the aggregate are hematite, magnetite, ilmenite, and steel. The density of aggregate alone is $5.0 - 5.6$ g/cm³. Firms such as Universal Minerals, Inc (UMI) typically ship the aggregate mixture to a nearby concrete block maker company which casts the blocks in a mold. UMI uses coarse and fine mineral aggregate material which has a 67% Fe content. This blend of cement and aggregate has a density of 4.0 g/cm³ or higher. With reinforcing rod included the block maker's cost is typically equal to that for the aggregate material. The quantity of vertical blocks required are 32 for the side walls, and 15 for the front plus back walls. Roof beams would total 18 blocks.

The steel enclosure shown in Figures 11.1 and 11.2 provides both a return path for the detector solenoid field, as well as both a passive external shield against cosmic rays and an internal shield against fast capture neutrons, $E > 0.5$ MeV, produced at the muon stopping target. The shape of the steel enclosure is rectangular with a wall thickness of 0.5 meters. It spans the length of the detector solenoid for about 12.4 meters. The plan is to cap off the ends of the steel box with stainless steel walls of 0.5 meters. Penetrations are limited primarily to ports totaling a size of 400 cm² for tracker and calorimeter cable runs, and for a smaller opening to run solenoid cables.

11.4 Active Cosmic Ray Shield

To achieve the design level of 10^{-4} for cosmic ray rejection for so large a detector and to do so economically we need to build upon proven developments in large scintillator array detectors. For this reason our expectation for good shield performance is based in large part upon results that have been obtained by the MINOS collaboration at Fermilab. They have developed long, high performance scintillator elements at modest cost [179, 180]. We have benefited from early work on embedded-fiber scintillation

Cosmic Ray Shield




Figure 11.5: Sectional view of a full Cosmic Ray Shield module mounted to the Detector Solenoid return yoke

detectors carried out by the TJNAF Detector Group in the early 1990s [181]. We will attach cosmic ray veto modules to the outside of the solenoid return yoke, as shown in Figure 11.5.

The basic idea is to circumvent the short attenuation length for a scintillator's blue light by bonding wavelength-shifting (WLS) fibers along the scintillator's long axis. This is illustrated in Figure 11.6 which shows a cross-sectional view of three fibers




Figure 11.6: Cross-sectional view of scintillator with embedded waveshifting fibers.

embedded in a scintillator coated with TiO_2 reflector. The objective is to obtain an adequate efficiency for a set of fibers to wave-shift a primary blue scintillation photon to a green photon that then travels internally along the hit fiber with a long attenuation length. This profile of a $10 \text{ cm} \times 1.0 \text{ cm}$ scintillator bonding six 1.5 mm diameter WLS fibers in its surface represents our current design for the *Mu2e* Active Shield. This is motivated primarily by the similar profile of the TJNAF scintillator which had five, 2.0 mm diameter fibers embedded in a $10 \text{ cm} \times 1.0 \text{ cm}$ scintillator and wrapped in Tyvek [181]. Their studies demonstrated a small 2% variation of signal response when moving a radiation source transverse to the fibers. Some signal roll off, however, was observed within 0.5 cm of the edge. When the number of incorporated fibers increased from one to five the signal dropped below a linear response by 50% in a plot of fiber light output versus the number of fibers. This effect indicates a good conversion efficiency to WLS photons. The study involved the Bicon fiber BCF92 bonded to various Bicon scintillators which had been grooved with a multi-tool carbide cutter. The PMT used in these studies was the Hamamatsu R580-17, a green-enhanced bi-alkali tube.

MINOS scintillator strips have a profile of $4.1 \text{ cm} \times 1 \text{ cm}$ and nominal lengths of 4.5 m and 8.0 m . The shorter strip is readout at one end; the longer one at both ends by a single embedded 1.2 mm diameter Y11 WLS fiber from Kuraray Corp. These signals are processed by a multi-anode, bi-alkali PMT, the Hamamatsu R5900U-00-M16. Performance variations among the 8 m strips for summed fiber CR signals have ranged from 6 – 9 photoelectrons in a steadily improving fashion. Attenuation length measurements for CR-induced fiber photons have achieved 6-8 m.

11.5 Scintillator Strips

The quantity of scintillator needed for the *Mu2e* CR Shield is mainly dictated by the outer area of the 4.4 m \times 4.4 m \times 13.7 m steel enclosure. In a design that orients scintillator strips transverse to the detector axis good overlap at the corners of the steel box will be provided by long strips with a nominal length of 4.7 m. Figures 11.1 and 11.2 show these scintillators only in edge view. A profile of 10 cm width and 1 cm thickness has been chosen to complete the strip dimensions. This width conveniently allows adjacent strips to be overlapped in a staggered pattern. This pattern eliminates a gap between strips which could reduce the veto efficiency significantly if the long extrusions show curvature along their sides. In addition, it compensates for any roll-off in performance near the edge of strips.

Altogether this amounts to 800 m² of scintillator which requires a minimum extrusion run of 9 km. Cut to length there would be about 2100 strips. In an attempt to maximize the number of photoelectrons resulting from a minimum ionizing particle passing through a scintillator layer we intend to embed six WLS fibers of 1.5 mm diameter along one surface in the manner studied by the TJNAF Detector Group [181]. This requires that the extrusion die for strip production be designed to produce six accommodating grooves along one surface spaced symmetrically at 1.667 cm intervals. During this same extrusion process a thin coating of TiO₂ will be applied to provide maximal surface reflectivity. The standard design with surface grooves is at the top; the bottom design showing six mid-plane holes for fibers.

A number of development studies will begin with the aim of projecting a level of performance for our prototype detector of six, 1.5 mm diameter fibers bonded in scintillator surface grooves. This will involve non-extruded scintillator from Bicron, and WLS fibers from Bicron and Kuraray. Comparisons between the Hamamatsu phototubes will be made. Both triggered cosmic rays and Bi-207 conversion electrons will be used to determine photoelectron response. A study will also be made of any roll off in performance near the edge of the strip detector. Response to ⁶⁰Co gammas will determine this source's suitability as a monitoring component for the CR Shield detector. Performance of the prototype module will also be tested in a beam such as the Fermilab MTBF.

11.6 Wavelength Shifting Fiber

Given that our longest scintillator strips are 4.7 meters long we intend to implement a single-end readout of the WLS fibers. This will facilitate an arrangement of strips which can best give a full, two-layer veto coverage of the detector solenoid. As has been noted MINOS has single-end readout of the 4.5 m strips in their Near Detector. Figure 11.6 depicts the smallest unit of three adjacent strips in one of the veto layers. This shows their staggered overlap in two views and the run of the WLS fibers. At the readout end these fibers are terminated 1-to-2 cm beyond the end of the strips.

This then sets the length of fibers at 4.72 m. (It should be noted that strip lengths are tied directly to the diameter of the detector solenoid which has yet to be finalized). Employing optical connectors, clear optical fibers of 1.5 mm diameter then carry light from the 18 waveshifting fibers of the three strips through a 190 degree low-loss bend to a PMT. Such a system will require 60 km of WLS fiber. The amount of optical fiber needed is about 7600 m. Studies are planned to compare detector performance with and without a far end fiber reflector. For a reflector the mean time for direct and reflected pulses is about 24 ns.

MINOS has compared 8 m long strip performance for Y11 fibers of 0.9, 1.0, 1.1, 1.2, and 1.5 mm diameter. The relationship between signal and fiber radius r was found to be $\text{signal} = kr^{1.4}$. Simply scaling the 1.2 mm MINOS results to those for a 1.5 mm diameter fiber suggests a signal enhancement of 36%. In addition there is reason to expect further signal improvement for our multi-fiber configuration given that MINOS studies showed they could double their signal if two 1 mm diameter fibers, spaced by 2 cm in separate grooves, were to replace a 1 mm fiber in a single groove. Based upon the JLab studies [181] of response versus fiber number our proposed arrangement of six fibers should give both an enhanced signal and a highly uniform response across the 10 cm width of the strips. Both references [179] and [180] have shown that using an adhesive to bond a fiber in a groove doubles signal yield relative to simply leaving an air interface. Bicon BC600 and Epon 815 have been shown to perform well in this regard. We intend to further determine if applying a vacuum to the two-part adhesive mixture provides an additional benefit by removing air trapped in the mixture. We already know from preliminary bonding tests with Bicon adhesive that numerous, small air bubbles can appear in the bonding layer if this is not done. Attempts to use a UV adhesive to bond WLS fibers would have the complication of the fiber absorbing the UV. We note also that MINOS has shown a 10% enhancement by covering each bonded fiber with a strip of reflective Tyvek.

11.7 Photomultiplier Tube and Signal Response

The MECO group considered two Photomultiplier Tubes (PMTs), both exhibiting a bi-alkali response. These are the Hamamatsu R580-17 and the Hamamatsu R7400U metal package PMT, which is available in 4-, 16-, and 64- anode versions. More recently, Hamamatsu provides the 7564B, a 64-anode MAPMT.

Recently there has been considerable research into novel photodetectors, such as Geiger-mode avalanche photodiodes (G-APD), which have better quantum efficiency than photomultipliers, similar gain, a smaller footprint, lower power consumption, and the potential of lower cost. They suffer from optical crosstalk and a large excess noise factor. Although our baseline design uses photomultipliers, should low-cost G-APDs become available that meet our specifications we would certainly use them.

Cosmic Ray Shield

1. Hamamatsu R580-17 PMT :

This is a standard 10-stage tube of 34 mm active diameter which operates at 1750 V and has a rise time of 2.7 ns. The label “17” indicates that its photocathode is green-enhanced to match the spectral emission of WLS fibers such as the Bicron BCF-92. It was the tube used in reference [181] studies for their multi-fiber investigations.

2. Hamamatsu R7400U PMT :

The R7400U series is the world’s smallest PMT mounted in a T0-8 metal package. It is an 8-stage electron multiplier using metal channel dynodes. The metal package has a 15 mm diameter and 10 mm length, with an effective photocathode diameter of 8 mm. This can accommodate a bundle of up to 19 fibers 1.5 mm in diameter. A gain of 7×10^5 is achieved at a nominal HV of 800 V. The compact dynode structure accounts for a very fast response both in rise time (0.78 ns) and fall time (1.15 ns), as well as a short transit time spread of 0.23 ns. The R7400U is the basic bi-alkali version with spectral response between 300 nm and 650 nm. An important feature of this tube for our application is its excellent immunity to modest magnetic field environments. A field of 100 Gauss would reduce signal by about 10%.

