

Letter of Intent

A Muon to Electron Conversion Experiment at Fermilab

The Mu2e Collaboration

28 September 2007

R.M. Carey, K.R. Lynch, J.P. Miller^{*}, B.L. Roberts
Boston University

W.J. Marciano, Y. Semertzidis, P. Yamin
Brookhaven National Laboratory

Yu.G. Kolomensky
University of California, Berkeley

C.M. Ankenbrandt, R.H. Bernstein, D. Bogert, S.J. Brice, D.R. Broemmelsiek,
D.F. DeJongh, S. Geer, M.A. Martens, D.V. Neuffer, M. Popovic, E.J. Prebys^{*}, R.E. Ray,
H.B. White, K. Yonehara, C.Y. Yoshikawa
Fermi National Accelerator Laboratory

D. Dale, K.J. Keeter, J.L. Popp, E. Tatar
Idaho State University

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

V. Lobashev
Institute for Nuclear Research, Moscow, Russia

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

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

R.S. Holmes, P.A. Souder
Syracuse University

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

^{*} Co-Contact Person

Letter of Intent

A Muon to Electron Conversion Experiment at Fermilab

Abstract:

We are writing this letter to express our interest in pursuing an experiment at Fermilab to search for neutrinoless conversion of muons into electrons in the field of a nucleus, which is a lepton flavor-violating (LFV) reaction. The sensitivity goal of this experiment represents an improvement of more than a factor of 10,000 over existing limits. It would provide the most sensitive test of LFV, a unique and essential window on new physics unavailable at the high energy frontier. We present a conceptual scheme that would exploit the existing Fermilab Accumulator and Debuncher rings to generate the required characteristics of the primary proton beam. The proposal requires only modest modifications to the accelerator complex beyond those already planned for the NOvA experiment, with which this experiment would be fully compatible; however, it could also benefit significantly from possible upgrades such as the “Project X” linac. We include the conceptual design of the muon beam and the experimental apparatus, which use the previously proposed MECO experiment as a starting point.

TABLE OF CONTENTS

1. INTRODUCTION AND MOTIVATION	4
2. EXPERIMENTAL TECHNIQUE	8
2.1 Event Topology	9
2.2 Principal Backgrounds	10
3. THE PROTON BEAM	13
3.1 Overview	13
3.2 Delivering Protons to the Accumulator (“Boomerang Scheme”)	14
3.3 Momentum Stacking and Rebunching	15
3.4 Resonant Extraction	17
3.5 Proton Extinction	19
3.6 Project Siting	20
3.7 Total Proton Delivery	21
3.8 Radiation Safety	22
3.9 Mu2e and Future Accelerator Upgrades	22
4. EXPERIMENTAL APPARATUS	24
4.1 Design Overview	24
4.2 Principal Subsystems	26
5. OUTLOOK	38
5.1 Near Term R&D Plan and Request for Support	39
5.2 Rough Estimate of Construction Cost and Timeline	41
BIBLIOGRAPHY	42

1. Introduction and Motivation

The Standard Model (SM) provides a remarkably good description of the basic constituents of matter and the interactions among them, from low energies to the mass scale of the weak vector bosons. Nevertheless, the SM appears to be incomplete, and the goal of world-wide experimental efforts is to uncover physics beyond the SM (BSM). BSM physics is motivated by the Standard Model's inability to satisfactorily account for several important puzzles, such as the resolution of the hierarchy problem that renders the Higgs boson radiative corrections unstable, the origin of CP violation required for baryogenesis, and the incorporation of cosmological evidence for the existence of cold dark matter.

The majority of the theoretical approaches that address the aforementioned issues predict new dynamics at the TeV scale. Examples include supersymmetry, technicolor, and extra dimensions. While the first evidence for such new physics might be discovered at the Tevatron or at the LHC, a comprehensive understanding will likely require both high precision measurements at lower energy as well as collider experiments beyond the LHC.

Desirable low energy precision measurements are those that are, first and foremost, sensitive to new TeV-scale dynamics in BSM models where collider signals would be weak or nonexistent. Additionally, should BSM physics be first observed at the Tevatron or the LHC, specific precision low energy measurements would play an important role in deciphering the underlying BSM dynamics. The most important measurements are those where experimental techniques can achieve the high sensitivity required to indirectly access phenomena at the multi-TeV scale and where theoretical uncertainties are under control at the required level of interpretability.

One class of experiments that satisfy the above criteria is the search for charged lepton flavor violating interactions (LFV). Specifically, we propose to search for the coherent conversion of a muon to an electron in the field of a nucleus ($\mu+N \rightarrow e+N$) with unprecedented sensitivity. As we elaborate below, this process occurs in most BSM theories at levels that can be accessed by the experimental technique described in this LOI (the Mu2e experiment). Further, this is one of a handful of rare processes that has a clear experimental upgrade path, thus emerging as the most promising LFV process to be pursued, when compared to searches for forbidden decays such as various neutrinoless decays of tau leptons and muons and LFV leptonic and semi-leptonic decays of K and B mesons.

With the discovery of neutrino oscillations, we now know that $\mu+N \rightarrow e+N$ must occur. However, the SM prediction for $R_{\mu e}$, the ratio of the probability of $\mu+N \rightarrow e+N$ relative to ordinary muon capture is highly suppressed by the tiny neutrino mass (to the fourth power), and is calculated to be of the order of 10^{-52} , well beneath the reach of any conceivable experimental measurement. Thus, BSM physics would be unambiguously indicated by the observation of the phenomenon, which has the spectacular signature of a single high-energy conversion electron of well-defined energy emerging from a target where low energy negative muons are stopped.

This reaction is closely related to the LFV process $\mu \rightarrow e+\gamma$, which is being searched for by the MEG experiment at PSI [1], except that the photon must be real in that case. While a clear signal

in either experiment would be an unambiguous sign of BSM physics, the relative size of the two measurements would provide a vital clue as to the details of that physics. In some new physics scenarios, the Mu2e and MEG experiments, at their design sensitivities, have comparable sensitivities. In many other scenarios, Mu2e is the more sensitive reaction, sometimes by many orders of magnitude.

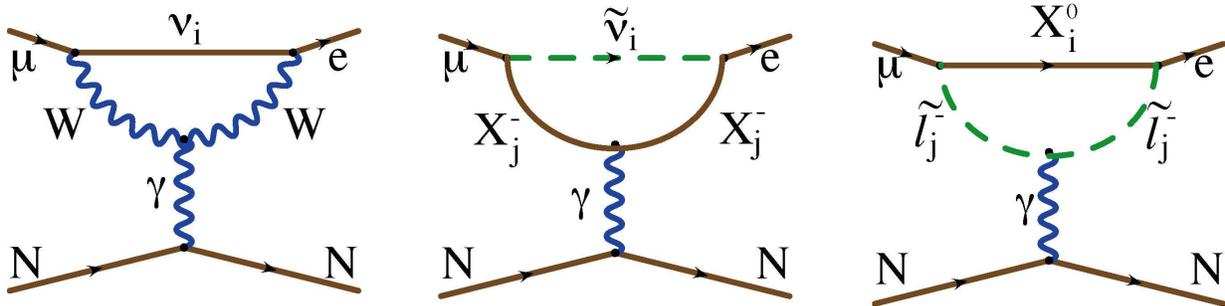


Figure 1.1: The leading Standard Model diagram for $\mu+N \rightarrow e+N$ is shown on the left. The center and right figures are the dominant SUSY diagrams.

It is instructive to see how the “transition dipole” SM diagram responsible for $\mu+N \rightarrow e+N$, as shown on the left side of Figure 1.1, receives additional contributions shown at the center and right of the same figure in a SUSY model [2]. This enhances $R_{\mu e}$ by up to 40 orders of magnitude to within the reach of experiments.

Under the assumption that the transition dipole SUSY diagrams, which can contribute to both reactions, dominate, the $\mu \rightarrow e+\gamma$ rate is 200 to 400 times larger than that for $R_{\mu e}$, depending on the choice of the Mu2e target nucleus [3, 4, 5]. The diagrams are similar to those that contribute to the muon g-2 anomaly, the latter being sensitive to the corresponding lepton flavor-conserving amplitude.

The additional mixing parameters introduced for the LFV diagrams are typically dependent on the fractional mass difference between first and second generation sleptons. The current best experimental limit on $R_{\mu e}$ is 4.3×10^{-12} (for a Ti target) [7] and 1.2×10^{-11} for $\mu \rightarrow e+\gamma$ [8]. Within the context of SUSY, these results can be combined with the latest muon g-2 result [9, 10] to constrain the slepton mass difference to be smaller than about 30 MeV for 500 GeV slepton masses. The goal for the first phase of MEG is a sensitivity of 10^{-13} for $\text{Br}(\mu \rightarrow e+\gamma)$, with a possible ultimate upgrade to 2×10^{-14} . Any further improvement beyond this will likely be very difficult to achieve. The initial goal for Mu2e is $R_{\mu e} \sim 10^{-16}$, with single event sensitivity of 2×10^{-17} . This would probe slepton mass differences down to about 1 MeV. Thus, the MEG and Mu2e goals both ensure sensitivity to a large portion of SUSY parameter space, as depicted in Figure 1.2. However, there are potential upgrade paths for the Mu2e experiment which could measure $R_{\mu e}$ to the level of 10^{-18} .

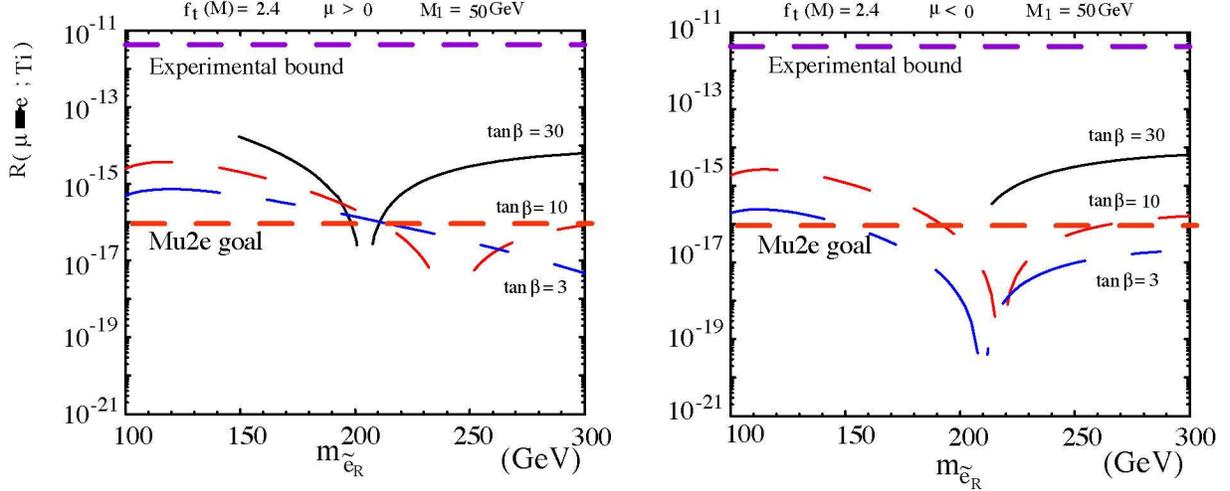


Figure 1.2: Predicted branching ratios for an SU(5) SUSY GUT Model are shown (from [6]), for positive and negative values of μ , the ratio of Higgs particle vacuum expectation values, as a function of the right-handed selectron mass. Also shown are the current experimental limit and the Mu2e goal.

Furthermore, since the mediator for coherent conversion couples to the nucleus and thus can be virtual, Mu2e possesses additional compelling discovery reach, with sensitivity to new classes of chiral changing amplitudes, like the SUSY case, as well as chiral conserving amplitudes [11]. Indeed, even within the SM there are chiral conserving Z exchange and W box diagrams, which are responsible for a factor of 100 enhancement of $R_{\mu e}$ over the $\mu \rightarrow e + \gamma$ rate, although the resulting rate is still too small to be observed. Such enhancement occurs in many BSM models for similar reasons. For example, while current limits on the mixing of a fourth generation heavy neutrino would make it unobservable in MEG, it is still possible to see 1000's of events in Mu2e.

In contrast to SUSY models, other models predict a weaker relationship for MEG and Mu2e. Figure 1.3 shows the accessible parameter space for MEG and Mu2e in the Littlest Higgs Model with T-Parity [12], where it can be seen that the relative enhancement between the two processes varies from 0.01 to 100. Note that the Mu2e reach spans virtually the entire parameter space. Another example with a significant relative enhancement is a Randall-Sundrum model with one compact extra dimension [13]. Even in the case where MEG reaches its ultimate goal after upgrades, both experiments would see signals of similar significance, as shown in Figure 1.4, and both have sensitivity well beyond the reach of direct observation at the LHC. Finally, the Mu2e measurement would access BSM amplitudes such as those due to leptoquarks [14], models with a Z' [15], or L-R symmetric models [16, 17], to which $\mu \rightarrow e + \gamma$ is insensitive, again with discovery reach well beyond what could be directly observed at the LHC.