The Hamamatsu R5900U offers 3 times the gain of the R7400U and has a window that is not as delicate as that for the R7400U. We had expressed a concern about how much pressure the face of the R7400U could take from fibers at the cathode, and it happens that this is an issue for that tube. Figures 11.7 and 11.8 show how we are planning to use the quad configuration R5900U-00-M4 which maps out four separate anodes in a single miniature-metal dynode structure. Each sector will conveniently handle the 18 Kuraray Y11 WLS fibers (1.5 mm dia.) plus a single optical fiber that can direct calibration light from a LED calibrator. This is an economical solution as well since high voltage is common to the four sectors.

This bi-alkali tube has a spectral response of 300-650 nm which is a reasonable match to the emission band of the Kuraray Y11 WLS fiber at 450-550 nm. As shown in Figure 11.7 the minimum effective photocathode area of 18×18 mm is well matched to the task of accepting the 18 fibers (1.5 mm diam.) of a three-strip module at the face of a single quadrant. In all a single tube can handle the output of four modules. The common HV is 900 V, while the typical gain is 2×10^6 . Like the R7400U this tube has a very fast response. Anode pulse rise time is 1.2 ns and Transit Time Spread(FWHM) is 0.32 ns. Crosstalk between adjacent 9×9 mm² aperture is 2% which should not affect our application given the distance of the fibers from the interfaces between sectors. Uniformity variations between each anode of 1-to-1.5 can be dealt with in the fast linear amplifier following each anode. Magnetic fields of




Figure 11.7: Isometric view of front of the H8711 PMT assembly. The Mu shield and 4×4 alignment grid are visible.




Figure 11.8: Isometric view of back of H8711 PMT assembly (without Mu shield). The pin connections for the 16 anodes and metal clip are shown.

10 Gauss are well tolerated. We intend to use the D Type Socket Assembly E7083 which provides separate 45 cm long signal cables for the four anodes.

11.8 Assembly and Installation of a Three-Layer Module

This arrangement is diagrammed in Figure 11.9 for a 55 cm assembly module. This module holds both layers of scintillators which are sandwiched between sheets of thin aluminum corrugated to capture the staggered array of strips. Besides confining the module the aluminum skin will help to provide a light-tight enclosure for the scintillators, optical fibers, and photomultipliers. To accommodate the latter two the aluminum sheet will advance 20 cm beyond the readout end of the strips, but only on one side. The other side will have a removable aluminum cove plate that serves as a light seal when installation is completed. The final size of this three-layer module will span $2.2 \text{ m} \times 4.9 \text{ m}$. This size is convenient for mounting on the Unistrut frame shown in Figure 11.1. While in principle these large modules can interlock with their neighboring modules, it remains to be seen whether this is feasible with the aid of the hall crane. If not, the modules themselves can be mounted in a staggered pattern that overlaps them for full veto coverage.

Covering the bottom of the steel enclosure with the scintillators is a more challenging task. In the side view of Figure 11.2 is shown a floor of steel, pallet-like structures in two layers into which 55 cm modules can be introduced. The load of the steel enclosure and the detector solenoid bears upon the top of this. This design is an alternative to a more standard series of short vertical posts which prevent the placement of scintillator modules.

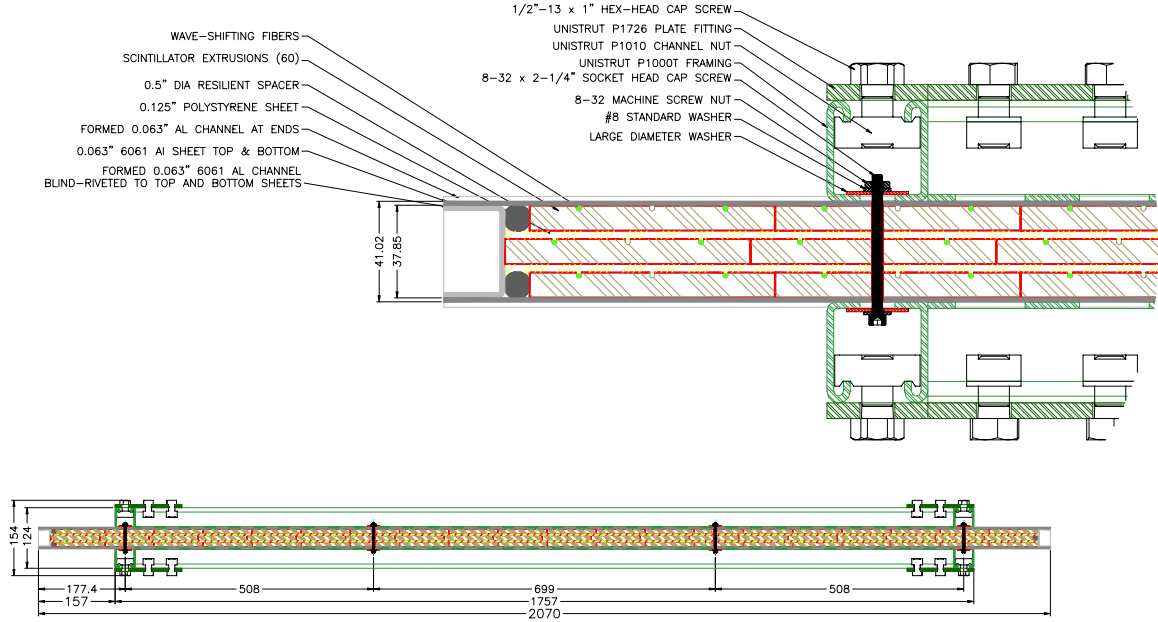


Figure 11.9: Path of optical fibers in transporting WLS fiber light to the pixels of multi-anode PMTs for a full module. Two H8711 assemblies (R7600-M16 PMT + Base) are required. Optical fibers only shown for 1 of 3 Layers, and then only one per scintillator bar. A Half Module is simply half of a Full Module.

11.9 Calibration of Active CR Shield

During the time periods between beam pulses cosmic ray muons offer a means of calibrating long term the readout amplitude and timing of the smallest unit of three adjacent strips in layer-1 of the CR Shield. These three strips are read out by a common PMT anode. In a typical CR event a coincident hit should occur in the three-strip unit which forms the neighboring unit of layer-2. This calibration is quite useful because the event used for calibration is precisely the signal seen by the detector during the beam pulse. The weak point in using this as a means of calibrating the Shield is its inability to quickly recognize a drop-off in response of individual strips. Cosmic ray fluctuations are large enough that other means of calibration are required for a fast indication of localized detector failure.

For short-term calibration and timing checks a pulsed N_2 laser will be used to access each strip scintillator by a fanout of its UV photon flux. An optical fiber connected to the far end of an individual strip will induce scintillator light that then simulates a regular charged particle event.

Consideration is also being given to a means of remotely drawing a ^{60}Co source across the sheet aluminum enclosing the strips. This will require a set of such source systems to reach all regions of the Shield. This has an advantage over a ^{137}Cs source in that the two coincident gammas from ^{60}Co , 1.17 MeV and 1.33 MeV, can produce

Mu2e Proposal

Compton signals in both scintillator layers beneath the source. The Compton Edge for the 1.17 MeV and 1.33 MeV gammas are 0.95 MeV and 1.1 MeV, respectively. The probability of a Compton signal in a 1 cm thick strip is 9%. Taking the fractional solid angles for both layers and the probability for Compton events a 20 μ Ci source can generate a coincidence rate of about 300 Hz. This offers the possibility of testing the coincidence between layers as well as the signal amplitude in individual layers.

12

Trigger & Data Acquisition System

12.1 Overview

THE DATA acquisition system (DAQ) must collect and organize data from various subsystems, make trigger decisions, build events, process potential physics events through a software filter and write a fraction of the triggered events to a storage medium. These broad tasks must be accomplished without significant dead time or other sources of inefficiency, while being robust against potential hardware failures. In addition, online monitoring and control of the experiment must be provided. A schematic diagram of the DAQ system is shown in Fig. 12.1. Much of the scheme is adapted from the MECO reference design [182].

The data from the tracker, calorimeter and cosmic ray shield will be processed, digitized and fed into pipelines. The tracker signals will be digitized and sparsified inside the vacuum chamber as is described in Sec. 12.2. Digital pipelines of the amplified and shaped analog signals from the calorimeter and cosmic ray shield will be made in the calorimeter digitizer modules (CDMs). The Level-1 trigger is generated when an energy cluster in the calorimeter is found to be above a preset energy threshold. This will cause the various subsystems to forward their data to event builders for further evaluation. Provision will be made for triggers to be issued occasionally to select data at lower energy thresholds for background studies and also for special calibration triggers.

Physics events will consist of the exceedingly rare DIO electron events with energies around 100 MeV, as well as copious background events which are expected to be associated mainly with lower energy clusters. The events will be built using a combination of data from the tracker, calorimeter and cosmic ray veto shield from the appropriate time slice (100-200 ns) of the pipelined data defined by the timing of the cluster that initiated the Level-1 trigger. The Event Builder will accept these data, build time-ordered events, and send them to a processor farm via a network switch. The online data analysis will have the capability for robust but simple track reconstruction with moderate resolution. This in turn will permit the DAQ to reject backgrounds by matching calorimeter clusters to tracks in space and time. The goal is to enable the DAQ to filter the data to reduce the rate at which data are written to tape to a manageable level.

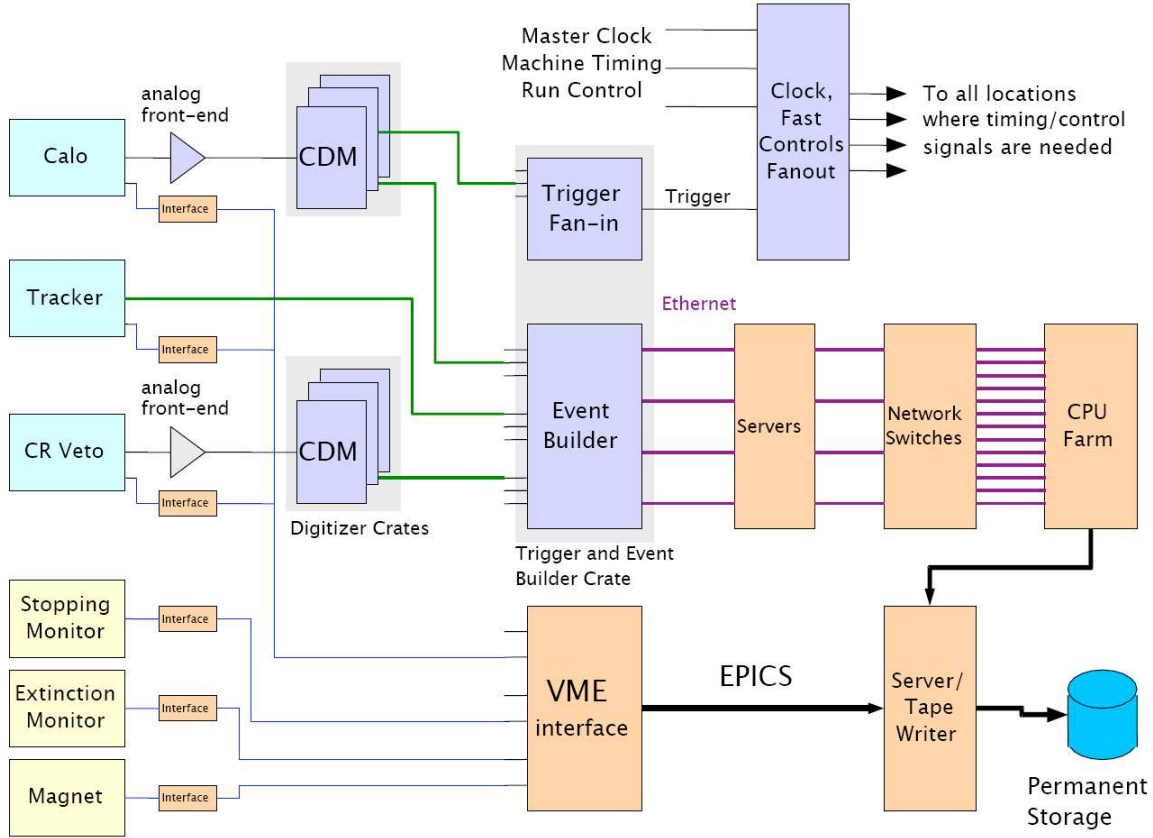


Figure 12.1: Schematic overview of the DAQ system.

A separate interface for experiment control as well as slower data flow to permanent storage will be based on the EPICS (Experimental Physics and Industrial Control System) framework. This will facilitate the recording of monitoring data, such as from the stopping and extinction monitors, as well as status data on all subsystems, such as the magnets, tracker, and calorimeter.

Data from the processor farm as well as from the EPICS slow control VME interface will be fed to a tape loader and drive with the capability of storing several TBytes of data. The DAQ system will also consist of a unified clock distribution system that will fan out the master clock, accelerator timing, trigger and other timing and control signals to all points in the online electronics where they are required. The DAQ system will also feature a clock and fast controls fanout subsystem that will fan out the trigger as well as a signal to indicate pulse injection, a system-wide clock (from which timing for various devices can be derived), an event counter and any other synchronization signals, system reset and initialization signals, etc.

12.2 Tracker Front End Electronics

The tracker readout electronics will be designed to work with either the anode wires or the cathode strips of the L-Tracker. The requirements and restrictions on the readout electronics include

- Amplitude and time information for the tracker hits, pipelined to allow trigger latency;
- Large dynamic range, to accommodate both electron hits and high ionization caused by low-energy protons;
- Low power consumption for the components operating in vacuum;
- Limited space available for the components and vacuum feed-throughs;
- High reliability;
- Low cost, large channel count.

The readout architecture will take advantage of modern integrated circuits technology, placing a multi-channel Application-Specific Integrate Circuit (ASIC) at the front end of the detector to read, digitize, and buffer the tracker hits. Such a fully digitized, pipelined, dead-timeless front end has been implemented in several modern DAQ systems. We plan to base the *Mu2e* tracker electronics on one of the successful ASIC designs.

Because of space limitations, the preamplifier will be placed on the detector plane near the anode wires and cathode strips. These preamplifiers will then feed an amplifier/discriminator placed on a digitizing board controlled through a local bus by a readout controller. A reasonable prototype is an amplifier-shaper-discriminator (ASD) chip which was designed for the ATLAS thin gap chamber [183]. Its threshold can be set remotely and it has both timing and analog outputs. Each ASD chip has 4 independent channels. The ASD chip will be redesigned for *Mu2e* because it is too old to be reproduced in modern foundries and does not have sufficient dynamic range to handle both signal electrons and highly ionizing, low energy proton backgrounds.

Most of the power in a readout system is expended on bus interfaces and impedance matching connections. These losses can be eliminated or reduced in a front-end application-specific ASIC. For example, an analysis of power consumption done for MECO [184] indicates that driving an analog signal out of the detector vacuum enclosure requires at least 150 mW/channel. In addition, locating digitizing electronics outside of the vacuum system has the disadvantage of requiring a very large number of vacuum feed-through connections.

An A/D ASIC, on the other hand, consumes about 30 mW/channel and will need a significantly smaller number of cables. A suitable ASIC, which provides digitized signals of both the analog waveform and threshold time signals, was developed for the

Table 12.1: Design parameters for the ELEFANT and an upgraded A/D chip

Parameter	ELEFANT	Upgraded chip
Channels/chip	8	8
Clock rate	60 MHz	15 – 60 MHz
ADC sampling rate	15 MHz	15 – 60 MHz
ADC resolution	6 bit flash	10 bit pipeline
TDC resolution	1.0 ns	1/64 of clock tick (0.2 – 1.0 ns)
Latency buffer	192×64	256×88
Latency time	12.6 μ s	3.4 – 17 μ s
Deadtime	32 clock ticks	1 clock tick
Output bus	8 bit	8/16 bit
Supply voltage	5.0 V	3.3 V
Total power	250 mW	<250 mW

BaBar experiment [185]. This ELEctronic Front-end Amplitude aNd Timing ASIC, (ELEFANT) has eight parallel flash ADC channels and 8 channels of delay lock loop (DLL) TDC, as well as a pipeline buffer to temporarily store the digitized data and send selected events to an event buffer if a trigger input line is activated. The chip can handle high single rates, and has low power consumption (250 mW for 8 channels). The ELEFANT chip is also able to output hit maps, which can be used to create seeds for the online (e.g. trigger) track reconstruction.

ELEFANT was designed and fabricated in a 0.8 μ m, 3-metal, 2-poly, CMOS process that is now obsolete. In addition its 5 V I/O is difficult to interface to modern FPGA designs, and its clock speed is insufficient for our purposes. With decommissioning of the BaBar detector, some number of chips may be available for early prototyping and tracker R&D, but for the actual *Mu2e* readout the chip needs to be redesigned. We can take advantage of new CMOS technology by rescaling the ELEFANT chip to match our requirements. Early prototyping was done for MECO [186]. Table 12.1 lists the design parameters for the ELEFANT chip and the upgraded version.

The block diagram of the digitization system is shown in Fig 12.2. All function blocks will be built around the latency buffer, which temporarily stores signal information before the trigger decision is made. The latency time, equal to the depth of the buffer times the clock period, must be longer than the time required to make the trigger decision. The sample rate of the ADC and the accuracy of the timing are also related to the clock rate.

The A to D block will process 8 data channels in parallel, with both the ADC and TDC encoding to 10 bit accuracy. Since the latency buffer alternately stores the ADC and TDC data, two characteristic bits will be used. In the present design, the

Trigger & Data Acquisition System




Figure 12.2: A block diagram of the front-end architecture for the readout electronics. The figure shows the preamplifier board, the digitizer board, the readout controller, and the sequencer.

chip has a 64 bit bus width, which will need to be increased to at least 88 bits. Two dual port SRAMs, (SRAM1 for data storage as an event buffer and SRAM2 for time stamp and hit map storage) are fed by the latency buffer. Upon arrival of the trigger, a pre-defined record length (8, 16, or 32 of data words each containing 64-88 bits) will be moved from the latency buffer into SRAM1, and a corresponding time stamp will be stored into SRAM2. When the data are moved, an 8 bit hit-map is formed, in which a bit is set when at least one TDC hit is recorded in the corresponding channel. The data in these SRAMs are then sequentially moved to the next buffer level under control of an FPGA, which uses the hit-map for zero suppression.

A unique component of the ELEFANT design is the hit occupancy map, which can be used in the generation of the event trigger orthogonal to the calorimeter trigger. An external trigger logic circuit can select appropriate data to be read from the latency buffer. In low occupancy running the system can be self-triggered with all hits read into the local bus. In normal (high occupancy) operation of *Mu2e*, tracker-generated trigger can be used in coincidence with the calorimeter, or a set of tracker-triggered events can be accepted with a prescale factor, to select events for calorimeter efficiency studies.

Readout for the Tracker will be organized as follows. Signals from the anodes and induction strips will be cabled to the preamps mounted along the outer radius of each vane which then feed frontend/digitizing boards. The digitizing boards will be connected via a backplane to the Readout Controllers (ROCs) with each ROC is connected to a readout sequencer. All modules will be placed in a shielded enclosure.

Eleven Front-End Boards (FEB) and two ROC modules will be required to read the anodes of a vane-plane (344 wires). Readout of the cathodes strips of a vane-plane (2×1040) will require 66 FEB, 11 ROC, and 3 sequencers. There will be eight vane-plane readouts per detector so 528 FEB, 88 ROC, and 24 sequencers are required.

According to simulations, the average signal hit involves 3 straws and 8 strips for a total of 11 channels. There will be an average of 10 crossings/hit resulting in 110 read channels per electron. The expected background level is 160 kHz/wire. We assume a 167 ns time gate for data acquisition. Each background hit would activate on average 4 cathode pads per resistive wire. Overall, we expect to read ~ 1000 channels per trigger, including both signal and background hits. Assuming the data size of 20 words/channel and 2 bytes/word, the total rate is ≈ 40 kb/trigger. A trigger rate of 1 kHz would require the transfer rate of 40 Mb/s. This can easily be transferred over 24 serial optical lines.

In addition to the data output, additional signals and voltage lines would need to be provided from outside:

1. High voltage (~ 1800 V). Each HV channel will feed one plane/vane. The current drawn during the active time gate will be approximately $1 \mu\text{A}/\text{wire}$. During the beam gate the current can increase by a factor of 10;
2. low voltage power of +3.3V (300A) and -3.3V (100A);
3. a slow control bus that sends control signals and environmental monitor information;
4. a system clock, to be regenerated by the local bus sequencer;
5. a trigger input which could be associated with a readout quadrant of wires;
6. a sync(trigger reset) counter that determines the time stamp placed on the data;
7. a system reset to return the readout to a standard operating condition.

12.3 Calorimeter Digitizer Modules

The calorimeter will measure the energy and position of the electron with modest resolution ($\sigma \sim 5\%$ and $\sigma \sim 1$ cm). As described in more detail in Chapter 10.4, the calorimeter will consist of four segmented vanes of lead-tungstate crystal detectors, placed just downstream of the tracker. The nominal plan is for each vane to consist of an 10×45 array of lead-tungstate crystals, or 1800 crystals total. Each column of 10 calorimeter crystals will be referred to as a tower. Each optically isolated crystal will be read out by two identical, side-by-side avalanche photodiodes (APDs) in

Trigger & Data Acquisition System

order to eliminate signals from charged particles traversing the APDs and to increase photostatistics, for a total of 3600 analog channels.