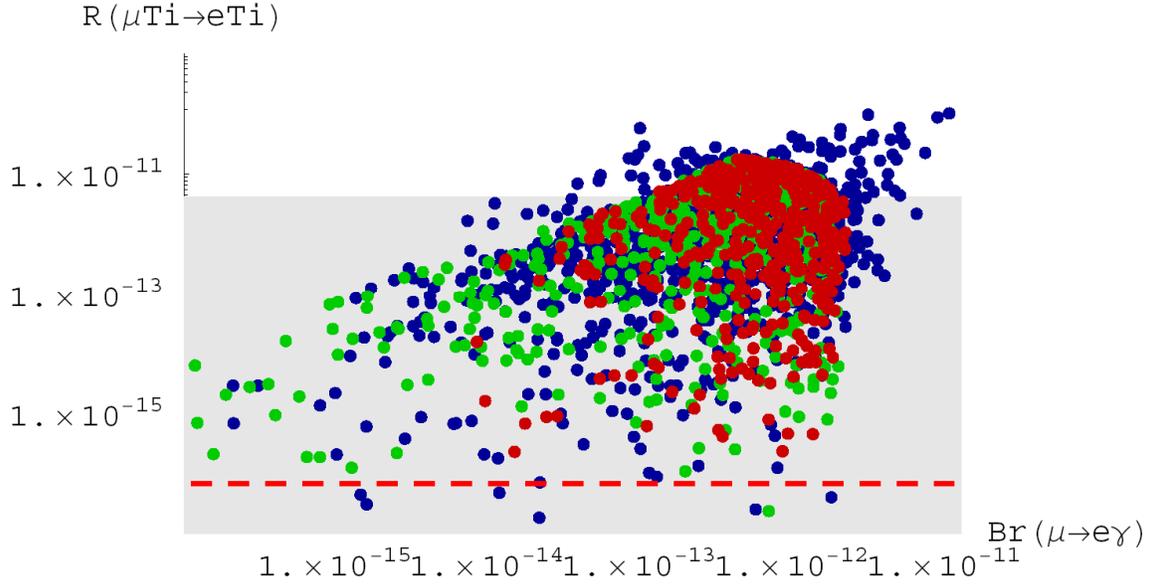


Figure 1.3: $\mu + {}^{48}\text{Ti} \rightarrow e + {}^{48}\text{Ti}$ rate as a function of $\text{Br}(\mu \rightarrow e + \gamma)$ for three (red, blue and green) scenarios of Little Higgs models with T-parity after imposing all existing LFV constraints (from [12]). The clear regions on the top and right are excluded. Also shown is the proposed goal for Mu2e.

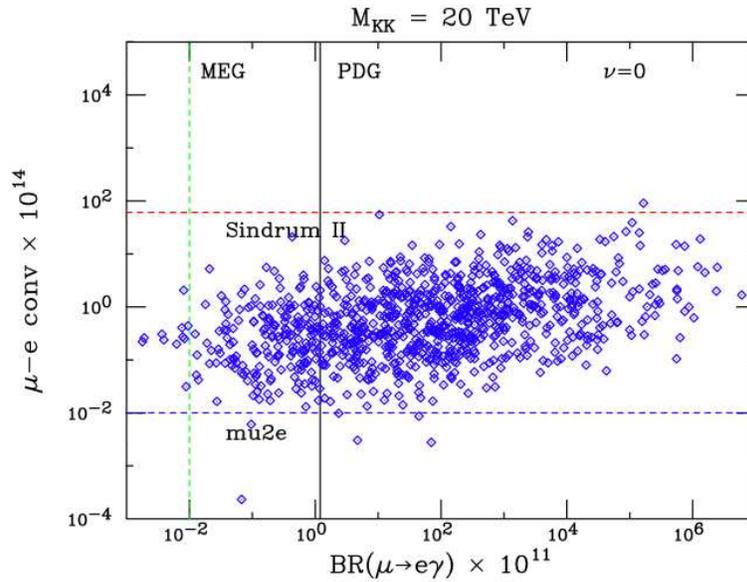


Figure 1.4: $\mu + {}^{48}\text{Ti} \rightarrow e + {}^{48}\text{Ti}$ rate as a function of $\text{Br}(\mu \rightarrow e + \gamma)$ for the Randall-Sundrum model with one warped, compact extra dimension, in the scenario where the Higgs boson is allowed to propagate in the bulk.

To summarize the potential of Mu2e, it is worth noting the range of possibilities for LFV discovery under various scenarios for MEG results. Should MEG see a signal at the level of 10^{-13} , Mu2e would see anywhere from roughly 100 to 20,000 events depending on the underlying BSM dynamics. Should MEG see no events, viable BSM scenarios still predict up to 20,000 events observable by Mu2e. It is also worth emphasizing that the event signature of a single particle separated from background is amenable to future upgrades, unlike other LFV signatures that require observation of particles in coincidence and thus often subject to irreducible combinatorial background. Finally, in the event of a positive Mu2e observation, additional information on BSM dynamics can be obtained by measurements using stopping targets with different atomic numbers.

The importance of this measurement has been widely acknowledged for some time, and an ambitious experiment was proposed for the AGS at Brookhaven. Called the Muon to Electron CONversion (MECO) experiment, it was based on a muon beam line concept that was first proposed for the MELC experiment [20] at the Russian Institute for Nuclear Research (INR) a few years earlier. In the end, the cancellation MECO was driven mainly by factors at Brookhaven which were outside of the control of the collaboration. The final review of the MECO experiment contained strong praise both for the physics goals and the viability of the proposed experimental technique [18]. More recently, a steering committee convened at Fermilab highlighted this experiment as a promising near term opportunity for the lab [19].

2. Experimental Technique

The baseline for the muon beam line and the detector system in this letter is the MECO design, which is illustrated in Figure 2.1. MECO was designed to search for the signature of a captured muon converting to an electron through the exchange of virtual particles with an aluminum nucleus. The MECO experiment is extensively documented in its technical proposal [21]. We do, however, expect to examine areas of possible improvement based on technical advances since that time.

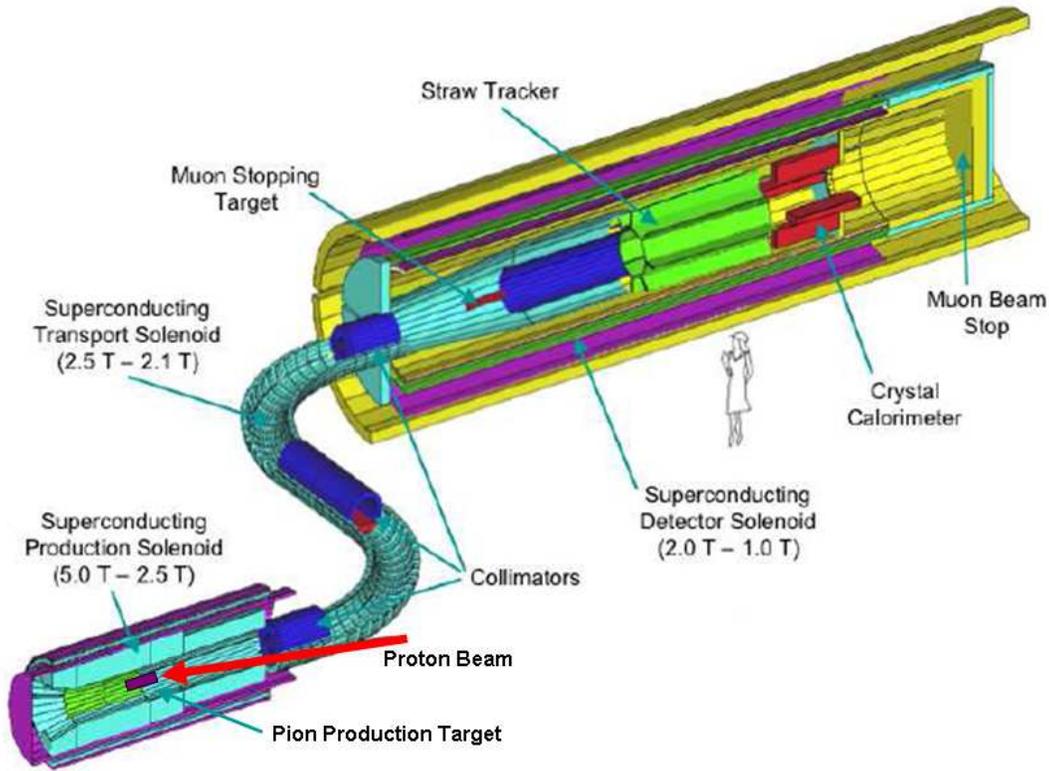


Figure 2.1: The Mu2e layout, taken directly from MECO. The primary proton beam enters from the right at the junction between the Production Solenoid and the Transport Solenoid, as indicated, and strikes the production target. The Transport Solenoid captures both backscattered pions/muons and those that are reflected from the high field region of the Production Solenoid. These are transported to the stopping target, within the Detector Solenoid. The detector is designed to kinematically reject electrons from ordinary muon decays.

2.1 Event Topology

Low energy negative muons are brought to rest in a stopping target. In a very rapid process, muons undergo an atomic cascade arriving in the $1s$ state of the aluminum target nucleus. To a good approximation, the muon has a hydrogen-like wavefunction bound to a charge Z , since its Bohr radius is far smaller than those of the atomic electrons. There are three main reactions: 1) Decay-in-orbit (DIO), where the muon decays while in atomic orbit, 2) the muon is captured by the nucleus, obeying lepton number conservation, and 3) the lepton flavor-violating (LFV) channel of interest, in which the muon undergoes neutrinoless muon to electron conversion in the field of the nucleus, producing a monoenergetic electron (105.0 MeV for muonic aluminum).

A *free* muon decays to electrons with a maximum energy of 53 MeV. A muon *bound in atomic orbit* can decay with the nucleus absorbing a significant portion of the momentum. In that case, the decay electron has a maximum energy at the same energy $E_e \sim 105$ MeV as the conversion electron, posing a background threat to the conversion electron signal. However the overwhelming majority of the decay electrons still have energies below 53 MeV. The probability

distribution falls very rapidly by the factor $(E_e - E)^5$ near the endpoint, and potential backgrounds can be overcome with sufficiently good electron energy resolution, as discussed below.

The probability of muon capture on the nucleus, $\mu + A(N, Z) \rightarrow A'(N', Z') + \nu_\mu + a\gamma + bn + cp$, is proportional to Z^4 , the number of protons in the nucleus (Z) times the probability of overlap between the nuclear and muon wavefunctions (Z^3 for the hydrogen-like 1s state). On average, $a \sim 2$, $b \sim 2$, and $c \sim 0.1$, and since the proposed muon capture rate is on the order of 10^{10} Hz, these gammas, protons and neutrons are a significant potential source of background to the detector system, part of the ‘delayed’ background described below. The muon decay and capture rates are roughly equal for muonic aluminum, leading to a lifetime in muonic aluminum of 0.88 μsec .

The $\mu + N \rightarrow e + N$ reaction occurs *coherently* on the nucleus, and consequently the probability varies as $Z^2 Z^3 \sim Z^5$ for low Z . Therefore $R_{\mu e}$, the rate of conversion divided by the rate of capture, increases proportionally to Z . This approximation breaks down at $Z > 40-50$ due to nuclear form factor effects, and at very high Z the ratio actually declines. Titanium has a factor 1.7 advantage over aluminum; however, its lifetime is smaller at 0.33 microseconds. Mu2e plans to use a pulsed muon beam. It has been determined in simulations that it is necessary to wait 600-700 ns after injection before data collection begins in order to suppress backgrounds, therefore aluminum is chosen as the stopping target because of its more suitable muonic atom lifetime.

The monoenergetic conversion electron has an energy well above most of the copious flux of electrons and gammas which arise from muon decays, greatly facilitating the isolation of the signal from backgrounds. This is a clear experimental advantage of the conversion process over the $\mu \rightarrow e + \gamma$ reaction. The latter reaction also produces a monoenergetic electron, but its energy is 53 MeV, where there is a large electron background from ordinary muon decay. A 53 MeV gamma ray must be detected in time coincidence with the electron, which limits data rates, and likely limits the ultimate statistical sensitivity of $\mu \rightarrow e + \gamma$ to LFV.

2.2 Principal Backgrounds

Some important information on expected backgrounds comes from the experimental experiences of previous $\mu + N \rightarrow e + N$ searches. As mentioned earlier, the SINDRUM II experiment at PSI established the current best limit, $R_{\mu e} < 4.3 \times 10^{-12}$ on titanium [22]. The electron energy spectra obtained before and after background suppression are shown in Figure 2.2. A simulated conversion electron peak, after folding with the SINDRUM II energy resolution and supposing that $R_{\mu e} = 4 \times 10^{-12}$, is indicated by the white dotted line. The ‘prompt background’ in Figure 2.2 within the region of the conversion peak consists of electron background from particles originating at the primary production target which cause high energy electron background when arriving in the region of the stopping target. Such a process can look like a conversion electron if it appears to come from near the target. The most important sources of prompt background are:

1. Electrons from in-flight muons which decay near the stopping target. These electrons can have $E > 100$ MeV if the muon has a momentum $p > 77$ MeV/c. This can look like a conversion electron if it scatters from the stopping target.

2. Secondary electrons produced when a beam pion stops in material in or around the stopping target and is radiatively captured; the radiated photons (energy up to ~ 139 MeV) can produce electrons in the energy range of the conversion electron via pair production in the stopping target or other surrounding material. This can look like a conversion electron if the pair production occurs in the stopping target.
3. Energetic electrons produced secondarily from antiproton annihilations along the muon beam line. The flux of antiprotons produced by the proposed 8 GeV proton beam is small, but not zero. Those antiprotons which are transported along the muon beam line with momenta similar to the muons have very low energies. A very thin window placed in the muon transport line would absorb these antiprotons while having little impact on the muons.
4. Beam electrons incident on the muon target and scattering into the detector region. Transmission of electrons along the muon beam line with $E > 100$ MeV must be highly suppressed.

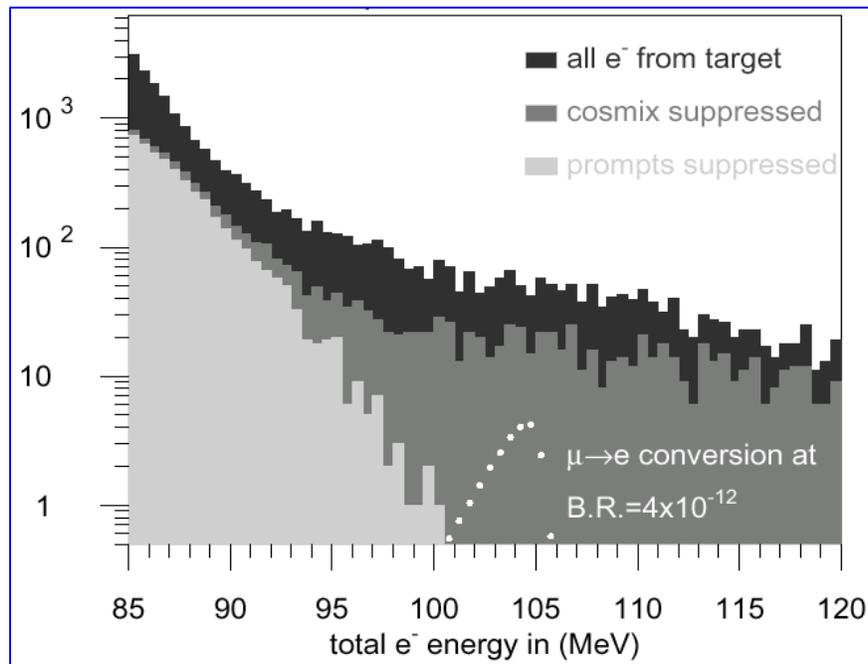


Figure 2.2: Histogram of the electron energy distribution from the SINDRUM II [22] $\mu + N \rightarrow e + N$ experiment. The darkly shaded region includes all electrons from the target, the medium shaded region represents the effect of the cosmic ray cuts, and the light gray region reveals the distribution after prompt and cosmic ray cuts. Also indicated is a simulated conversion electron peak, assuming $R_{\mu e} = 4 \times 10^{-12}$. SINDRUM II is not background limited.

SINDRUM II eliminated prompt backgrounds by vetoing any candidate conversion electrons that were in time coincidence with particles entering the detector from the muon beam line. This coincidence requirement limited the rate and therefore the ultimate statistical sensitivity of the SINDRUM II experiment. They had no antiprotons since their proton energy was $E < 600$ MeV. As shown, the veto is quite effective at eliminating the prompt background. Similarly, veto counters and event reconstruction cuts are very effective at eliminating the cosmic ray

background. The resulting spectrum has *no counts* under the conversion electron peak. Clearly, this suggests that the limit on $R_{\mu e}$ could be much improved if a much larger data sample were available.

There is a third component of the background apparent in Figure 2.2: the electrons from muon decay in orbit (DIO). As previously mentioned, the flux of such electrons varies with $(E_e - E)^5$ near the endpoint energy E_e . This background can be eliminated with sufficiently good electron energy resolution. For the proposed Mu2e sensitivity, $R_{\mu e} \sim 10^{-16}$, simulations show that 1 MeV FWHM electron energy resolution will give 0.25 DIO background counts compared to 5 conversion events at the design value of 1×10^{18} integrated muon stops (corresponding to 4×10^{20} protons), in the energy range 103.6-105.2 MeV.

Apart from the DIO electrons, other kinds of ‘delayed’ background consist mostly of low-energy neutrons, protons and gammas due to muon captures in the target. These can produce a large flux of low energy background in the detectors. There are two potential detector problems introduced by this background: the possibility of fake tracks leading to false conversion electrons, and detector dead time. The detector design described in later sections addresses these potential problems.

Occasionally, muon capture can produce a high energy gamma ray via radiative muon capture. The gamma could convert in material into a high energy electron, with a maximum energy of 102.5 MeV in aluminum, which is ~ 2.5 MeV below the conversion electron energy. This background is added to the DIO background. If the detector energy resolution is sufficient to handle the DIO background then it is automatically sufficient to handle the lower energy muon radiative decay background. Adequate suppression of cosmic ray background will require passive shielding as well as an active hermetic scintillator veto counter system.

To achieve the muon flux required by Mu2e in a reasonable amount of beam time, it is necessary to go to much higher muon stopping rates than previous experiments. Yet it is also necessary to continue to control the prompt background. The Mu2e experiment plans to inject muons onto the stopping target in narrow (< 200 ns) time bursts, separated by intervals of about 1.5 microseconds, somewhat larger than the lifetime of muonic aluminum. The veto counter used in SINDRUM II and its associated rate limitation will be eliminated. Muon to electron conversion data would be taken between bursts, after waiting a sufficient time following an injection burst (~ 700 ns) for the prompt background to subside. A suppression of the primary proton beam between bursts by a factor of 10^9 relative to the burst (extinction) is necessary to control the prompt background. The Mu2e experiment outlined in the next sections will produce and be able to handle much higher muon stopping rates than SINDRUM II: about 10^{11} Hz versus about 10^7 Hz.

3. The Proton Beam

3.1 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 [23] are implemented, there should be an excess proton capacity of this order even during the NOvA era [24]. 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 3.1. 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 μsec 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.

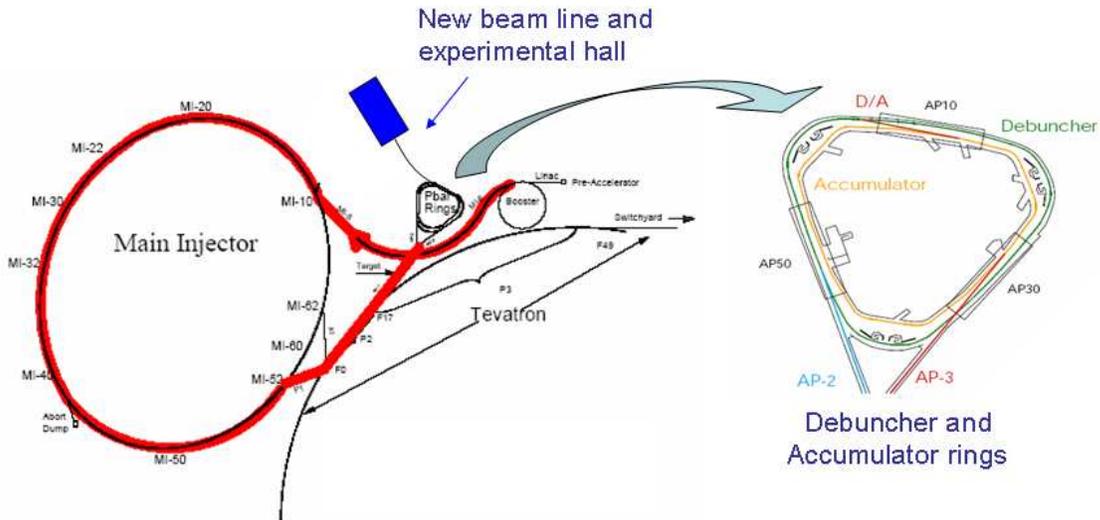


Figure 3.1: The relevant parts of the accelerator complex are shown. At right is a detail of the antiproton ring, showing the Debuncher (outer ring) and Accumulator (inner ring). The complete path taken by the protons in the proposed "boomerang scheme" is shown at left.

3.2 Delivering Protons to the Accumulator (“Boomerang Scheme”)

Initially, it was believed that injecting protons into the Accumulator during NOvA running would require a new transport line to be built from the Booster to the Debuncher/Accumulator enclosure. However, an ingenious “boomerang scheme” [26] 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 kind. This will greatly reduce the burden on the accelerator complex of supporting this experiment.

A key part of the NOvA accelerator upgrades will be a modification to the MI-8 beam line to allow beam to be directly injected into the Recycler [25]. 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 NOvA plan, 12 Booster batches will be stacked into the Recycler during every 1.333 second Main Injector cycle. The NOvA timeline is shown in Figure 3.2. 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.

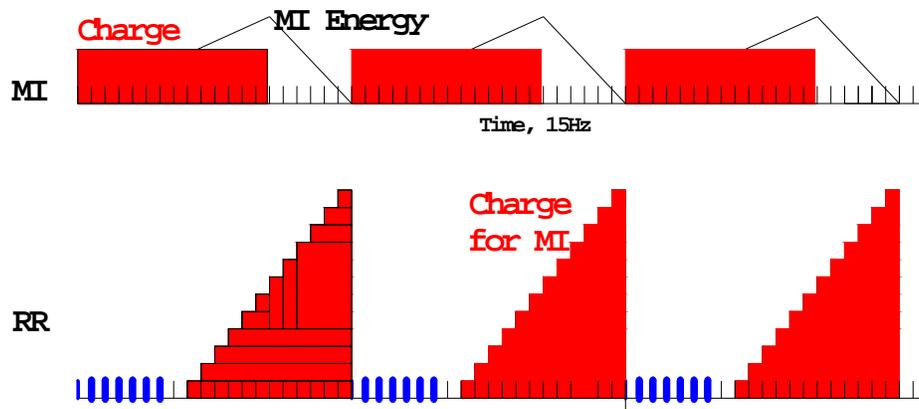


Figure 3.2: A schematic illustration of the timeline for 15 Hz Booster batches in the NOvA era. NOvA proton batches are shown in red. Twelve Booster batches are stacked in the Recycler and then transferred all at once to the Main Injector, eliminating the loading time and increasing protons to the NuMI line. Shown in blue are the unused Booster batches available while the Main Injector is ramping.

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 [26]. 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 3.3 shows a proposed location for the extraction from the Recycler as well as beam matching to the existing P150 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 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 this¹. 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 year², and we will assume this in our baseline planning.

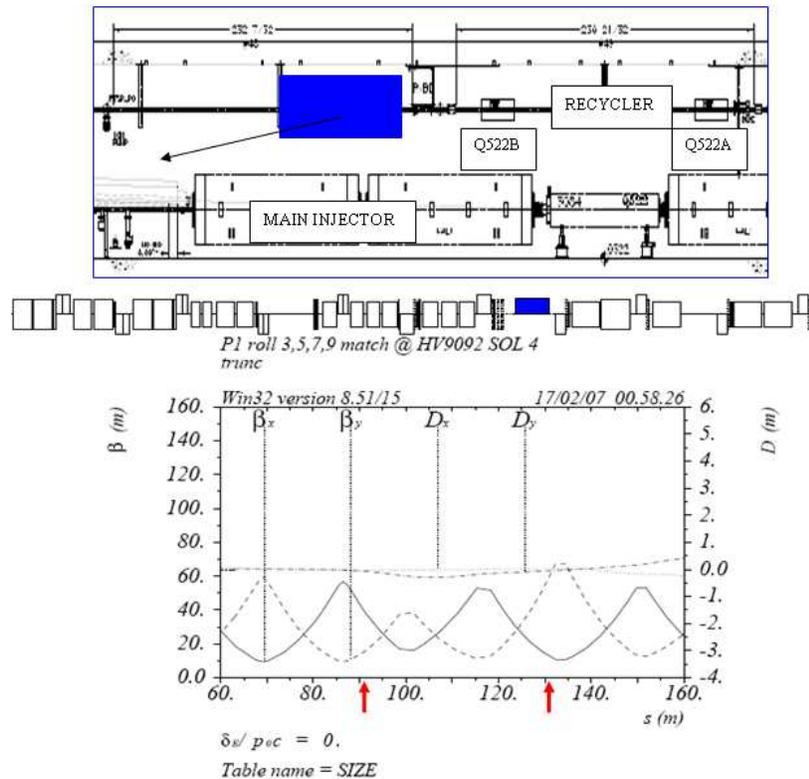


Figure 3.3: The top figure shows an elevation view of the proposed new extraction area in the Recycler. The bottom part of the figure shows the insertion section and lattice matching into the P150 line.

3.3 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 [27]. Figure 3.4 illustrates the scheme. Protons are injected at an

¹ The total Booster output is limited by the maximum average repetition rate that the RF system can support and by radiation issues due to beam loss. Making the Booster robust to 15 Hz will require at most modest improvements over what is planned for the NOvA project. Radiation will be reduced by improving Booster efficiency through improved orbit control and extraction handling.

² These annual proton rates are calculated based on the stated batch sizes and timelines, assuming the canonical “Snowmass Year” of 2×10^7 seconds, and should be used only to set the scale. The Proposal will include more accurate projections.

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.

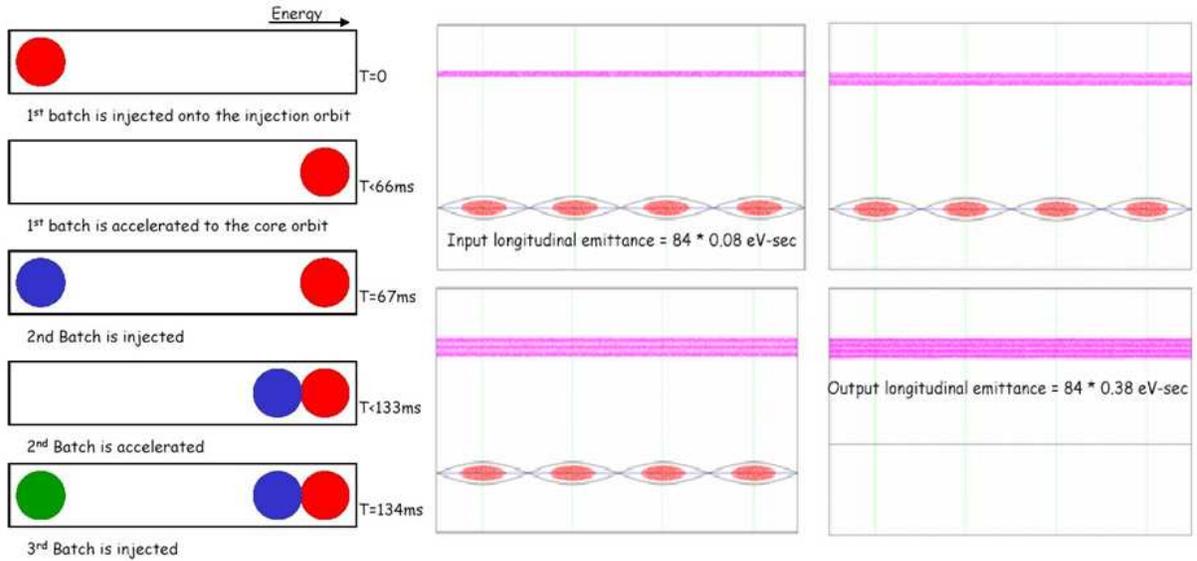


Figure 3.4: The left figure shows a conceptual schematic of momentum stacking. On the right is shown a simulation of the capture and momentum stacking of four Booster batches.

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 3.5 [28]. 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.