Pre-processed (amplified and shaped) analog signals from the crystal calorimeter will be fed into calorimeter digitizer modules (CDM) where the pulse heights will be digitally sampled with 8 bit resolution at a continuous 100-500 M samples/s. Each CDM will receive 40 APD analog signals from two towers of eight crystals. Thus, there will be a total of 90 CDMs for the calorimeter, mounted in 5-6 custom VME crates.

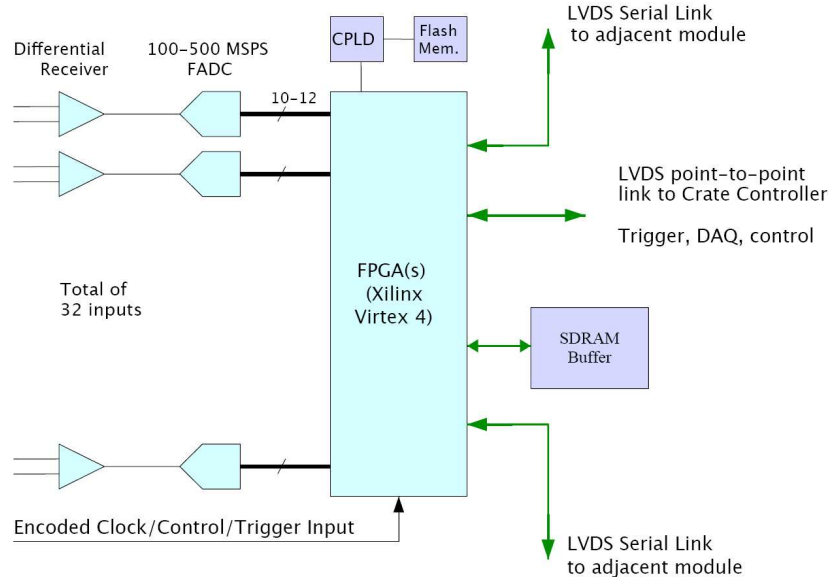


Figure 12.3: A block diagram of the Calorimeter Digitizer Module.

Identical digitizer modules (though with different firmware) will be used for the cosmic ray veto. The 40 differential analog signal inputs to each CDM module will have an absolute signal range of less than $\pm 5\text{VDC}$, and 100 ohms differential termination. The input signals must be shaped appropriately in order to derive energy and time information with good resolution. The digitizer input will be DC coupled. The gain and DC offset will be programmable over a reasonable range. The input signals will be accepted on a single multi-pin connector with individually-shielded pair cable.

A block diagram of the digitizer module is shown in Fig. 12.3. Each analog signal will be buffered by a high-speed buffer amplifier, with programmable DC offset and gain. The signals will be digitized continuously by a flash ADC. The ADC clocks will be synchronized and phase-locked to an experiment-wide master clock. All 32 channels of digitized data will feed directly into a large field-programmable gate array (FPGA). FPGAs will be chosen with sufficient logic and memory resources to accommodate the Level-1 trigger logic and DAQ buffering for the calorimeter readout, or alternatively local coincidence logic for the cosmic ray veto. Direct point-to-point links will be provided on the backplane between each module and adjacent modules

for triggering purposes. Sufficient bandwidth will be provided for event building and transmitting trigger information from each module to the crate controller

Pre-processed (amplified and shaped) analog signals from the Cosmic Ray photo-multipliers will be fed into the same CDM hardware as for the calorimeter, however with FPGA's reprogrammed for this specific task. The signals would be shaped to around 50 ns wide in order for the greater than 100 MSPS sampling to be effective in producing high resolution energy and time determinations for each pulse. Each CDM will receive 40 APD analog signals. There will be about 50-60 CDM modules for 1600–1900 channels in the cosmic ray veto mounted in 3-4 custom crates.

12.4 Level-1 Trigger

The initial processing in the Level-1 trigger logic will be contained in the CDMs. Each CDM sums energy in groups of 3x3 crystals using its own digitized energy values and energy values received from adjacent CDM modules. A trigger output will be generated when one of the sums, which is updated every 10 ns (assuming a 100 MHz clock in the CDM), exceeds a threshold.

A vane trigger (see Fig. 12.4) would be formed by the CDM crate controller as an 'or' of the CDM trigger outputs. The pattern of summed crystals will be programmable, and other options such as a 4x4 group could be implemented if required. The central trigger / DAQ crate contains a single trigger fan-in module which would receive trigger signals from the four CDM crate controllers (one for each calorimeter vane). It will produce as output an encoded trigger signal (trigger accept/reject plus particle location) at the CDM clock rate. Trigger time resolution better than 10ns, if required, could be extracted from the trigger sum by means of a digital filter.

No a priori limit is placed on the Level-1 trigger rate at the upstream end, except that at most one Level-1 trigger is permitted per pipeline clock period (nominally 10ns). Each Level-1 trigger causes a certain amount of data to be inserted into the DAQ pipeline, and since the bandwidth of this pipeline is finite, this does impose a practical limit on the trigger rate which can be accommodated. Simulation results suggest, assuming reasonable background rates, a Level-1 trigger rate (based on 3x3 crystal energy sum and an energy threshold in the range of 80 MeV) of nominally 1kHz. The electronics will be designed to handle trigger rates up to around 5 kHz.

The Level-1 trigger is the signal indicating the first recognition of a potential Physics Event. This signal tells all Physics Event-related systems (calorimeter, tracker, cosmic ray shield) to send their data to the DAQ for further evaluation. The Level-1 trigger is based solely on whether the energy in a cluster of crystals in the calorimeter exceeds a pre-assigned threshold. As discussed above, the first stage of the trigger is formed in the CDM crates: each CDM will have a fast link to the data in the CDMs of the two adjacent calorimeter towers and the digital sum of energy in any combination of pre-programmed clusters (e.g. a 3x3 array) of crystals is calculated.




Figure 12.4: Level-1 vane trigger

If the cluster sum exceeds a threshold, the CDM signals would then be delivered to the Crate Controller and then on to the Trigger Fan-in.

The Trigger Fan-in will form a logical OR of the CDM signals, and would issue a Level-1 trigger, which includes the trigger signal itself, plus the number of the calorimeter vane which has been hit (as required by the tracker electronics readout). The final Level-1 trigger, formed in the Trigger fan-in, is then fed to the Clock and Fast Controls Fanout, which in turn delivers the trigger to all systems.

12.5 Event Building and Processing

The Event Builder (EB), which collects data from the calorimeter, cosmic ray veto system and the tracker, will have 3-level hierarchy, as shown in Fig. 12.5. The level of event building (“Stage 1”) occurs in the CDM crates for the cosmic ray veto and calorimeter systems. Complete event fragments are output from each crate controller from crates containing CDMs, marked with an experiment-wide event number. A tracker interface module will be designed to accept digitized tracker data on fiber optic links and sort the data based on the time-stamp into event fragments.

The second level of event building (“Stage 2”) occurs in the central trigger/DAQ crate. Two stage 2 event builder modules and the tracker interface module will receive CDM data, check for errors and build sub-events. The “Stage 3” event builder module gets input from the three Stage 2 modules and combines all data into a single stream. The output of the Stage 3 module will be a single link to the memory buffer/Router

Mu2e Proposal

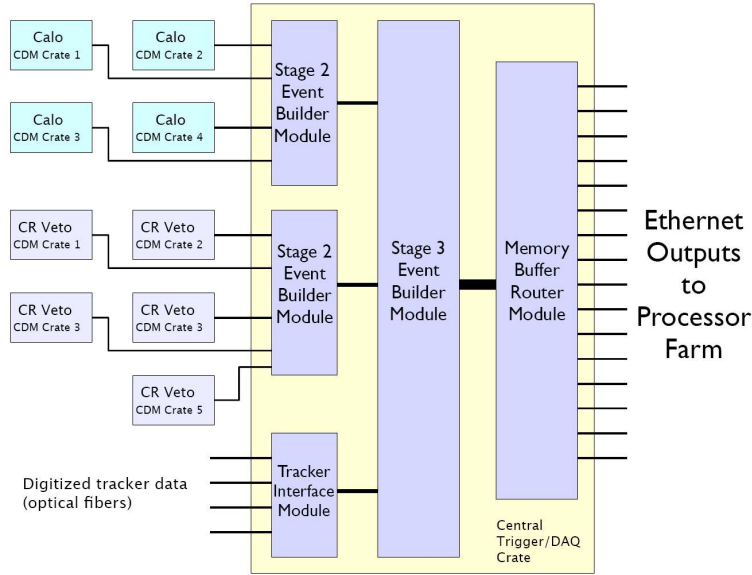


Figure 12.5: Schematic of the data flow in the Event Builder

module, where events are buffered into several output queues. Each queue will feed an Ethernet link, which feeds data to one of several computers in the processor farm.

The main task of the processor farm will be to process the Level-1 triggered events from the EB, reconstruct simple tracks and electromagnetic clusters and thus have the capability to select only those events that have a topology similar to the rare ~ 100 MeV electron tracks created at one of the target foils. These data, along with other data in response to auxiliary triggers to monitor the quality of the beam and detector subsystems, must be logged to permanent storage with minimal dead-time.

The data flow in and out of the processor farm is shown in Fig. 12.6. Data from the EB will be sent as 4 or 8 output streams using point-to-point Ethernet links. Diskless PCs will accept these data and farm them out via network switches. The processor farm will comprise about 50 PCs, each consisting of dual 3 to 4 GHz CPUs. Selected events will be sent to a central PC that will collect the data and write to permanent storage.

The online analysis must have the capability, in case of an unexpectedly large Level-1 trigger rate due to a noisy environment, to reconstruct tracks and reject background efficiently. A simplified algorithm that would provide 0.5–1.0 MeV momentum resolution would suffice; the target would be to process events at a rate greater than 50 Hz. This would easily allow the processor farm to handle a 2 kHz Level-1 trigger rate with about 100 CPUs. A factor of ten suppression in background would imply tape-writing at 200 Hz, corresponding to about 10 MByte/s, though the system can be designed to handle tape writing at a factor of 5 higher rate. A data storage facility of 5 to 10 TBytes should suffice for the online DAQ system.