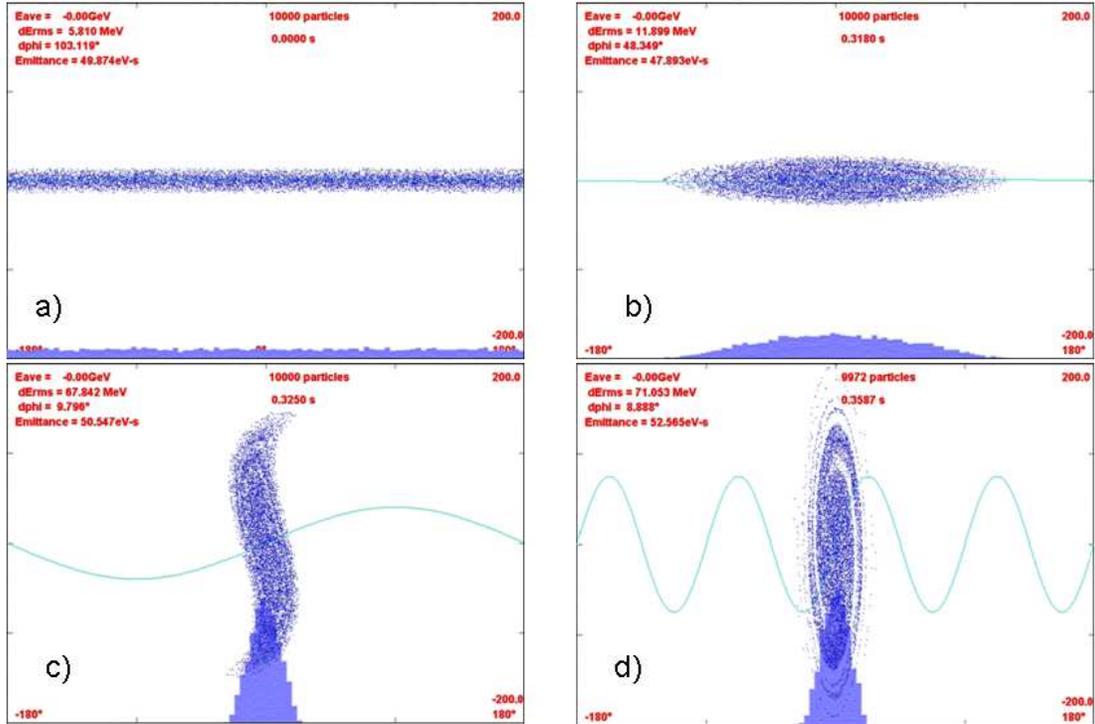


Figure 3.5: The hybrid rebunching scheme. Figures a) and b) show the $h=1$ capture in the Accumulator. After this step, particles are transferred to the Debuncher, where c) shows the $h=1$ phase rotation. The final $h=4$ capture is shown in d).

3.4 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 much more 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 [29].

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 3.6 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 ($\pm 5''$ extraction channel) followed by a 2 m C-magnet.

The existing tune working point of the Debuncher is $v_x/v_y=9.764/9.785$, making the $v_x=29/3$ a logical resonance to exploit for slow extraction. Table 3.1 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 3.7 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 extraction septum is superimposed.

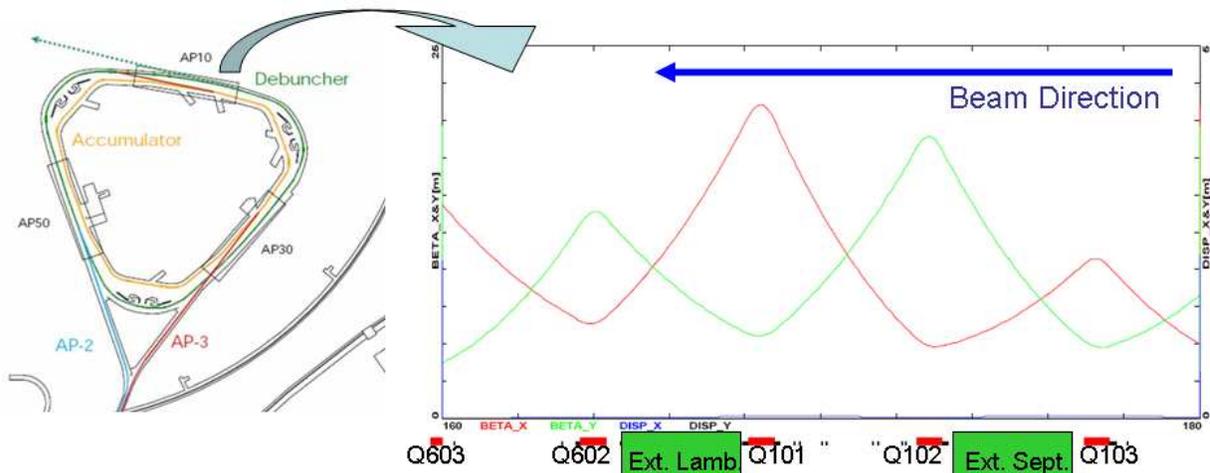


Figure 3.6: The proposed extraction region beneath the AP-10 service building. The extraction septum would be located between the Q103 and Q102 quadrupoles followed by Lambertson and C-magnet between the Q101 and Q602 quadrupoles.

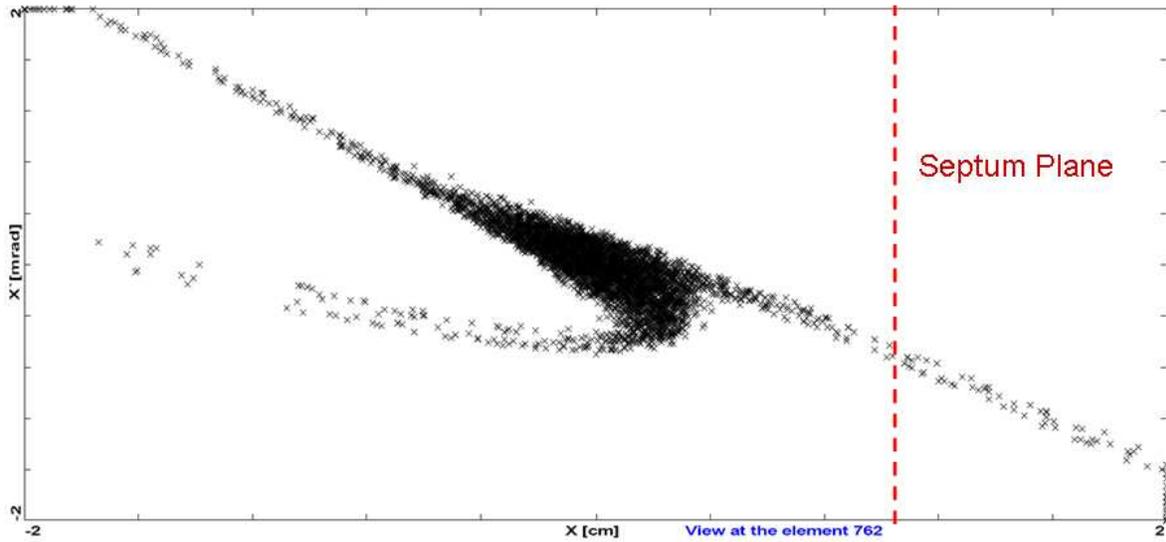


Figure 3.7: Preliminary simulation of a third order resonance in Debuncher. The position of the extraction septum is superimposed.

Resonant Extraction Parameters	
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 at electrostatic septum (m)	15
β at Lambertson (m)	22
β at harmonic quads (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	.048
Septum field (MV/m)	8
Septum length (m)	3

Table 3.1: The approximate parameters of the third order resonant extraction, with the septum located between the Q101 and Q602 quadrupoles.

3.5 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 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 $\sim \tau_\mu$, ensuring near perfect proton extinction between bunches, and restricting the search for the conversion of stopped muons to the inter-bunch period.

Radiative π Capture	0.08
Scattered Electrons	0.04
μ Decay in Flight	0.08
π Decay in Flight	<0.004
Total	0.2

Table 3.2: The most significant beam backgrounds, for 4×10^{20} protons on the primary target and an extinction factor of 10^{-9} . There would be 5 signal events if $R_{\mu e} \sim 10^{-16}$.

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 $\sim 1 \mu\text{s}$ intervals in which the experiment is live. Extensive GEANT simulations of these background processes were performed for MECO [30]. Table 3.2 shows the numbers of events produced by each of these backgrounds for 4×10^{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 5 events for $R_{\mu e} = 10^{-16}$.

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 3.8. 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. We are also investigating other ideas for achieving the necessary levels of extinction. For example, an electron lens, similar to the one used in the Tevatron [31] could be used to selectively move only the tune of the extracted bunch near resonance during extraction.

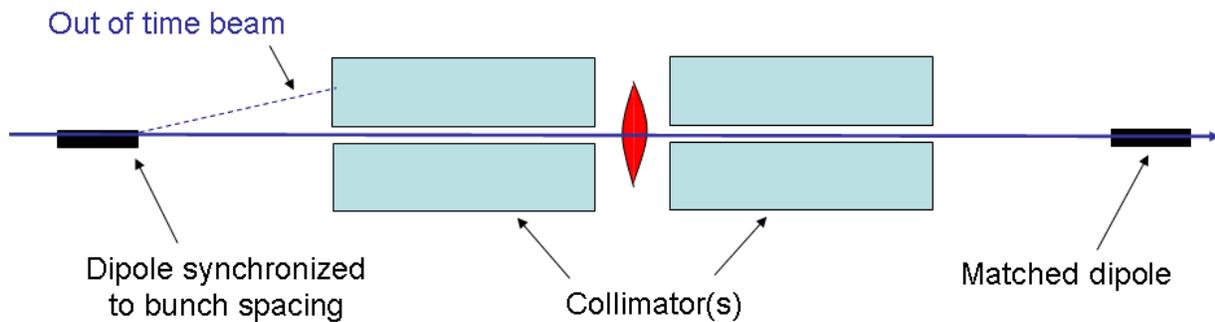


Figure 3.8: The AC dipole scheme for beam extinction. A matched pair of dipoles is synchronized with extracted beam bunches such that only in-time beam can pass through a collimator or series of collimators.

3.6 Project Siting

Unfortunately, there is no readily useable space to locate the detector for this experiment, so in all likelihood, this project will involve developing a green-field site. Figure 3.9 shows the proposed location, just off of Giese Road, near the SciBooNE detector hall. Wetland mitigation issues should be minimal for this location [32]. Also, the proximity to the SciBooNE hall and the MiniBooNE target building should reduce the cost of providing conventional services to the building. This effort could be coordinated with other potential users of this facility to develop a more general use experimental space.

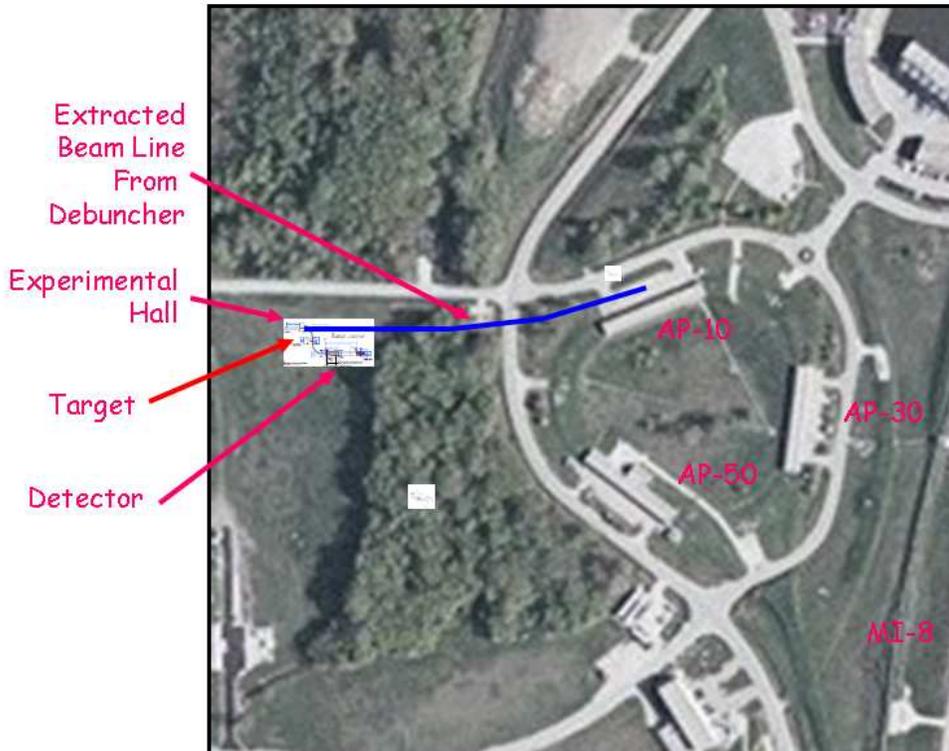


Figure 3.9: Proposed location of the experimental hall.

3.7 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 NOvA program. This will require two things:

- The Booster will have to run beyond the 10.5 Hz average repetition rate required by the NOvA 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 [33] 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 department'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 NOvA era. At that point, protons to the NOvA 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 [34]. 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 4×10^{20} protons needed by the experiment in something on the order of two years.

3.8 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 13' in most outside areas and is 10' or less in the service buildings. Tunnel roof loading limits preclude increasing this overburden, which is more than 10' short of the passive shielding requirements (“Cossairt Criteria”) 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 the documentation needed for the shielding assessment of such a configuration is immense, it is important to begin working on it as soon as possible.

3.9 Mu2e and Future Accelerator Upgrades

All of our discussions so far assume that the accelerator complex implements the baseline modifications proposed to support the NOvA experiment [35] following the termination of the collider program. In this scenario, Mu2e uses the excess capacity of the Proton Source without impacting the NOvA program.

It is reasonable to assume that this will remain the running configuration at the lab until at least 2014 or so. In the event the ILC is built on an aggressive timescale, and sited at Fermilab, the lab will likely remain in the basic NOvA configuration at least through the 2010's.