Trigger & Data Acquisition System

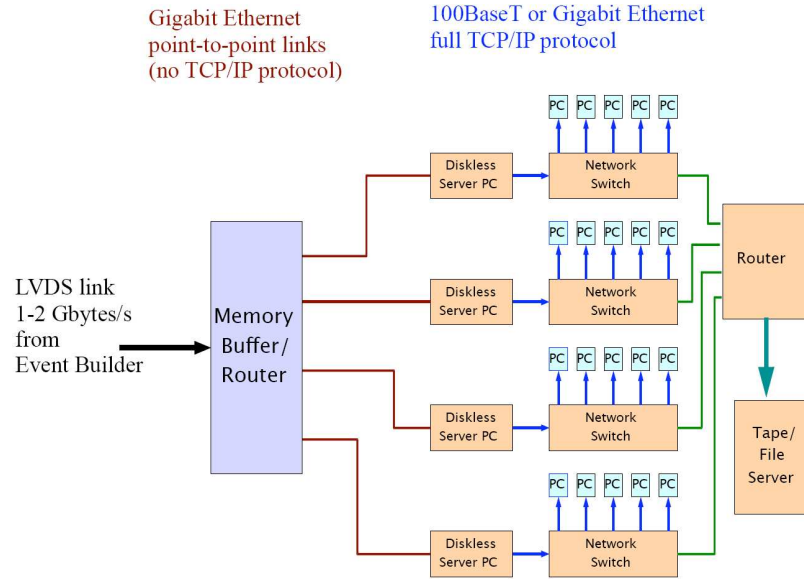


Figure 12.6: Schematic of the data flow in the Processor Farm

12.6 Slow Control and Monitoring

An important function of the DAQ system is monitoring all subsystems, setting up a framework for alarm handling, providing an interface for control and calibration, and recording monitor data asynchronous with the main triggered data. In addition, the slow control and monitoring system will record data from the stopping target monitor and the extinction monitoring system.

A software framework based on EPICS, will allow communication between control screens, alarm panels and device drivers (which can perform hardware manipulation and record data via interface cards). The central PC that writes the triggered data to permanent storage will also generate data records periodically containing monitor and auxiliary data.

The front-end hardware such as passive sensors and input/output controllers can be housed in a VME crate, with a crate controller providing access to software control processes running on remote PCs. Other devices such as high-voltage controllers can be directly accessed via Ethernet.

13

Offline Computing and Data Analysis

THE role of the offline computing is to produce fully reconstructed, calibrated, and classified samples of events for the final publication-quality results. While offline computing is typically outside the scope of a construction project, the offline processing of the data needs to be thought through and coordinated with the online to ensure efficient and expedient operations.

We envision the following algorithms to run in the offline system:

- *Full event reconstruction*, including calorimeter reconstruction, tracking, muon veto event reconstruction, and extinction monitor, among others.
- *Calibration and alignment* algorithms on specific subsets of data
- *Final event selection* and production of condensed *Event Summary Data* (ESD, also known as data summary tapes).

Calorimeter cluster reconstruction will expand on the algorithms in the Level-3 trigger, providing measurements of cluster energies, positions, timing, as well as cluster lateral moments to aide in background rejection. The calorimeter clusters would serve as the first point in the tracking reconstruction. Tracking algorithms will involve pattern recognition and fitting, and will provide track momentum and specific ionization measurements. Alignment and calibration algorithms will be run periodically on the subsets of data, with the results stored in the calibration database for use in the event processing. Fully reconstructed information will be condensed into an ESD format appropriate for keeping on the analysis disks. We estimate that the ESD format will use $\mathcal{O}(100 \text{ bytes/trigger})$, such that a year worth of ESD data would occupy less than 0.5 TB of disk space.

We plan to adopt well-established programming standards and practices for software development. Several appropriate software management tools are already adopted and maintained by Fermilab experiments, such as **SoftRelTools** (SRT), developed by CDF and BaBar, and **SCRAM**, used by CMS. We intend to take advantage of the local expertise and adopt a build system that would be maintained at Fermilab. The software code will be maintained in a versioning system, such as **CVS** (used with SRT

by CDF and BaBar) or somewhat more modern **SVN** (already adapted to SRT and SCRAM).

MECO developed a very detailed Monte Carlo simulation package **GMC**, based on **GEANT3**. **GMC** provided detailed simulation of physics and background processes, but did not include detector digitization modeling. **GMC** is serving as a benchmark for early development of *Mu2e* simulations. However, since **GEANT3** is no longer being developed, we will need to migrate to a **GEANT4**-based simulation. Similarly, MECO reconstruction software existed as a mixture of Fortran and C++ codes, with loose interfaces to the simulation. The reconstruction code will be reimplemented for *Mu2e*, and integrated with the simulation in the same software framework. This capability will be needed before CD2, in order to be able to fully evaluate the detector design and possible options.

Detector Enclosure, Civil Construction, and Infrastructure

THERE is no available space in the vicinity of the Antiproton Source to house an experiment the size of *Mu2e* so a new facility must be developed. The Fermilab Facilities Engineering Services Section (FESS) has done a preliminary design and cost estimate for a detector enclosure and the civil construction required for a beamline from the Antiproton source to the detector enclosure [74]. The new *Mu2e* detector hall will be located in an area west of Kautz Road and South of Giese Road near the SciBooNE detector Hall. The proposed location is shown in Fig 14.1. A more detailed layout is shown in Fig.14.2.

14.1 Beamline

The *Mu2e* Collaboration has worked with the Accelerator Division (AD) and FESS to develop a proposal for delivering beam to a new *Mu2e* detector enclosure. FESS has provided a preliminary design and cost estimate for the required civil construction [74]. The proton beam will be extracted downstream of the AP-10 region. The precise location and angle were chosen to minimize the amount of civil construction required in the Accumulator/Debuncher rings by bringing the beam out at the largest angle possible, thus minimizing the length of disturbed enclosure.

The beam then drops underneath a creek and wetland. A double-walled beam pipe will pass under the creek, requiring no disturbance of it and the surrounding wetland. At the border of the wetland we will place cast-in-place concrete below grade sloping down to the experimental hall. The cast-in-place sections will be roughly the same size as the Main Injector tunnels. A 150 foot long section will enclose the extinction channel.

Following the extinction channel the proton beam will enter the detector hall. An elevation view of the beamline is shown in Fig. 14.3.



Figure 14.1: A conceptual view of the site location.

14.2 Detector Enclosure

The detector enclosure has been designed to contain the solenoid system and the surrounding veto shield while providing sufficient space for installation and service of equipment. Space for staging the detector prior to its insertion into the detector solenoid has also been provided. The enclosure also satisfies our requirements for passive shielding from cosmic rays.

The idea is to have an S-shaped trench following the shape of the solenoid with sufficient room for installation and maintenance. Using this trench is much cheaper than excavating a large area and then stacking concrete shielding. While this certainly limits flexibility it is consistent with the current minimal design philosophy. There is enough room underground for cryogenic infrastructure. We are studying the details of accessing the veto counters and it may be necessary to widen the trench a few feet.

The trench will be covered with barite and steel to lower cosmic ray backgrounds from sky-shine to an acceptable level. We note MECO assumed 2.0 m of heavy concrete and 0.5 meters of steel (the latter also served as a flux return and cannot

Detector Enclosure, Civil Construction, and Infrastructure




Figure 14.2: A preliminary drawing of the site relative to the Accelerator Complex. Note the two already-existing pipelines for Helium crossing the Accumulator/Debuncher. This preliminary drawing does not include the access road from Giese Road or the hardstand for the liquid N_2 dewars, although these are included in the costs.

Mu2e Proposal




Figure 14.3: A preliminary drawing of the beam from the Accumulator/Debuncher through to the Detector Hall.

change.) The mean density of heavy concrete is $\approx 5.8 \text{ gm/cm}^3$ and steel is 7.87 gm/cm^3 . Earth is 2.2 gm/cm^3 . The sum of this shielding is 21 ft earth-equivalent which this design provides (for all but cosmic rays from near-overhead) without any additional shielding in the trench. We can certainly stack this much material above the solenoid in the trench and then slightly above floor level using the cheapest combination of materials.

The above-ground area of the hall is large enough to assemble the pieces of the solenoid and then lower them into the trench. We believe a 40 ton crane will suffice.

A beam absorber in a separate enclosure will absorb the remnant proton beam downstream of the production solenoid. It is possible, and we are considering, forming a muon test beam in this area but it is beyond the scope of this project and experiment.

14.3 Electrical Power

The electrical power for the detector enclosure will be provided by a new 1500 kVA transformer. The transformer has been sized to accommodate the anticipated needs to maintain the facility and operate the experiment. A new concrete encased power duct bank will be installed to connect the detector enclosure to the existing Fermilab electrical grid. These costs are included in the estimate from FESS.

14.4 Cryogenics

The superconducting solenoids require both liquid nitrogen and helium. The design has a road leading from Giese Road to a hardstand where we will place nitrogen dewars. Helium must be brought from the Accumulator/Debuncher region to the experiment and then return. Convenient supply lines exist to join to local supply and return lines. These lines will follow the beamline from the Accumulator/Debuncher tunnel through the below-grade pipe and into the cast-in-place enclosure.

Our current thought is that after the Tevatron ceases operation, and is no longer kept cold, the most efficient way to liquefy the Helium is to use Tevatron satellite refrigerators. We estimate that three satellite refrigerators will satisfy our needs. It is an interesting question to determine whether the minimum life-cycle cost will be associated with using these refrigerators as relatively inefficient liquefiers or to purchase new ones but for now we plan on recycling existing equipment.

Two modified satellite refrigerators would have to be installed to support operations of the *Mu2e* cryogenic components. To sustain reliable operation, these refrigerators will be modified to have redundant sets of dry and wet expanders. New valve boxes to interface refrigerators, expanders and the distribution system must be built.

The cryogenic distribution system includes a bayonet can to interface refrigerators, a horizontal and vertical transfer line, a valve box-phase separator located in the




Figure 14.4: A plan view of the trench for the solenoid.

Detector Enclosure, Civil Construction, and Infrastructure

tunnel, a tunnel transfer line, and a tunnel bayonet can. In this estimate, it is assumed that the P-bar Debuncher valve box with a phase-separator will be available for use at *Mu2e*. The distribution and control systems for superconducting magnets or for any other tunnel cryogenic components are not included in this estimate.

Helium and nitrogen inventory systems would need to be added. The current system gets its inventory from the Tevatron, which is assumed, will not be operating. Three 30,000 gallon Tevatron gaseous helium storage tanks will be relocated from their current location between A3 and A4 to the new *Mu2e* refrigerator building. Furthermore, one of the CHL nitrogen dewars assumed to be available will need to be moved to the new refrigerator area

It is assumed that at least three Tevatron helium screw compressors located at F0 will be available for the *Mu2e* cryogenic system. An estimated 900 ft. of suction and discharge headers will have to be installed to connect the F0 compressors and new refrigerators.

It is estimated that the new *Mu2e* cryogenic system will cost approximately \$1.6M and require approximately 10 FTE to conduct engineering, design and installation of the *Mu2e* cryogenic system. The cost breakdown is given in Chap.4.