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 [36]:

- 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 more 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 H^- 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.

3.9.1 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 NOvA [37]. 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 NOvA 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 NOvA (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 NOvA 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.

3.9.2 Project X

The 8 GeV linac proposed for Project X would be based on ILC cryomodules and RF systems [38]. 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 NOvA 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 [39], 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 [40]. 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. This issue will be studied in more detail for the proposal.

Increased beam flux also has implications for the target and detector. During the next year, we will study these issues to determine whether increased capacity will be designed into the detector from the beginning, or staged in later as an upgrade.

4. Experimental Apparatus

4.1 Design Overview

As a baseline and for the purposes of this LOI, we have adopted without change the MECO apparatus, as originally proposed for the Brookhaven AGS [21]. The slightly different beam conditions at Fermilab do not change the conclusions significantly.

The design of the Mu2e experiment centers on maximizing the stopped muon flux and optimizing the detection efficiency and energy resolution for the 105 MeV (muonic aluminum) conversion electrons, while minimizing the acceptance of background particles. As described in the previous section, a pulsed (<200 ns) 8 GeV proton beam with a period of around 1.5 μ s produces bunches of muons which are transported and stopped in the aluminum stopping target. The measurement period would start about 700 ns after the proton pulse, when backgrounds from the proton pulse have died away sufficiently, and would continue until the next proton bunch injection. The measurement period of ~ 700 -1700 ns after injection is a good match to the 0.88 μ s lifetime of muonic aluminum.

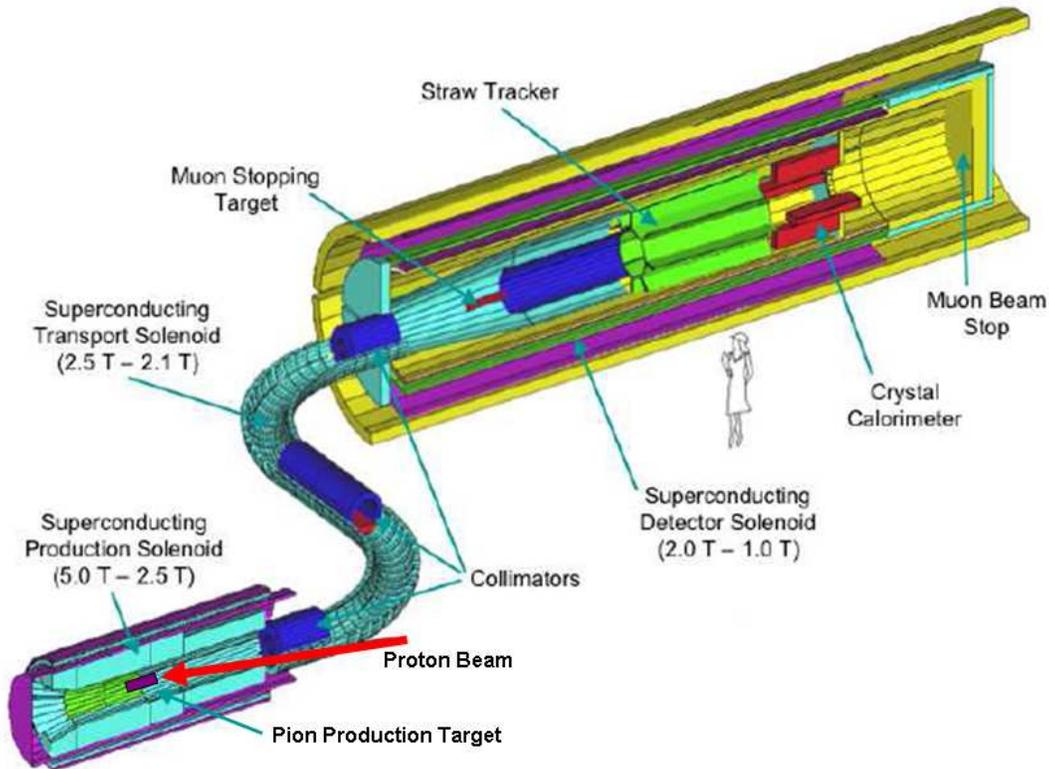


Figure 4.1: Layout of the Mu2e apparatus, showing the major components. Produced pions and decay muons are reflected by the high field end of the Production Solenoid. The Transport Solenoid and collimation system serve to transport momentum-selected muons for efficient absorption by the stopping target. The detector is designed to reject ordinary muon decay electrons by confining their trajectories to a lower radius than the tracking elements. Not shown is the antiproton stopping window, located where the upstream and downstream Transport Solenoids meet.

Figure 4.1 shows the main components of the Mu2e apparatus. Negative muons are brought from the production target to the stopping target along three successive solenoids: the Production Solenoid (PS), the Transport Solenoid (TS), and the Detector Solenoid (DS). A continuous negative magnetic field gradient is maintained along all straight solenoid sections. For adiabatic motion in a straight solenoid with a field gradient, p_t^2/B is a constant. Thus p_t/p decreases as the particle moves downstream, in effect accelerating them toward the stopping target. When an upstream moving muon spirals from a low field to a high field, it can be reflected back toward the low field region if the gradient is large enough.

The production target is placed near the middle of the Production Solenoid, with the proton beam directed upstream in order to avoid beam spray toward the downstream magnets. The Production Solenoid has a strong field gradient, varying from 5 T upstream to 2.5 T downstream. This follows the scheme originally proposed for the MELC experiment [20] in Russia (the precursor to the MECO experiment). The gradient collects and accelerates muons and pions in the downstream direction, and reflects a portion of muons and pions initially traveling upstream back downstream.

The S-shaped Transport Solenoid connects the downstream end of the Production Solenoid to the upstream end of the Detector Solenoid. It consists of three straight solenoidal sections connected

by two 90-degree toroidal bends. The upstream field is 2.5 T and the downstream is 2.0 T. The gentle gradient in the straight sections serves to prevent temporary local trapping of particles with small longitudinal momentum components. These could lead to particles arriving at the stopping target late enough to fall into the measurement time window. The bends in the Transport Solenoid prevent a straight line path between the production and stopping target/detector region. Moreover, the S-shape serves to displace the helical trajectories of charged particles vertically in a direction which depends on the sign of the charge and by an amount which depends on the transverse and longitudinal momentum components. Circular collimators are placed in the entrance and exit straight sections to greatly reduce the flux of electrons and muons with high transverse momenta. An asymmetric collimator is placed in the central straight section. The collimators remove almost all of the positively charged particles and reject high momentum electrons, muons, and pions with high efficiency. A thin beryllium window is placed at the center of the Transport Solenoid to absorb antiprotons.

The upstream portion of the Detector Solenoid contains the aluminum stopping target, in the center of a field which decreases from 2 T to 1 T, upstream to downstream. The gradient serves to reflect some of the candidate conversion electrons initially spiraling upstream back toward the detectors, thus increasing the detector acceptance. The aluminum target is a series of seventeen 0.02 cm thick disks placed perpendicular to the beam with 5 cm spacing, and radii tapering from 8.30 cm to 6.53 cm toward the downstream direction. The geometry is chosen in order to minimize energy straggling of the conversion electrons in the target, a major contributor to the electron energy resolution, and to maximize the number of muon stops and the detector acceptance of conversion electrons.

The detectors are located downstream from the stopping target in a region of a very uniform 1 T field. The displacement helps reduce the gamma and neutron background emanating from the stopping target. The detectors consist of a particle tracker and an electromagnetic calorimeter. There is no detector material at radii less than 38 cm. Charged particles with transverse momenta less than 55 MeV/c miss the detectors completely, spiraling into a beam dump downstream. The vast majority of the DIO electrons have momenta less than 55 MeV/c, as do most of the beam electrons and muons.

4.2 Principal Subsystems

4.2.1 Muon Beam Line

4.2.1.1 Introduction

Central to the muon beam line is the Superconducting Solenoid Magnet System: the Production Solenoid, shown in Figure 4.1, houses the primary pion production target, heat and radiation shield, upstream vacuum enclosure and proton beam exit window. The Transport Solenoid contains the collimators and antiproton stopping window, and the Detector Solenoid contains the muon stopping target, proton and neutron absorbers, beam stop, downstream vacuum enclosure and instrumentation feed-through bulkhead. The warm bore of the system is held at high vacuum and consists of two separate volumes; the Production Solenoid and upstream half of the Transport Solenoid (TSu), and the downstream half of the Transport Solenoid (TSd) and the

Detector Solenoid. These two volumes meet at the anti-proton stopping window. Vacuum is required in the detector solenoid mainly to limit backgrounds from muons stopping on gas atoms, followed by either muon decay or capture.

4.2.1.2 Solenoid Magnet System

The magnet system itself comprises the coils and cryostats, cryostat vacuum system, mountings, liquid helium system and cryogenic controls, power supply, quench detection and control system. The magnetic field specifications have been developed to meet the physics requirements discussed earlier; the field intensity varies from 5 T at the upstream end of the Production Solenoid to a constant 1 T at the detectors. The falling or constant field along the beam line has been designed to ensure that after the production target is hit, transit time for all particles is less than 650 ns, and that no particles are trapped in field minima leading to possible late arrivals at the stopping target. Thus, proton beam extinction is of paramount importance in ensuring the beam is off during the 1 μ s observation time window of the detectors.

A comprehensive Conceptual Design Report for the solenoid system [41] has been produced by the MIT Plasma Science and Fusion Center for the MECO experiment. The physics simulation, based on GEANT3, MARS, FLUKA, etc. of the full Mu2e setup uses the most recent magnetic field maps available, calculated from the engineering designs of the superconducting coils.

The magnet system is designed to: (1) maximize the number of low-energy muons brought to rest in the stopping target; (2) minimize the interactions of the exiting primary beam with the system; (3) minimize the flux of neutrons, protons, electrons, and other particles in the Detector Solenoid that could cause unacceptably high rates in detector elements; (4) minimize the number of particles, particularly pions and muons above 76 MeV/c and electrons above 100 MeV/c, that have a transit time to the DS greater than 650 ns from the time the primary beam hits the target; (5) maximize the acceptance for conversion electrons to intercept the tracker and trigger, while minimizing the flux of low energy photons, neutrons, and protons produced in the stopping target interacting in the detectors; and (6) Measure the momentum of conversion electrons with high precision without extensive corrections for inhomogeneous magnetic fields in the detector region.

4.2.1.3 Production Region

The primary 8 GeV/c proton beam enters the Production Solenoid pointing away from the Transport Solenoid and strikes the production target situated in an axially graded field, from 5 T to 2.5 T along the beam line. The target [42] is a solid 16 cm long gold cylinder of radius 3.0 mm and cooled by a turbulent water jet that flows through a 0.3 mm wide annular gap. This channel is formed by a 0.5 mm thick Ti cylindrical shell with its ends closed except for a Ti inlet and outlet pipe with 2.1 mm ID and 3.2 mm OD and 25 cm lengths which also serve as supports.

The steady state power limit of the Production Solenoid superconducting coils is 150W and instantaneous local heating must be $< 25 \mu$ W/g. The heat and radiation shield [43] of the Production Solenoid is designed to protect the superconducting coils from the punishing environment expected there. Engineering designs exist for dividing the over 40 ton shield, made of copper and tungsten, into a system of components weighing about 200 lbs each. The shape is

roughly a hollow cylinder with its narrowest inner radius at 25 cm. The anticipated heat load is 16 kW; with eight 1 cm diameter water-cooling channels located against the inner cryostat wall, a flow of three gallons per minute is expected to produce no more than a 5 K difference between the coolant inlet and outlet, while the maximum radial temperature difference will be no more than 40 K.

4.2.1.4 Transport Region

The Transport Solenoid system filters the particle flux producing a momentum- and charge-selected muon beam, with good reduction of μ^+ , π^+ , π^- , electrons, positrons, protons, and anti-protons, and favors low-energy μ^- . The negative muon beam has a high efficiency for stopping in the aluminum target, and the electron spectrum from muon decay-in-flight cuts off well below the muon conversion momentum (105.0 MeV/c in Al and 104.3 MeV/c in Ti). The S-shape of the Transport Solenoid also eliminates line-of-sight paths for neutrals in the Production Solenoid and Detector Solenoid regions.

4.2.1.5 Collimators

The first collimator [44], at the entrance to the Transport Solenoid, is a hollow copper cylinder with inner radius varying from 15-17 cm, outer fixed at 25 cm, and length 1 m. Significant shielding is required to protect the TSu (upstream section of TS) coils from heat and radiation. The stainless steel Transport Solenoid inner cryostat wall will be 2 cm thick and have a 24 cm radius, adding to the heat and radiation protection. As charged particles gyrate with small radii about field lines within the warm bore of the first quarter toroid magnet, positives drift perpendicular to the magnet plane to intercept the solid portion of the second copper collimator, distinguished by its off-center channel; negatives drift in the opposite direction. This second collimator channel is 0.8 m long and has a 20 cm radius at its widest point. The port geometry only permits passage of particles less than 5 cm above the magnet axis and less than 19 cm below it.

The downstream half of the S-shaped Transport Solenoid magnet (TSd) contains an identical center collimator making the off-center channel 1.84 m long. There is a 24 cm gap which is occupied by the Kapton Antiproton Stopping Window [45]; this device provides a barrier to antiprotons, a potentially serious source of physics background, and isolates the two vacuum volumes PS+TSu and DS+TSd, preventing radioactive atoms from contaminating the detector region. The second quarter toroid TSd restores the axial symmetry of the beam by imposing opposite drifts for each charge, now traveling in the same direction as it entered the Transport Solenoid. The transition from the Transport Solenoid to Detector Solenoid has a hollow cylindrical collimator of inner radius 12.8 cm, outer 25 cm, and length 1 m, made of boron- or lithium-loaded polyethylene, to protect the detector region from neutron backgrounds. There are also twelve 0.06 mm thick annular copper foils with 8.5 cm inner radius, 12.8 cm outer radius that serve to reduce particle noise rates in the tracker.