15

Environment, Safety and Health

g 10000

MANY of the same ES&H issues as other Fermilab experiments are shared by *Mu2e*, including fire, ionizing radiation, RF radiation, oxygen deficiency and electrical hazards are all relevant safety concerns for *Mu2e*. Environmental issues include potential disturbance of wetlands and groundwater activation. We discuss specific issues in the Sections that follow.

15.1 Antiproton Accumulator and Debuncher Rings

Radiation levels are the most significant issue in the Accumulator/Debuncher tunnel. There are two parts to the issue:

1. Losses within the tunnel causing damage to components and raising the radiation levels to a point where they are too high to allow for work to be done in the enclosure.
2. Above-ground radiation. It is worth noting that the area under the tunnel is relatively impermeable clay and we expect no significant problem from groundwater contamination although a formal study must be done.

The passive shielding requirements are not met: there are only 10–13 ft of earth shielding. As noted in Chap. 6.7 the tunnel roof loading limits preclude adding the 10 ft of additional earth overburden required for open-occupancy. The area around the Accumulator/Debuncher will be fenced, making it a radiation area with a 5 mr/hr limit at the boundary. The enclosures themselves will be interlocked.

An important step is to understand losses within the Accumulator/Debuncher tunnels. We note they are fixed-momentum rings—hence configuration control on the magnets can be used to quickly catch failures and prevent prolonged accidental losses. In addition, the magnets provide some level of self-shielding for the tunnel and above-ground regions. Finally we expect to deliberately mis-steer the beam and measure losses at selected points. After these measurements, like those done for the Booster, have been performed, we will install “chipmunks” to monitor radiation levels inside the tunnels.

These are far from idle concerns. The expected particle flux is roughly 300K times the usual operating levels in the tunnel. Nonetheless we are optimistic this will not be a significant problem. The Accumulator/Debuncher has been used to store protons at 1/20 of the experiment intensity, and held them for prolonged periods. The limiting apertures are likely the portions of the RF that will no longer be needed. Only the region around the slow-extraction section is likely to require significant new shielding.

While none of this substitutes for a detailed ES&H study of the radiation levels we believe our preliminary assessment shows the scheme has no significant issues that make the use of the Accumulator and Debuncher tunnels unworkable.

15.2 Proton Beamline

The proton beamline exiting the Accumulator/Debuncher is below grade, involves no wetland mitigation by design, and will be sufficiently deep such that the natural passive shielding will exceed 23 ft and allow open access above.

15.3 Detector Enclosure

The shielding of the production solenoid is sufficient to mitigate groundwater concerns. The enclosure itself will be interlocked. The beam absorber will be standard; 3.6×10^{20} protons at 8 GeV involves no new technology or particular challenges from either a groundwater or residual activation standpoint.

16

Mu2e in the Project X Era

16.1 *Mu2e* and Future Accelerator Upgrades

ALL OF our discussions so far assume that the accelerator complex implements the baseline modifications proposed to support the NO ν A experiment [85] following the termination of the collider program. In this scenario, *Mu2e* uses the excess capacity of the Proton Source without impacting the NO ν A program.

It is reasonable to assume that this will remain the running configuration at the lab until at least 2014 or so. In the unlikely event the ILC is built on an aggressive timescale, and sited at Fermilab, the lab will likely remain in the basic NO ν A configuration at least through the 2010s.

Recently, a steering group was organized at Fermilab to discuss a range of options for cases in which the ILC is delayed and/or not built in the US. Two general scenarios emerged from these discussions [187]:

- In the event that the ILC is moderately delayed (2 years or so) or built on a short timescale outside the US, it becomes attractive to pursue the “SuperNuMI” (SNUMI) program: a series of rather aggressive upgrades to the existing proton complex to increase the beam power to the NuMI line to roughly 1.2 MW.
- If the ILC is significantly delayed or not built at all, then it is felt that the most promising project for the lab is so-called “Project-X”, an 8 GeV linac based on ILC technology which would inject into the Recycler for loading into the Main Injector. This would provide up to 2.3 MW to the NuMI line, as well as up to 200 kW of 8 GeV protons.

Below we consider the implications of each of these scenarios for the *Mu2e* experiment.

16.1.1 SNUMI

As originally conceived, the SuperNuMI project involved building a new 8 GeV proton line from the Booster to the Accumulator, where protons are momentum stacked prior to being loaded into the Recycler, from which they are loaded into the Main Injector

in a manner similar to NO ν A [188]. Under the SNuMI plan, 18 Booster batches are sent to the Accumulator every 1.33 second Main Injector cycle, leaving at most 2 additional batches for use by *Mu2e* or other experiments. This corresponds to reduction in available flux to 1.2×10^{20} protons per year, if protons are not diverted from the NO ν A program.

Recall that beam loss issues are one of the problems which have to be solved for this experiment. Under our baseline proposal, all of the protons in the Accumulator/Debuncher enclosure are used by the *Mu2e* experiment. A radiation limitation, particularly early in running, might result in a reduced, but still useful, flux. On the other hand, in the SNuMI scenario, three times as many protons would be loaded into the Accumulator/Debuncher enclosure just for the needs of NO ν A (18 Booster batches as opposed to 6). The radiation problems associated with these would have to be solved before any protons were available to *Mu2e*, which represent a small perturbation to the NO ν A needs.

It is important to note that the *Mu2e* experiment, if implemented first, could be a valuable step towards realizing the SNuMI goals. First of all, proton transport via the Recycler could be made to work even with proton stacking in the Recycler by upgrading the extraction magnet to a fast kicker that would allow Booster batches bound for the Accumulator to sneak in between stacked batches. This would eliminate the need for the new beam line from the Booster to the Accumulator, significantly reducing the scope of the SNuMI project. Also, initial *Mu2e* commissioning would be an important proof of principle for the proton momentum stacking in the Accumulator.

16.1.2 Project X

The 8 GeV linac proposed for Project X would be based on ILC cryomodules and RF systems [189]. It would produce 1 ms long 9 mA pulses of H^- ions at a 5 Hz repetition rate. Three of these would be loaded into the Recycler for transfer to acceleration in the Main Injector, corresponding to 2.3 MW at 120 GeV to the NO ν A program. The Main Injector cycle time would be 1.4 seconds, so this would leave up to 4 Linac pulses available to an 8 GeV program. This represents 200 kW of 8 GeV beam power, or over 3×10^{21} protons per year. This is almost 10 times more proton flux than the *Mu2e* baseline scenario. Obviously, it is desirable to take advantage of as much of this increased flux as we can; however, doing so presents some challenges.

First and foremost, there is the problem of getting the protons from the Recycler to the experiment. Originally, it was thought that the most straightforward approach would be to rebunch the protons in the Recycler, slow extract them from there, and transport them directly to the experiment; however, slow extraction of 8 GeV protons from the Recycler may not be practical [189], so we will be forced to fast extract beam from the Recycler to the Accumulator and then handle it in the same way we handle the Booster beam. A number of suggested techniques to do this have been discussed [189, 190]. They all appear to be quite straightforward. Once beam is

in the Accumulator/Debuncher enclosure, we face the challenge of keeping radiation losses within acceptable limits with the dramatically increased proton flux.

16.2 Detector upgrades

The Phase One experiment described here reflects many years of planning and design work. In comparison, the preliminary explorations of the many possible upgrade paths to a Phase Two experiment are much less detailed. (It is envisioned that Phase One would use Booster beam whereas Phase Two would probably occur in the Project X era.) Flexible planning is required for Phase Two to facilitate responding to multiple possible future scenarios. That planning must anticipate the various factors that might limit the performance of the Phase One experiment as well as the experimental results that might emerge. For example, by the time that the experiment has reached a single-event sensitivity of order 10^{-16} , the situation might be rosy or grim. At one extreme, the rosy scenario, there would be a clear-cut signal with very little background; at the other extreme, the grim scenario, there would be considerable background in the signal region with no hint of a signal. Intermediate scenarios can obviously also be imagined.

The rosy scenario would engender a desire for more data on various stopping target materials. Then Phase Two might involve increasing the data collection rate by taking advantage of Project X to provide higher proton intensity on the target and/or improving the muon/proton ratio. Some upgrade of the targeting system and the shielding for the production solenoid would probably also be necessary. The use of higher-Z stopping targets would discriminate among the various new physics possibilities. That in turn would necessitate measures to reduce the 700 ns dead time between the arrival of the proton beam on target and the detection of candidate events. One possible way to achieve that is described below.

In the grim scenario, various measures to reduce backgrounds would be considered. For example, a bend could be provided between the stopping target and the detector as in the PRISM/PRIME proposal.

Fundamentally new concepts for pion production and for muon collection and transport have also been developed. If these preliminary concepts prove viable, they would be useful whether the Phase One results are rosy or grim. Some of these ideas are inspired by recent work on muon cooling for neutrino factories and muon colliders. See, for example, Ref. [191] and Ref. [192]. For pion production, a dipole plus a wedge would create a nearly monochromatic pion beam. The muons resulting from the decays of the monochromatic pion beam would be collected and cooled in a helical cooling channel, thereby enhancing the stopping flux and reducing the dispersion of muon stopping times. This scheme also promises to reduce backgrounds because the hadrons contaminating the muon beam would stop in the cooling absorber and not reach the stopping target. In addition, a brighter muon beam would allow use of thinner stopping discs, thereby reducing systematic errors in event reconstruction.

Mu2e Proposal

References

- [1] Y. Fukuda *et al.*, Phys. Rev. Lett. **77**, 1683 (1996).
- [2] K. Hirata *et al.*, Phys. Lett. **B280**, 146 (1992).
- [3] B. Cleveland *et al.*, Astrophys. J. **496**, 505 (1998).
- [4] W. Hampel *et al.*, Phys. Lett. **B447**, 127 (1999).
- [5] J. Abdurashitov *et al.*, Phys. Rev. C **60**, 055801 (1999).
- [6] Y. Fukuda *et al.*, Phys. Rev. Lett. **81**, 1562 .
- [7] S. Fukuda *et al.*, Phys. Rev. Lett. **86**, 5651 .
- [8] Q. Ahmad *et al.*, Phys. Rev. Lett. **87**, 071301 (2001).
- [9] Q. Ahmad *et al.*, Phys. Rev. Lett. **89**, 011301 .
- [10] E. Aliu *et al.*, Phys. Rev. Lett. **94**, 0811802 (2005).
- [11] K. Eguchi *et al.*, Phys. Rev. Lett. **90**, 021802 (2003).
- [12] T. Araki *et al.*, Phys. Rev. Lett. **94**, 081801 (2005).
- [13] D. Michael *et al.*, Phys. Rev. Lett. **97**, 191801 (2006).
- [14] P. Minkowski, Phys. Lett. **B67**, 421 (1977).
- [15] M. Gell-Mann, P. Ramond, and R. Slansky, *Supergravity* (North Holland, Amsterdam, 1979).
- [16] T. Yanagida, in *Proceedings of the Workshop on Unified Theory and Baryon Number in the Universe*, edited by O. Sawada and A. Sugamoto (KEK, Tsukuba, Japan, 1979).
- [17] S. Glashow, *1979 Cargèse Lectures in Physics — Quarks and Leptons* (Plenum, New York, 1980), No. 707.

- [18] W. Marciano and A. Sanda, Phys. Lett. **B67**, 303 (1977).
- [19] M. Brooks *et al.*, Phys. Rev. Lett. **83**, 1521 (1999).
- [20] U. Bellgardt *et al.*, Nucl. Phys. **B299**, 1 (1988).
- [21] C. Dohmen *et al.*, Phys. Lett. **B317**, 631 (1993).
- [22] W. Bertl *et al.*, Eur. Phys. J. C **47**, 337 (2006).
- [23] S. Ahmed *et al.*, Phys. Rev. D **38**, 2102 (1988).
- [24] MEG Collaboration, <http://meg.web.psi.ch/>.
- [25] C. Yaguna, Int. J. Mod. Phys. A **21**, 1283 (2006).
- [26] Y. Kuno and Y. Okada, Rev. Mod. Phys. **73**, 151 (2001).
- [27] R. Kitano, M. Koike, and Y. Okada, Phys. Rev. D **66**, 096002 (2002), erratum-
ibid. D76, 059902 (2007).
- [28] Y. Farzan, JHEP **0707**, 054 (2007).
- [29] S. Treiman, F. Wilczek, and A. Zee, Phys. Rev. D **16**, 152 (1977).
- [30] A. Zee, Phys. Rev. Lett. **55**, 2382 (1985).
- [31] Y. Okada, K. i. Okumura, and Y. Shimizu, Phys. Rev. D **58**, 051901 (1998).
- [32] G. Bennett *et al.*, Phys. Rev. D **73**, 072003 (2006).
- [33] J. Hisano and K. Tobe, Phys. Lett. **B510**, 197 (2001).
- [34] V. Cirigliano, B. Grinstein, G. Isidori, and M. Wise, Nucl. Phys. **B728**, 121 (2005).
- [35] L. Calibbi, A. Faccia, A. Masiero, and S. Vempati, Phys. Rev. D **74**, 116002 (2006).
- [36] M. Blanke *et al.*, JHEP **705**, 013 (2007).
- [37] A. de Gouvêa, S. Lola, and K. Tobe, Phys. Rev. D **63**, 035004 (2001).
- [38] M. Kakizaki, Y. Ogura, and F. Shima, Phys. Lett. **B566**, 210 (2003).
- [39] P. Langacker, The Physics of Heavy Z' Gauge Bosons, 2008.
- [40] G. Giudice and R. Rattazzi, Phys. Rept. **322**, 419 (1999).
- [41] M. Fukugita and T. Yanagida, Phys. Lett. **B174**, 45 (1986).

REFERENCES

- [42] S. Petcov, S. Profumo, Y. Takanishi, and C. Yaguna, Nucl. Phys. **B676**, 453 (2004).
- [43] G. Giudice *et al.*, Nucl. Phys. **B685**, 89 (2004).
- [44] W. Buchmuller, P. Di Bari, and M. Plumacher, Annals Phys. **315**, 305 (2005).
- [45] S. Davidson, E. Nardi, and Y. Nir, Leptogenesis.
- [46] S. Pascoli, S. Petcov, and W. Rodejohann, Phys. Rev. D **68**, 093007 (2003).
- [47] S. I. Collaboration, Eur. Phys. J. C **47**, 337 (2006).
- [48] MECO Collaboration, Technical report, Brookhaven National Labs, (unpublished), see MU2E-doc-358.
- [49] A. Badertscher *et al.*, Phys. Lett. **B79**, 371 (1978), erratum-ibib.B80:434,1979.
- [50] V. Abadjev *et al.*, Technical report (unpublished).
- [51] T. Suzuki, D. Measday, and J. Roalsvig, Phys. Rev. C (1987).
- [52] K. Yonehara and V. Balbekov, Technical report, Fermilab (unpublished).
- [53] COMET Collaboration, The COMET Proposal to JPARC, JPARC Proposal, 2007.
- [54] R. Kitano *et al.*, Phys. Rev. D **66**, 096002 (2002), erratum-ibid.D76:059902,2007.
- [55] W. Marciano, private communication, Marciano and his collaborators Czarnecki and Sirlin have already performed much of the required calculations for the MUCAP experiment. See A. Czarnecki, W. Marciano, and A. Sirlin, Phys. Rev. Lett. **99**, 032003 (2007).
- [56] O. Shanker, Phys. Rev. D **25**, 1847 (1982).
- [57] T. Liu, Pattern Recognition and Background Studies, MECO Internal Note 038, 1999, also MU2E-docdb-154.
- [58] A. Frischknecht *et al.*, Phys. Rev. C **32**, 1506 (1985).
- [59] T. Stoudl *et al.*, Nucl. Phys. **A91**, 520 (1967).
- [60] R. Eramzhyan *et al.*, Nucl. Phys. **A290**, 294 (1977).
- [61] T. Liu, MECO \bar{p} Background Studies, MECO Internal Note 026, 1998, also MU2E-docdb-97.

- [62] O. Benhar *et al.*, Nucl. Phys. **A579**, 493 (1994).
- [63] B. Cork *et al.*, Phys. Rev. **107**, 248 (1957).
- [64] A. Vaisenberg *et al.*, Sov. Phys. JETP Lett. **29**, 661 (1979).
- [65] B. Kopeliovich and F. Niedermayer, Phys. Lett. **B151**, 437 (1985).
- [66] R. Armenteros *et al.*, Phys. Rev. **119**, 2068 (1960).
- [67] G. J. Marmer *et al.*, Phys. Rev. **179**, 1294 (1969).
- [68] T. Liu, MECO Late Arriving Particles Backgrounds, MECO Internal Note 029, 1998, also MU2E-docdb-100.
- [69] T. Liu, Current Distribution in Transport Solenoid, MECO Internal Note 061, 2001, also MU2E-docdb-155.
- [70] D. Koltick, Cosmic Ray Veto System, MECO Internal Note 014.
- [71] Yoshida, http://www-ps.kek.jp/jhf-np/NP08/presentations/muon/pdf/NP08_Muon_Yoshi%da.pdf.
- [72] RSVP, RSVP Cost Summary, [http://rsvp.bnl.gov/Project_Office/Estimates/April6-82005/RSVP_Level_3_%Matl_Labor_Cost_Summary_AY%20\\$%Sheet.pdf](http://rsvp.bnl.gov/Project_Office/Estimates/April6-82005/RSVP_Level_3_%Matl_Labor_Cost_Summary_AY%20$%Sheet.pdf).
- [73] A. Klebaner and J. Theilacker, Technical report, FNAL (unpublished).
- [74] F. Engineering, Technical report, Fermi National Accelerator Lab (unpublished).
- [75] Wojcicki review, RSVP cost estimate from information presented at the Wojcicki review, http://rsvp.bnl.gov/Project_Office.
- [76] RSVP, RSVP Schedule and Milestones, http://rsvp.bnl.gov/Project_Office/Reviews/RSVP_Summary_Milestones_April_13_v3.pdf.
- [77] Technical report, MIT Plasma Science and Fusion Center (unpublished).
- [78] Y. Cui, K. Lan, and E. Hungerford, IEEE Trans. Nucl. Sci **53**, 2281 (2006).
- [79] Y. Cui, K. Lan, and E. Hungerford, IEEE Trans. Nucl. Sci. **53**, 2281 (2006).
- [80] J. Krider and H. Nguyen, Technical report, Fermilab, (unpublished), mu2e-docdb-360.
- [81] V. Kashikhin *et al.*, Technical report, Fermilab, (unpublished), MU2E-doc-263.

REFERENCES

- [82] Fermilab Proton Plan, http://www-accel-proj.fnal.gov/Proton_Plan/index.shtml.
- [83] NOvA Experiment, <http://www-nova.fnal.gov/>.
- [84] C. Ankenbrandt *et al.*, Technical report, Fermilab (unpublished).
- [85] Accelerator and NuMI Upgrades (ANU) for the NOvA Experiment, http://tdserver1.fnal.gov/AcceleratorSupport/NOvA_ANU/.
- [86] D. McGinnis, Technical report, Fermilab (unpublished).
- [87] D. Neuffer, Technical report (unpublished).
- [88] N. Alexeev, S. Bereznitsky, and A. Bolshakov, Proceedings of the 1997 PAC, 1997.
- [89] MECO Collaboration, DRAFT MECO Technical Proposal, http://www.bnl.gov/npp/mu-e_docs/Draft_MECO_Tech_Prop.pdf, 2001.
- [90] E. Prebys, Technical report, Fermilab (unpublished).
- [91] E. Prebys, Technical report, Fermilab (unpublished).
- [92] R. Ducar and J. Reid, Technical report, Fermilab (unpublished).
- [93] D. McGinnis, Technical report, Fermilab (unpublished).
- [94] *Accelerator Shielding and Radioactivation*, Fermilab Radilogical Control Manual, ch. 7.
- [95] R. Djilkibaev and V. Lobashev, Sov. J. Nucl. Phys. **49**, 384 (1989).
- [96] R. Palmer, Nucl. Phys. Proc. Suppl. **51A**, 61 (1996).
- [97] D. Armutliiski *et al.*, preprint submitted to Sov. J. Nucl. Phys. (unpublished).
- [98] R. Djilkibaev, Technical report, New York University (unpublished).
- [99] M. Bachman and R. Lim, Technical report, University of California, Irvine (unpublished).
- [100] V. Tumakov, Technical report, University of California, Irvine (unpublished).
- [101] A. Warren, in *Mathematics Applied to Electrical Engineering 2 ed. Monographs on Electrical Engineering* (Chapman & Hall LTD, London, 1958), Vol. 9, p. 324.