4.2.1.6 Muon Stopping Target and Absorbers

The goal of the stopping target design is to maximize the probability for beam muons to stop and for conversion electrons to be detected in the tracker and calorimeter. At the same time, the target should be designed to minimize both the energy loss of conversion electrons as they exit

the target and the number of electrons from muon DIO that reach the tracking detector. Furthermore, detector rates, e.g. beam electron bremsstrahlung, and backgrounds, e.g. cosmic ray interactions in the target, are minimized with the smallest possible target mass. The transverse size, thickness, spacing, and number of thin disks that comprise the target were optimized to best achieve these goals.

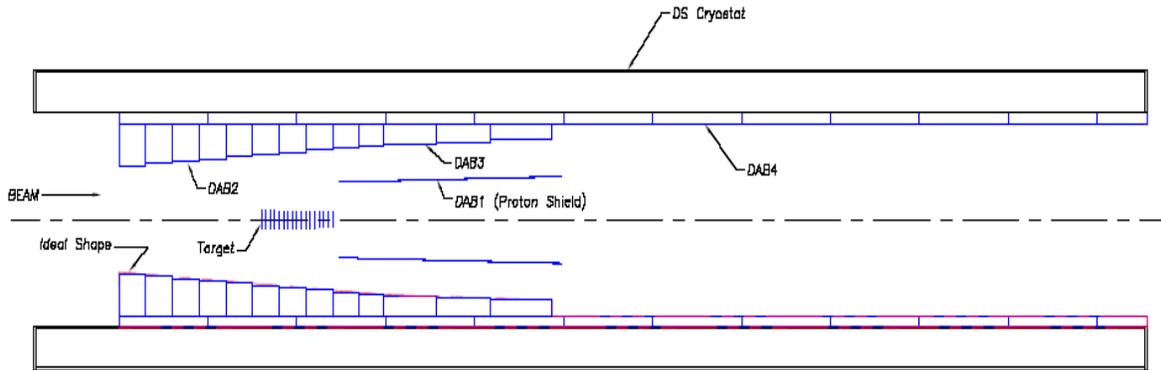


Figure 4.2: Schematic drawing of the stopping target and absorbers in the Downstream Solenoid. The proton shield is designated by DAB1 and the three neutron shields by DAB2-4.

The target, shown in Figure 4.2, with mass 159 g, has seventeen 0.02 cm thick aluminum disks. They are arranged parallel to each other, centered on the solenoid magnet axis and with each face perpendicular to it. The target is tapered in the downstream direction, with 5.0 cm disk spacing and radii from 8.30 cm to 6.53 cm. Titanium is also being considered, however a complete conceptual design for Ti does not yet exist. The detectors are at least 2.5 m away from the target, greatly reducing backgrounds from neutrons and gammas.

The target is centered in a graded axial magnetic field; with the first disk at 1.57 T and the last at 1.30 T. The gradient reflects electrons emitted upstream back to the detectors, resulting in about 60% of all conversions falling within the geometrical acceptance of the tracker and calorimeter. Conversion electrons with transverse momenta > 90 MeV/c are swept downstream by the decreasing magnetic field into the range 75 MeV/c to 90 MeV/c transverse momenta at the detectors. Beam particles with lower transverse momenta that do not scatter in the target pass down the center of the solenoid without intercepting the detectors. The graded field also ensures that electrons originating upstream of the gradient with 105 MeV/c arrive at the detectors with transverse momenta < 75 MeV/c, eliminating potential backgrounds.

The absorbers in the Detector Solenoid area [46] are designed to moderate protons and neutrons. Figure 4.2 shows a schematic view of the Detector Solenoid, revealing the boron or lithium loaded polyethylene proton and neutron absorbers in the muon beam line vacuum volume. The largest potential contribution to the tracker rate is from protons from muon capture in the stopping target. A proton shield consisting of a hollow conic section, 250 cm long, with average radius 36 cm, and 0.05 cm thick would extend from the end of the stopping target to the front of the tracker.

The neutron absorbers rest against the inner solenoid cryostat wall, at a radius 0.94 m. Neutrons from the stopping target are a potential source of background in the active cosmic ray veto counters that surround the Downstream Solenoid. A second neutron absorber between the outer DS cryostat wall and the 0.5 m thick iron cosmic ray shield (and magnetic flux return) is required to protect the cosmic ray veto shield.

4.2.2 Primary Detectors

The detector system is comprised of two units: a tracker and a calorimeter, as shown in Figure 4.1 and Figure 4.3. The tracker makes a highly precise measurement of the momentum of the rare decay electrons with energy near 105 MeV. After passing through the tracker, an electron deposits its energy in the calorimeter, which serves as the trigger for the data acquisition system. The detector system is designed to distinguish the desired electrons from the copious lower energy electrons from DIO. In addition, the system must operate in a high-rate environment created by the products from ordinary muon capture.

4.2.2.1 Tracking Detector

The tracker consists of straw tubes arranged in vanes as illustrated in Figure 4.3 and Figure 4.4. The Detector Solenoid magnetic field is designed so that traps electrons with $p_T < 55$ MeV/c at the detectors pass through the central vacuum region, bypassing the detectors. This, together with the magnetic gradient, insures that no electrons produced with $E < 53$ MeV can reach the tracker. Particles of higher p_T pass through the straw tubes and have their local position and direction measured as they helically travel downstream. The vanes ensure a minimum of three position measurements along a circular projection of the helix, and extend out to the largest radius for conversion events. The length of the tracker is designed to provide redundant position information for at least two helical orbits, a powerful advantage for pattern recognition. A view along the beam direction is shown in Figure 4.4 showing how particles of different momenta hit or miss the tracking detector, and show the structure of the straw tubes within the frames.

The tracker itself is made up of 2.6 m long straw tubes that reside in the vacuum oriented nearly parallel to the 1 T uniform magnetic field. The straws are 5 mm in diameter and 25 μ thick. The inner part of the tracker is an octagon, and there are eight vanes that extend to the largest radius where there might be signal events. The detector planes consist of three layers of hexagonal close packed straws, shown in the insert in Figure 4.4, to provide mechanical support and to facilitate pattern recognition. The outer straws are resistive so that the z-coordinate can be determined by signals induced on cathode strips etched on a kapton foil adjacent to the straws. The entire system has about 3000 straws and 16,000 cathode strips.

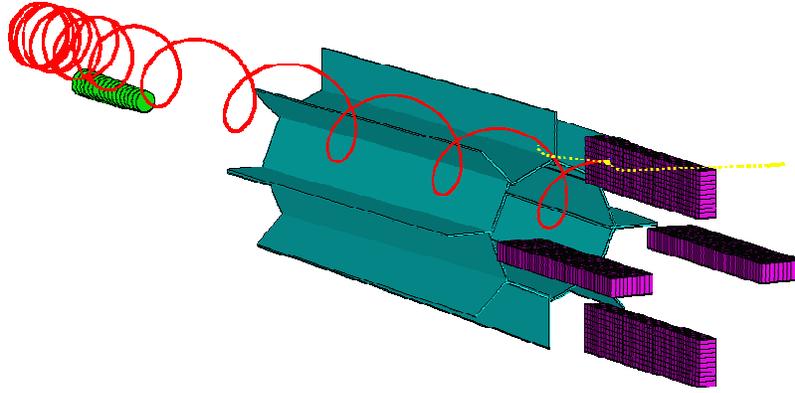


Figure 4.3: A perspective view of the MECO L-tracker showing the straw frames as they are azimuthally rotated and the position of the calorimeter.

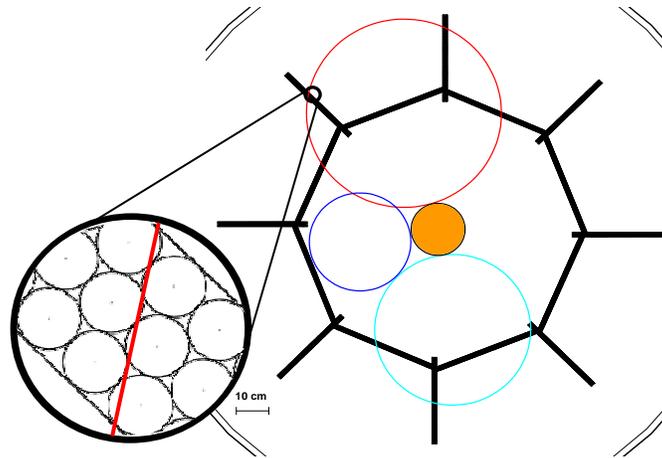


Figure 4.4: The L-tracker viewed along the axis of the Detector Solenoid. The solid yellow circle is the stopping target. The blue circle represents the largest circle from background electrons ($p_T = 55$ MeV/c). The red and cyan tracks are candidates of interest that will trigger the apparatus.

The above design, called the L-tracker, is one of two possibilities that we are evaluating. A second design, denoted the T-tracker because the straws are transverse to the magnetic field, is shown in Figure 4.5. There are 260 sub-planes made up of sixty 5 mm diameter conducting straws of length from 70-130 cm, totaling 13,000 channels. The motivation for the T-tracker is that the fabrication is more straightforward and the absence of cathode strips makes the electronics more robust. The disadvantage of the T-tracker is that the pattern recognition is more challenging.

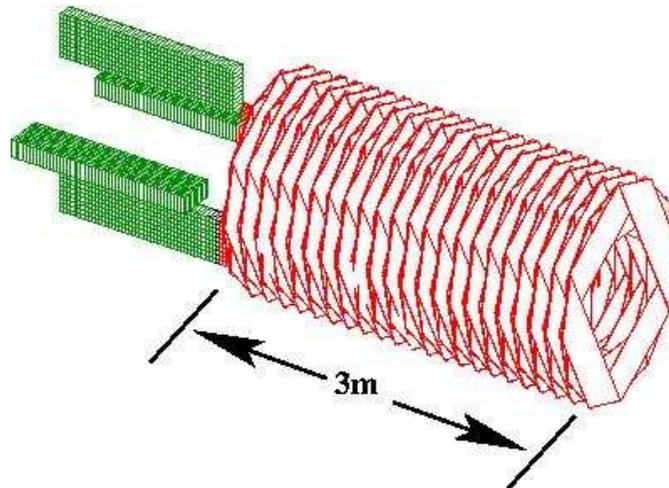


Figure 4.5: A perspective view of the T-tracker showing the straw frames as they are azimuthally rotated and the position of the calorimeter.

A GEANT simulation of the detector was performed to evaluate the ability of either tracker to (a) identify the signal and (b) measure the momentum with sufficient resolution. There are a number of processes that yield significant backgrounds, and these background rates depend critically on the details of the apparatus.

A summary of backgrounds vs. the elapsed time from when the proton beam impacts the production target is given in Figure 4.6. In Figure 4.6, there are very high rates at “early” times when the primary beam is on and the detector is dead. It is likely that the voltages on the chambers and calorimeter avalanche photodiodes (APD’s) will be reduced during injection in order to improve their recovery from the high rates encountered during injection. After 700 nsec, the rates are considerably lower, but the majority of the hits in the tracker are still due to background. One important example is low energy protons from the muon capture, which are heavily ionizing and produce large pulses which will provide a challenge for the electronics; crosstalk and extra dead time must be avoided.

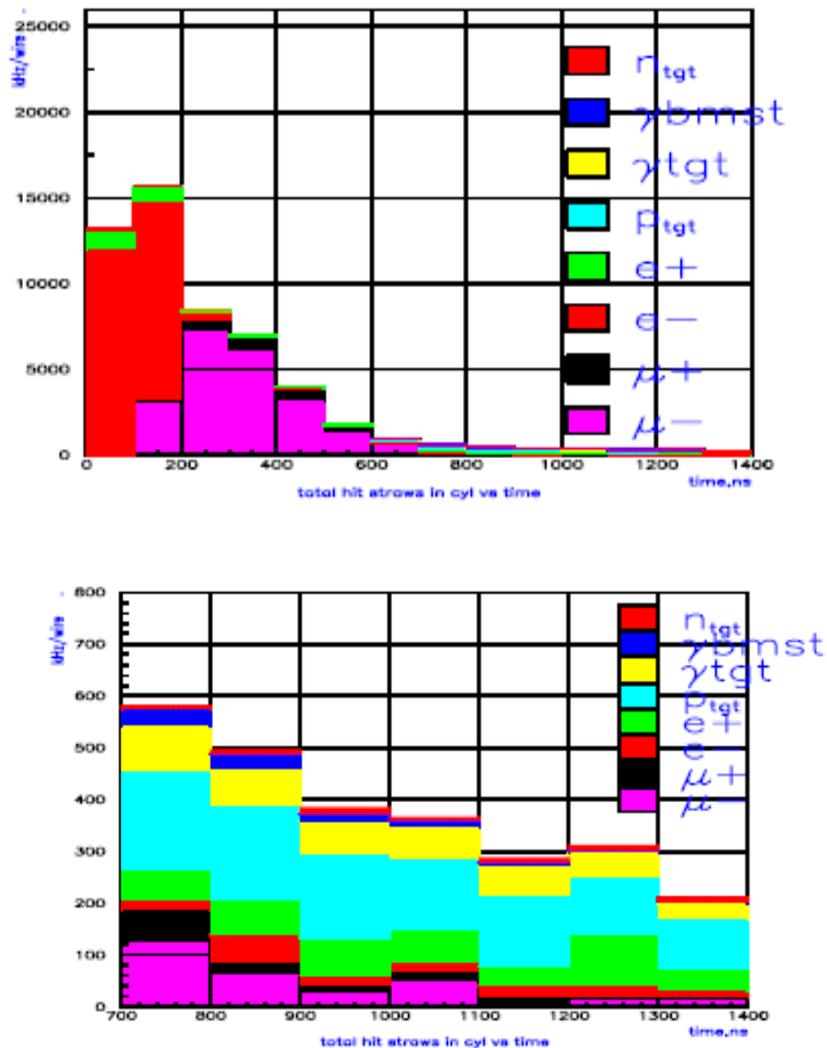


Figure 4.6: Background rates in the L-tracker. The top panel shows the rates at early times when the trigger is inactive. The bottom panel shows the rates during the live time of the apparatus. Photons, neutrons, and protons all make significant contributions. Note that “rtgt” is the total rate contribution from the stopping target, “ γ bmst” is the rate contribution from the muon beam stop, “ γ tgt” is the contribution due to gammas from the stopping target, “ptgt” is that of protons from the stopping target.