- [102] C. Tranter, in *Bessel Functions with Some Physical Applications* (The English Universities Press LTD, London, 1968), p. 148.
- [103] in *Induction Heating Handbook* (McGraw-Hill, London, 1979), p. 426.
- [104] V. Tumakov, Technical report, University of California, Irvine (unpublished).
- [105] M. Hebert, V. Tumakov, W. Molzon, and B. Christensen, Technical report, MECO (unpublished).
- [106] R. Djilkibaev and A. Mincer, Technical report, MECO (unpublished).
- [107] V. Tumakov, Technical report, MECO (unpublished).
- [108] P. Yamin, Technical report, BNL (unpublished).
- [109] V. Tumakov, Technical report, MECO (unpublished).
- [110] A. McUmbert and W. Morse, Technical report, MECO, (unpublished), MU2E-doc-229.
- [111] D. Weiss *et al.*, Technical report, MECO, (unpublished), MU2E-doc225.
- [112] J. Popp and M. McKeown, Technical report, MECO, (unpublished), MU2E-doc-214.
- [113] T. M. Collaboration, DRAFT MECO Technical Proposal, Available at <http://mu2e-docdb.fnal.gov/cgi-bin/ShowDocument?docid=5>, 2001.
- [114] A. Radovinsky, MECO Coil Sizing with DS Iron, Available at <http://mu2e-docdb.fnal.gov/cgi-bin/ShowDocument?docid=324>, 2001, MECO-Magnet-Document-MM040.
- [115] W. V. Hassenzah, MECO Magnet Field Specifications, Available at <http://mu2e-docdb.fnal.gov/cgi-bin/ShowDocument?docid=291>, 2001, MECO-Magnet-Document-MM005.
- [116] A. Radovinsky, Tolerances Sensitivity Study, Rev.1, Available at <http://mu2e-docdb.fnal.gov/cgi-bin/ShowDocument?docid=314>, 2001, MECO-Magnet-Document-MM030.
- [117] A. Radovinsky, MECO Energy, and Inductances Update, Available at <http://mu2e-docdb.fnal.gov/cgi-bin/ShowDocument?docid=316>, 2001, MECO-Magnet-Document-MM032.
- [118] B. Smith, PS Coil Design with Increased Margins, Available at <http://mu2e-docdb.fnal.gov/cgi-bin/ShowDocument?docid=322>, 2001, MECO-Magnet-Document-MM038.

REFERENCES

- [119] B. Smith, MECO Conductor Joint Design and Analysis, Available at <http://mu2e-docdb.fnal.gov/cgi-bin/ShowDocument?docid=332>, 2001, MECO-Magnet-Document-MM048.
- [120] B. Smith and P. Titus, PS Coil Quench Analysis, Available at <http://mu2e-docdb.fnal.gov/cgi-bin/ShowDocument?docid=327>, 2001, MECO-Magnet-Document-MM043.
- [121] J. Brisson and J. Smith, Thermal studies of MECO PS magnets, Available at <http://mu2e-docdb.fnal.gov/cgi-bin/ShowDocument?docid=323>, 2001, MECO-Magnet-Document-MM043.
- [122] P. by the Massachusetts Institute of Technology Plasma Science and F. Center, MECO Superconducting Solenoid System Conceptual Design Report, Available at <http://mu2e-docdb.fnal.gov/cgi-bin/ShowDocument?docid=6>, 2002.
- [123] B. Smith, Mu2e Magnet Status, Changes to the magnet design include: Updating the field specifications, changing the number of coils in the production solenoid from 9 to 11, adding shielding around the production solenoid, adding endcaps on the production solenoid. Available at <http://mu2e-docdb.fnal.gov/cgi-bin/ShowDocument?docid=351>, 2008.
- [124] C. Chen, T. Liu, B. Christensen, and M. Hebert, Revised Return Yoke and Coil Build Design Information, Available at <http://mu2e-docdb.fnal.gov/cgi-bin/ShowDocument?docid=367>, 2004, MECO-SOL-03-001.
- [125] P. Titus, Displacement Input for the Field Quality Study, MECO Transport Solenoid, Available at <http://mu2e-docdb.fnal.gov/cgi-bin/ShowDocument?docid=348>, 2004, MECO-Magnet-Document-MM065.
- [126] D. Weiss *et al.*, Anti-Proton Stopping Window Reference Design, MECO-MUB-03-001, 2003.
- [127] H. Herzog and K. Alder, Helvetica Physica Acta (1980).
- [128] O. Shanker, Phys. Rev. D **25**, 1847 (1982).
- [129] P. Wintz, AIP Conf. Proc. **698**, 789 (2004).
- [130] V. Balashov and R. Eramzhyan, Atomic Energy Reviews **5**, 1 (1967).
- [131] Y. Budyashov, Soviet Physics JETP **33**, 11 (1971).
- [132] P. Singer, Springer Trackst in Modern Physics **71**, 39 (1974).
- [133] S. Sobottka and E. Wolls, Phys. Rev. Lett. **20**, 596 (1967).

- [134] L. Vilgel'mova *et al.*, Sov. J. Nucl. Phys. **13**, 310 (1971).
- [135] A. Wyttenbach *et al.*, Helv. Phys. Acta. (1976).
- [136] N. Mukhopadhyay, Rep. Prog. Phys. **30**, 1 (1977).
- [137] T. Kozlowski, Nucl. Phys. **A436**, 717 (1985).
- [138] T. Liu, Technical report, University of California, Irvine (unpublished).
- [139] V. Bondarenko *et al.*, Technical report, CERN (unpublished).
- [140] S. Majewski *et al.*, Nucl. Instrum. Methods (1994).
- [141] D. Ambrose *et al.*, Phys. Rev. Lett. **81**, 4309 (1998).
- [142] D. Ambrose *et al.*, Phys. Rev. Lett. **81**, 5734 (1998).
- [143] G. DeCataldo, Nucl. Instrum. Methods **A409**, 73 (1998).
- [144] E. Barbarito, Nucl. Instrum. Methods **A381**, 39 (1996).
- [145] J. Gordon and E. Mathieson, Nucl. Instrum. Methods **A227**, 267 (1984).
- [146] J. Gordon and E. Mathieson, Nucl. Instrum. Methods **A235**, 505 (1985).
- [147] J. Gordon and E. Mathieson, Nucl. Instrum. Methods **A270**, 602 (1988).
- [148] K. Lan, Y. Cui, and E. Hungerford, IEEE Trans. Nucl. Sci. **53**, 1390 (2006).
- [149] J. Butler *et al.*, Nucl. Instrum. Methods **A290**, 122 (1990).
- [150] Arai *et al.*, Nucl. Instrum. Methods **A381**, 355 (1996).
- [151] R. Djilkibaev and R. Konoplich, arXiv:hep-ex/0312022 (unpublished).
- [152] M. Kubantsev, I. Larin, and A. A. Gasparian, JLAB-PHY-06-528, arXiv:physics/0609201 (unpublished).
- [153] CMS Collaboration, Technical report, CMS (unpublished).
- [154] ALICE Collaboration, Technical report, CERN/LHCC 99-4, ALICE TDR 2 (unpublished).
- [155] D. Aleksandrov *et al.*, Nucl. Instrum. Methods **A550**, 169 (2005).
- [156] J. Chen, L. Zhang, and Z. R.Y., Nucl. Instrum. Methods **A572**, 218 (2007).
- [157] M. Kotulla *et al.*, Technical report (unpublished).

REFERENCES

- [158] P. Semenov *et al.*, Nucl. Instrum. Methods **A582**, 575 (2007).
- [159] R. Novotny, talk given at CALOR 2006 conference, 2006.
- [160] A. Borisevich, Nucl. Instrum. Methods **A537**, 101 (2005).
- [161] SICCAS, Technical report, Shanghai SICCAS High Technology Corp., 1295 DingXi Rd., Shanghai, 200050 China (unpublished).
- [162] JSC, Technical report, JSC Bogoroditsk Technical Chemical Plant (BTCP), Bogoroditsk, Tula oblast, Russia (unpublished).
- [163] Technical report, 21 Promyshlennaya Street, Apatitys, Murmansk region, 184209, Russia (unpublished).
- [164] Technical report, Materials Research Laboratory, Tsukuba 305-0856, Japan (unpublished).
- [165] Technical report, Kharkov 61001, Ukraine (unpublished).
- [166] Radiation Monitoring Devices, Inc., Technical report, 44 Hunt St., Watertown, MA 02472 (unpublished).
- [167] Advanced Photonics, Inc., Technical report, 151 Trade Zone Drive, Ronkonkoma, NY 11779 (unpublished).
- [168] Hamamatsu, Inc., Technical report, Hamamatsu, Inc., 360 Foothill Rd, Bridgewater, NJ 08807 (unpublished).
- [169] D. Renker, in *9th International Conference on Calorimetry in Particle Physics* (Calor 2000, Annecy, 2000).
- [170] R. Djilkibaev, J. Sculli, and A. Toropin, Technical report, New York University (unpublished).
- [171] A. Bartoloni *et al.*, Nucl. Instrum. Methods **A582**, 462 (2007).
- [172] P. Lecoq *et al.*, Nucl. Instrum. Methods **A365**, 291 (1995).
- [173] R. Zhu, Nucl. Instrum. Methods **A413**, 297 (1998).
- [174] R. Djilkibaev and J. Sculli, Technical report, New York University (unpublished).
- [175] V. Batarin *et al.*, Nucl. Instrum. Methods **A556**, 94 (2006).
- [176] M. Anfreville, Nucl. Instrum. Methods **A594**, 292 (2008).
- [177] M. Castoldi *et al.*, Nucl. Instrum. Methods **A403**, 22 (1998).

- [178] J. Kremer *et al.*, Phys. Rev. Lett. **83**, 4241 (1999).
- [179] M. Collaboration, Technical report, FNAL (unpublished).
- [180] B. Choudhary *et al.*, Technical report, Caltech (unpublished).
- [181] R. Wojcik *et al.*, Nucl. Instrum. Methods **A342**, 416 (1994).
- [182] E. Hazen *et al.*, Technical report, MECO (unpublished).
- [183] O. Sasaki and M. Yoshiro, IEEE Trans. Nucl. Sci. **46**, 1871 (1999).
- [184] K. Lan, Y. Cui, and E. Hungerford, IEEE Trans. Nucl. Sci. **51**, 2362 (2004).
- [185] S. Dow *et al.*, IEEE Trans. Nucl. Sci. **46**, 785 (1999).
- [186] E. Hungerford, Technical report (unpublished).
- [187] Report from the Project X steering group, http://www.fnal.gov/directorate/Longrange/Steering_Public/report/Steering-Group-Report-2007-09-18.pdf.
- [188] D. McGinnis, Technical report, Fermilab (unpublished).
- [189] McGinnis *et al.*, Accelerator Issues of Project X, http://www.fnal.gov/directorate/Fermilab_AAC/AAC_July_07/ProjectX.pdf.
- [190] C. Ankenbrandt, Technical report, Fermilab (unpublished).
- [191] Yonehara *et al.*, in *Proceedings of EPAC 2006*, edited by K. Yonehara *et al.* (WEPLS016, Edinburgh, Scotland, 2006).
- [192] M. Cummings *et al.*, Intense Stopping Muon Beams, Proceedings of EPAC 2008, Genoa, Italy, MOPP071.