Typical results from the reconstruction of 105 MeV tracks are shown in the left panel of Figure 4.7. The resolution for the Gaussian part of the spectrum is 0.3 MeV FWHM. There is also a noticeable high energy tail. To convince ourselves that the tail will not bump DIO electrons into the signal window, we generated 10^5 events with energy above 102 MeV. The result was that <0.2 events appear in our acceptance window, as shown in the right panel of the figure. Similar results have been obtained for the L-tracker.

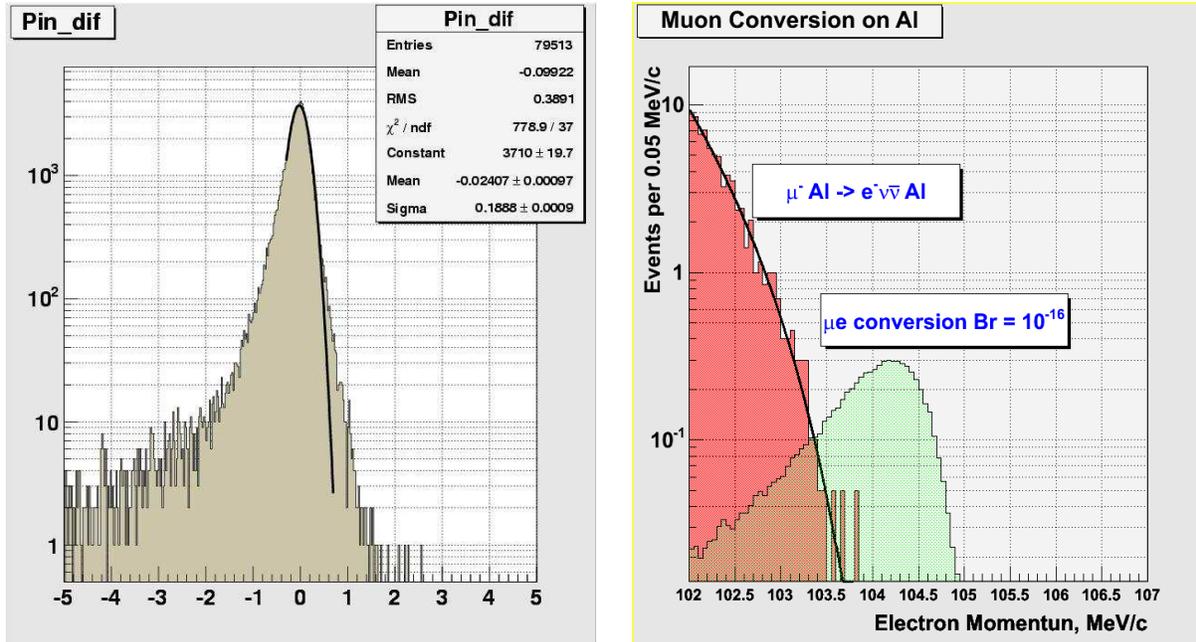


Figure 4.7: Left panel: Difference in the reconstructed momentum of events from 105 MeV/c. Right panel: Simulation of the experiment with ten times more running time and backgrounds shown. Only two events are in our fiducial region.

4.2.2.2 Electron Calorimeter

The Electron Calorimeter is designed to provide the event trigger for the experiment. This initiates data acquisition, in particular the readout of the tracker which provides the high resolution momentum measurement of the helical electron track. In addition, the calorimeter provides an energy measurement for the electron and a 3-dimensional position measurement for one point on the helical track. This calorimeter space-point measurement is essential for the pattern recognition of the T-tracker. While the energy resolution does not compete with the momentum resolution of the tracker, it does provide a redundant measurement.

The geometry of the calorimeter is shown in the right side of Figure 4.3. It is comprised of four vanes 120 cm long and 30 cm in the radial direction. The size is a compromise between cost and detection efficiency. The vanes are made of (lead tungstate) PbWO_4 crystals $3.75 \times 3.75 \times 12 \text{ cm}^3$. The depth corresponds to 13.5 radiation lengths and the transverse size matches the 2.2 cm Moliere radius of the showers. Lead tungstate is chosen as a compromise between cost, speed, and photon yield. A total of 1024 crystals make up the entire detector.

In order to detect the photons in the 1 T magnetic field, a pair of avalanche photodiodes (APD's) is mounted on each crystal. To improve the light yield of the lead tungstate and to minimize electronics noise, the system will operate at -24 degrees Celsius. The net resolution of the system is about 5%, with roughly equal contributions from shower statistics, photoelectron counting statistics, amplifier noise, and pileup from the copious low energy photons from the target. Tests of a crystal with cosmic rays have verified this level of performance. The dominant source of triggers is DIO events which have a rate which increases rapidly away from the endpoint energy. The design resolution of the calorimeter allows a trigger threshold of 80 MeV, low enough so

that it is very efficient yet high enough so that the trigger rate is manageable. The rapid falloff of the acceptance of the calorimeter with decreasing electron energy, shown in Figure 4.8, is a key feature that prevents tails of DIO events from swamping the trigger.

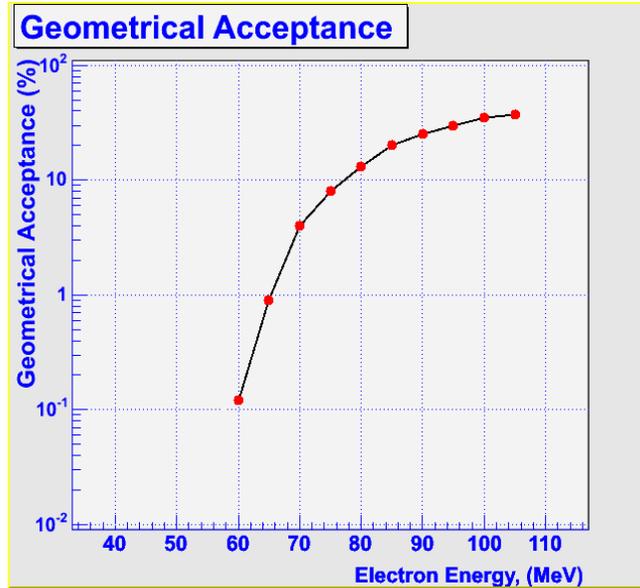


Figure 4.8: Acceptance of the calorimeter as a function of the energy of the electron. The rapid falloff with decreasing energy ensures that the high energy tails of the DIO electrons do not swamp the trigger.

4.2.2.3 Cosmic Ray Shield

Cosmic rays are a possible source of 100-110 MeV electrons that can look like valid conversion events in the detector. In order to reject electrons produced by cosmic rays, the apparatus will be surrounded by a hermetic Cosmic Ray Shield (CRS) constructed from scintillator bars. Monte Carlo calculations indicate that the veto inefficiency must be about 10^{-4} . In addition, the yoke of the Downstream Solenoid magnet and a meter of concrete blocks provide passive shielding.

Typical scintillator shields have an inefficiency of about 1%. To obtain the required rejection, a logical OR between two layers is planned. One of the challenges for the CRS is the presence of neutrons, which can be captured and produce a gamma ray which in turn Compton scatters in the scintillator.

The major sources of neutrons are the primary production target as well as the stopping target. The rates from these neutrons are so high that if the apparatus were vetoed by single hits, the live-time would be small. To reduce the sensitivity to the neutron-induced background, three layers of scintillator are planned and only events in which two out of the three layers fire will be vetoed.

A diagram of the apparatus with the CRS is shown in Figure 4.9. The entire outside of the magnet return yoke is covered. The basic element is a $1 \times 10 \times 460 \text{ cm}^3$ extruded scintillator bar, which is read out by three wavelength-shifting fibers. The scintillators are grouped into “modules” of 60 bars forming three 20-bar layers covering an area of $200 \times 460 \text{ cm}^2$. There would be a grand total of 3120 bars that cover the entire apparatus. The three shifter-fibers from

each scintillator bar are connected by optical fibers to a multi-anode PMT with 16 outputs. The $4 \times 4 \text{ mm}^2$ pixels for each anode are large enough to read out the three fibers for each bar. Prototype modules have been assembled and tested. The design has passed numerous tests.

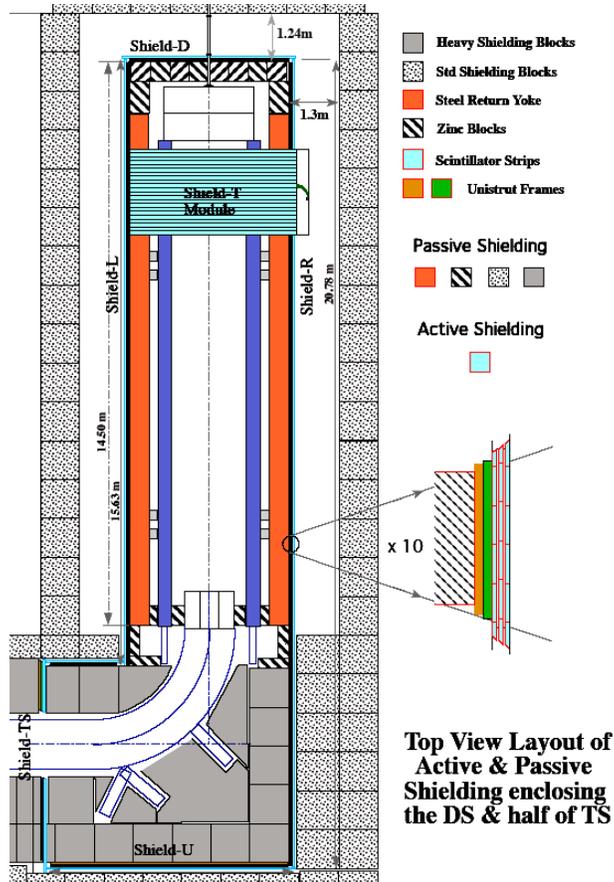


Figure 4.9: Diagram of the Cosmic Ray Shield (CRS). One full module is shown on the top. The expanded view on the right shows in detail the three layers and how the bars are overlapped.

4.2.3 Trigger, Data Acquisition and Computing

The Trigger and Data Acquisition (DAQ) system collects and organizes data from the detector subsystems and makes a decision about logging events to the storage medium. In addition, it provides online monitoring and control of the experimental apparatus and conditions.

A general schematic of the DAQ system is shown in Figure 4.10. The main elements of the system are the:

- **Master Clock:** Distribution system, which provides synchronization signals for the detector systems and the clock signals to the digitizers, and distributes the trigger signals.
- **Tracker Electronics:** Mounted inside vacuum near the detector, provides both pulse height and timing information about detector hits. The digitizing electronics is interchangeable between either L-Tracker or T-Tracker designs, and takes advantage of modern integrated circuit technology, placing a multi-channel IC at the front end of the

detector to amplify, shape, digitize, and buffer the signal. MECO pre-amplifier and digitizer designs were based on successful chips developed for ATLAS and BaBar experiments. These will need to be updated to modern CMOS processes. In total, tracker electronics consists of 12k to 20k channels depending on the detector architecture and digitizes signals at the clock rate of 50-100 MHz. The pipelined system allows for the extraction of 100-200 nsec buffers from the detector following a Level-1 (L1) trigger decision.

- **Calorimeter and CRS Electronics:** Provides amplified and pulse-shaped signals from the two detector systems. These signals are continuously digitized at 100 MHz clock rate and pipelined in a custom Calorimeter Digitizer Module (CDM).
- **Clock Trigger Distribution System:** Fans out the L1 trigger, a beam extraction signal, a system-wide clock (from which the Tracker digitizer and CDM sampling clocks are derived), an event counter and any other synchronization signals, system reset and initialization signals, etc.
- **Event Builder:** Accepts data from the detector subsystems, builds time-ordered events, and sends them to a processor farm via a network switch. The event builder will have to be able to cope with the total data of about 1 GByte/s.
- **Slow Data Monitors:** Include the Stopping Monitor, the Extinction Monitor, and the Magnet Monitor, would send their data asynchronously to the Data Logging Manager, to be merged into the data stream and also forwarded to online monitoring processes.
- **CPU Farm:** Consisting of the off-the-shelf processors, will run Level-3 (L3) trigger algorithms based on improved calorimeter reconstruction and calibration and limited tracking information. The purpose of the L3 systems is to reduce the logging rate by about a factor of 10.

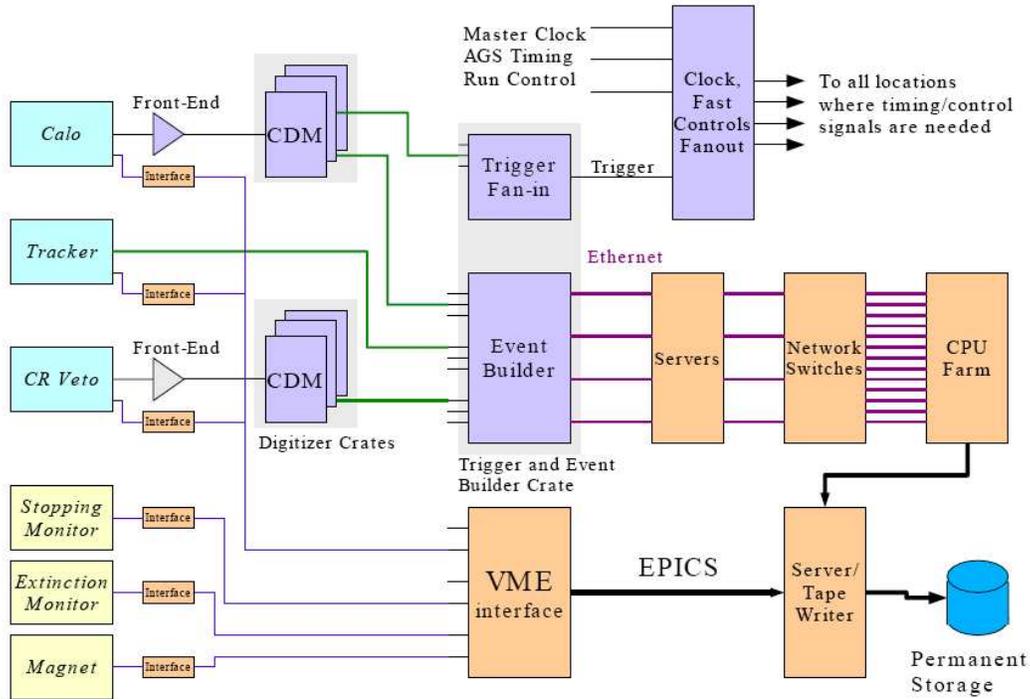


Figure 4.10: Schematic of the DAQ system.

The trigger system consists of two levels, hardware L1 trigger, based on the calorimeter information, and software L3 trigger, where simplified version of the tracking algorithms could be run. The present Monte Carlo estimates of event rate suggest the output rate of about 1 kHz from L1, which for a typical even size of 50 kB would correspond to the data rates of 50 MB/sec. L3 is expected to reduce this rate by a factor of 10 to a modest 5 MB/sec logging rate to the permanent storage.

A pipelined design of the Event Builder would allow significant headroom in data rates, should the detector occupancy fluctuate due to larger background rates or in case of an increase in beam intensity. The MECO design called for a 2 GB/sec capacity in the Event Builder, which was achievable with the current technology a few years ago. In addition, the pipelined architecture would allow us to insert an intermediate (L2) trigger layer should the need arise.

Allowing for about a factor of two headroom in data rates, the overall size of the data in permanent storage would be about 150 TB. A cluster of about 100 CPUs would be required to process the data through the full offline reconstruction in real time with minimal latency. We anticipate the need for a similar amount of computing power for the full Monte Carlo simulations of the apparatus.

5. Outlook

Our plan is to develop a technically defensible proposal to the Fermilab PAC in about one year, which would allow us to compete for funds to move on to a full technical design and ultimately

construct and operate the experiment. We will use the MECO apparatus as a starting point, but we expect that improvements will be possible; these will be incorporated into our final proposal. Significant resources will be required in order to complete the proposal in a timely manner over the next 12 months. In this section, we first describe our near-term R&D plan and list the resources that will be required. Our concluding section provides a very rough estimate of the construction cost and possible timeline.

5.1 Near Term R&D Plan and Request for Support

Our near term R&D will focus on producing a formal proposal over the next year. Much of this work will be done by the members of the collaboration, as the research fraction of their time, but some work will require some level of formal support by the lab and/or funding agencies. We focus here on a one year plan to work toward a proposal, emphasizing design and simulation. Following the proposal, our efforts will ramp up to include significant prototype work.

5.1.1 Accelerator Design Work

Work needs to be done on the details of the required accelerator modification, including transport from the Recycler, beam rebunching, slow extraction, and the experimental proton beam line. We believe the collaboration can perform most of this work, with the exception of the proton beam line to the experiment. We would therefore like to request a fraction of an FTE, probably from the External Beams Department, to help in the beam line design.

Accelerator design work also includes the plan for beam extinction. Fermilab is currently working with the University of Osaka to develop the AC dipole scheme and to investigate techniques for measuring extinction. This work is currently supported by approximately \$50K from the “US-Japan” fund, and a similar amount from the Fermilab Muon Collider Department. This effort should produce a workable design for the proposal.

5.1.2 Civil Construction

Although the Mu2e experiment will require a significant operational reconfiguration of the accelerator complex, it has been designed in such a way as to require no civil construction up to the point where the beam is extracted from the Debuncher.

Unfortunately, there is no existing enclosure suitable for the experiment, so a new proton beam line and experimental hall will need to be built, as discussed in Section 3. FESS has estimated that they will require about \$100K to work up a preliminary design for the building.

An important consideration in designing the experimental hall is whether to design it specifically for this experiment, or to make it a more general purpose facility for other potential users of this beam. We will work with the Accelerator Physics Center to come up with a recommendation in this matter.

5.1.3 Radiation Safety

One cannot overemphasize the importance of addressing radiation safety issues from the beginning of this project. For the envisioned beam intensities, the overburden on the Accumulator/Debuncher enclosure is far short of the requirement for passive shielding, even if we reclassify the entire area as a Radiation Area. We will therefore have to rely on an active monitoring system similar to the one used in the Booster. Establishing the effectiveness of such a system requires a great deal of work. We would like to enlist a fraction of an FTE from the ES&H department to aid us in working up a plan to produce the requisite Shielding Assessment Document (SAD) for the Mu2e experiment.

5.1.4 Muon Transport Line

Because of time constraints, we have restricted ourselves to the original MECO design for the muon transport channel for this letter. However, it would be a mistake not to consider advances in the understanding of muon transport that have occurred over the last decade, thanks largely to the work of those interested in cooling of muon beams, either for neutrino factories or muon colliders. For the proposal we will investigate the feasibility of more aggressive designs for the transport channel which might result in a significantly higher muon yield. This work is being largely undertaken by Muons, Inc, a collaborator in this experiment.

The magnet design will be extremely complex, so we would like to take advantage of the Technical Division expertise in superconducting magnet technology to aid us in producing a working design. A CDR for the muon beam line has been produced by MIT for the MECO experiment. Ideally experts from the Technical Division can begin to take a role in the further development of this design.

5.1.5 Personnel Needs

In addition to the specific R&D needs outlined, we hope to get support for some additional personnel. In particular, we hope that Fermilab could provide support for two Postdocs, with at least one of them dedicated to modeling efforts. There are several areas which need to be studied.

1) There is the possibility that intensity upgrades will occur at Fermilab. In that case we need to investigate the ultimate rate handling capabilities of the proton target and the detector system, and to evaluate new ideas to handle higher rates which have appeared since the MECO experiment was proposed. For example, how much could backgrounds be reduced, and rate capabilities increased, if the detectors were displaced further downstream, or at the end of a toroid, relative to the muon stopping target?

2) We need to evaluate the usefulness of locating the detectors underground in order to reduce cosmic ray background. In addition, there are issues of neutron background from the stopping target which cause background in the active cosmic shield which need to be addressed.

3) We need to work with the muon beam line team to evaluate new beam designs for backgrounds and number of stopped muons.

4) We need to develop an effective extinction monitor soon in order to evaluate in the near-term new schemes to achieve a 10^{-9} extinction factor, or perhaps even better in case future intensity upgrades allow us to reach an even smaller statistical error limit on the measurement. One idea promoted by the MECO experiment is to install a port which views the proton target. A magnetic spectrometer would select ~ 2 GeV scattered protons and detect them in scintillators.

5.2 Rough Estimate of Construction Cost and Timeline

While the final Mu2e proposal may differ from the MECO baseline, the cost and schedule developed for MECO are our best present estimates for Mu2e. The MECO cost and schedule was reviewed many times. The final and most definitive review was performed by the “Wojcicki Panel” [18]. The cost of the experimental apparatus was estimated to be \$27M in 2005 dollars, including an average contingency of 24%. Separately, the MECO solenoids and cryogenics were estimated to be \$58M including contingency. A cost estimate for the proton beam line and experiment hall will be developed for the proposal, but it is expected to be small on the scale of the experiment itself.

The critical path for MECO construction was defined by the Solenoid System and its associated cryogenics. A comprehensive conceptual design for these had been prepared by the MIT Plasma Science Fusion Center. The plan was to solicit an engineering design and procurement from a commercial vendor or vendors. This process could have begun as soon as funding was in hand and would end with solenoid installation in about 5 years. We believe the schedule will be driven by the solenoid construction; the detector system will take less time. If staged properly, we believe the experiment could be ready to take data within 5 years from the start of significant funding.

Bibliography

- 1 G. Signorelli, *J. Phys. G* 29, 2027 (2003).
- 2 Y. Kuno and Y. Okada, *Rev. Mod. Phys.* 73, 151 (2001).
- 3 R. Barbieri and L. J. Hall, *Phys. Lett. B* 338, 212 (1994).
- 4 R. Barbieri, L. J. Hall and A. Strumia, *Nucl. Phys. B* 445, 219 (1995).
- 5 J. Hisano, T. Moroi, K. Tobe, M. Yamaguchi and T. Yanagida, *Phys. Lett. B* 357, 579 (1995).
- 6 J. Hisano, T. Moroi, K. Tobe and M. Yamaguchi, *Phys. Lett. B* 391, 341 (1997). [Erratum-ibid. B397, 357 (1997).]
- 7 C. Dohmen, *et al.*, [SINDRUM II Collaboration.], *Phys. Lett. B* 317, 631 (1993).
- 8 M.L. Brooks, *et al.*, [MEGA collaboration], *PRL* 83, 1521.
- 9 G.W. Bennett, *et al.*, [Muon g-2 Collaboration], *Phys. Rev. D* 73, 072003 (2006).
- 10 J.P. Miller, *et al.*, *Rep. Prog. Phys.* 70, 795 (2007).
- 11 M. Raidal and A. Santamaria, *Phys. Lett. B* 421, 250 (1998).
- 12 M. Blanke, A. J. Buras, B. Duling, A. Poschenrieder and C. Tarantino, *JHEP* 0705, 013 (2007).
- 13 K. Agashe, A. E. Blechman and F. Petriello, *Phys. Rev. D* 74, 053011 (2006).
- 14 E. Gabrielli, *Phys. Rev. D* 62, 055009 (2000).
- 15 J. Bernabeu, E. Nardi and D. Tommasini, *Nucl. Phys. B* 409, 69 (1993).
- 16 G. Barenboim and M. Raidal, *Nucl. Phys. B* 484, 63 (1997).
- 17 V. Cirigliano, Kurylov, M. J. Ramsey-Musolf and P. Vogel, *Phys. Rev. Lett.* 93, 231802 (2004).
- 18 RSVP Review Panel (“Wojcicki Panel”) Report,
http://mu2e.fnal.gov/selected-documentation/rsvpbseline_reportfin.doc.
- 19 Fermilab Steering Group Report, p 17, September 18, 2007,
http://www.fnal.gov/directorate/Longrange/Steering_Public/report/Steering-Group-Report-2007-09-18.pdf.
- 20 R.M. Djilkibaev and V.M. Lobashev, *Sov. J. Nucl. Phys.* 49(2), (1989) 384; V.S. Abdjey, et al, INR Preprint 786/92 (1992).
- 21 MECO Collaboration, “DRAFT MECO Technical Proposal” (2001),
http://www.bnl.gov/npp/mu-e_docs/Draft_MECO_Tech_Prop.pdf
- 22 C. Dohmen, *et al* [SINDRUM II Collaboration], *Phys.Lett.B317:631-636, 1993*. Note that the published limit on the conversion branching ratio is 4.3×10^{-12} , and this is the number quoted by PDG. The best limit

from SINDRUM II, 6.1×10^{-13} , was presented as a preliminary result in a talk, http://www.triumf.ca/ichep98/talks_6/talk601.pdf.

23 Fermilab Proton Plan, http://www-accel-proj.fnal.gov/Proton_Plan/index.shtml

24 NOvA Experiment, <http://www-nova.fnal.gov/>

25 Accelerator and NuMI Upgrades (ANU) for the NOvA Experiment. http://tdserver1.fnal.gov/AcceleratorSupport/NOvA_ANU/

26 C. Ankenbrandt, *et al*, Fermilab BEAMS-DOC-2678 (2007).

27 D. McGinnis, Fermilab BEAMS-DOC-1782 (2005). (*Note: this scheme involved a new transfer line not required by our proposal.*)

28 D. Neuffer, BEAMS-DOC-2787 (2007).

29 N.N. Alexeev, S.L. Bereznitsky, and A.E. Bolshakov, Proceedings of the 1997 PAC, Vol 1., pp 261-263 (1997).

30 MECO Collaboration, “DRAFT MECO Technical Proposal”, pp 25-27 (2001), http://www.bnl.gov/npp/mu-e_docs/Draft_MECO_Tech_Prop.pdf

31 See, for example, K.Bishofberger, Ph.D. Thesis, UCLA (2005).

32 D. Bogert, Fermilab, private communication.

33 R.J. Ducar and J.S. Reid, Fermilab BEAMS-DOC-2883-v1 (2007).

34 D. McGinnis, Fermilab BEAMS-DOC-1783 (2005).

35 These are captured by the “Accelerator Neutrino Upgrade” (ANU) subproject within the NOvA proposal, <http://www-nova.fnal.gov/>

36 The report of this steering group, found at http://www.fnal.gov/directorate/Longrange/Steering_Public/report/Steering-Group-Report-2007-09-18.pdf, is useful for framing the discussion in this section.

37 See detailed discussion in D. McGinnis, BEAMS-DOC-1782-v7 (2005).

38 McGinnis, *et al*, “Accelerator Issues of Project X”, http://www.fnal.gov/directorate/Fermilab_AAC/AAC_July_07/ProjectX.pdf

39 *ibid.*, Section 6.2

40 *ibid.*, Section 6.3 and C. Ankenbrandt, *et al*, Fermilab BEAMS-DOC-2812-v1

41 Magnet System CDR, http://mu2e.fnal.gov/selected-documentation/mit_meco_magnet_cdr.pdf

42 “Production Target Reference Design”, http://mu2e.fnal.gov/loi/meco_documents/meco-tgt-03-001v1-01.pdf

43 “Heat and Radiation Shield Reference Design”, http://mu2e.fnal.gov/loi/meco_documents/meco-tgt-02-001v1-07.pdf

44 “Collimators and Transport Solenoid Shielding Reference Design”,
http://mu2e.fnal.gov/loi/meco_documents/meco-mub-03-002v1-01.pdf

45 “Anti-Proton Stopping Window Reference Design”,
http://mu2e.fnal.gov/loi/meco_documents/meco-mub-03-001v1-00.pdf

46 “Detector Shields and Absorbers Reference Design”,
http://mu2e.fnal.gov/loi/meco_documents/meco-mub-03-005v1-00.pdf