8-2016

A study of the radiative pion capture process as a background to the search for muon to electron conversion with the mu2e experiment.

Jacob Colston
University of Louisville

A STUDY OF THE RADIATIVE PION CAPTURE PROCESS AS A BACKGROUND TO THE SEARCH FOR MUON TO ELECTRON CONVERSION WITH THE MU2E EXPERIMENT

By

Jacob Colston
B.S. Physics, University of Louisville, 2014

A Thesis
Submitted to the Faculty of the
College of Arts and Sciences of the University of Louisville
in Partial Fulfillment of the Requirements
for the Degree of

Master of Science
in Physics

Department of Physics and Astronomy
University of Louisville
Louisville, Kentucky

August 2016

Copyright 2016 by Jacob Colston

A STUDY OF THE RADIATIVE PION CAPTURE PROCESS AS A BACKGROUND TO THE SEARCH FOR MUON TO ELECTRON CONVERSION WITH THE MU2E EXPERIMENT

By

Jacob Colston
B.S. Physics, University of Louisville, 2014

A Thesis Approved On

April 20, 2016

by the following Thesis Committee:

_____
Thesis Director
David Brown

_____
Christopher Davis

_____
Jacob Wildstrom

_____
Serban Smadici

ii

ACKNOWLEDGEMENTS

First and foremost, credit must be paid to the Blabe her-very-self, the Lady Balbanash, Queen Babelnasty, who rules over all of Blabedom: Shannon Drown. The goonery of existence itself would be utterly unmanageable without her companionship, nevermind the writing of this thesis. She has become both the foundation from which I spring forth, and my comrade in reckless ruckus. The union of Blabedom and Goondom will be magnificent to behold, and our enemies shall tremble in fear. I owe her for my content and happiness, my stability and (in?)sanity.

My family forms one of the primary boundary conditions for the differential equation describing what it is to be Jake. Striking out into a field of study that none of my family has even been tangentially involved with has been hard at times, but they've been with me through all of it, and shown no signs of quitting. On the contrary: my parents, Dan and Denise Colston, the Makers of Goons, bore my burdens when I was lowest, sat through lectures over subjects they don't *necessarily* understand or care about (making sure to nobly lie about it when that was the case), chauffeured me back and forth to school and the lab, and even went with me to different states to help with apartment searches and grad-school errands. My brother, Jared, the Gooniest of Bros, has also helped with the first and the last (and a little with the third), and, more than that, been a roommate and best friend to me. My sister, Maggie, the Royal Goon-Court Sorceress, has become a friend and a mirror for me for parts of myself I desperately needed to know. I owe them collectively for who I am and the support they've given me.

Dr. David Norvil Brown is the individual holding the most responsibility for my coming to physics as both a career and a school of thought. I met him when taking my first

ABSTRACT

A STUDY OF THE RADIATIVE PION CAPTURE PROCESS AS A BACKGROUND TO THE SEARCH FOR MUON TO ELECTRON CONVERSION WITH THE MU2E EXPERIMENT

Jacob Colston

April 20, 2016

This thesis will introduce radiative pion capture (RPC), a process which can produce a fake signal in a search for the coherent conversion of a muon to an electron in the presence of a nucleus. There will be a brief introduction to standard model (SM) physics, as well as some more in-depth discussion of the relevant high energy physics at the Mu2e experiment. We will discuss charged lepton flavor violation (CLFV), as well as Supersymmetry, which predicts CLFV at higher intensities than the SM prediction. A description of the RPC process follows, including the external and internal conversions in pion captures, and these processes' contributions to background at Mu2e. We will conclude with an estimate of the background contributions as obtained with analysis of Monte Carlo simulations.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

1.1 Intro to High Energy Physics and the Standard Model

Humankind, from as early as anyone can tell, has always sought to know what all things are made of. A quick tour through human history is very illustrative of this, especially as we approach modern times and the Standard Model (SM). The so-called Four Elements of earth, wind, water, and fire are a classic example of man's effort to classify the fundamental building blocks of the universe. Even as far back as the Ancient Greeks, however, we see such advanced concepts as the "atom", which was hypothetically the smallest unit into which things could be divided. Skip through history two millenia, and we have Rutherford's atom, and the Periodic Table (Whyte, 1961). With Rutherford's discovery, "atom" became a misnomer, due to its composition of protons, neutrons, and electrons. The idea of there being basic building blocks of matter was far from dead, however. There was simply a tinier, more fundamental family of entities to explore and study.

A decade *before* Rutherford's discovery, Quantum Mechanics was born when Maxwell Planck successfully described the Black Body Radiation spectrum with his assumption of the intrinsic quantization of the energy of light (for which he later won the Nobel Prize in 1918) (The Nobel Foundation). With the early experiments of Quantum Mechanics, alongside Einstein's new theory of Relativity, the scientific community endured a substantial blow to its classical intuitions (Born, 1962). However, endure it did,

and, with these great bastions of modern physics, there was finally a basis for understanding this new family of possibly fundamental entities. After the establishing of the groundwork of modern physics, and decades of high energy physics (HEP) experiments, we find ourselves in the 1970's, when knowledge of sub-atomic particles and their interactions was compiled into the still-prevailing theory of particle physics, the SM. More than just a compilation of existing sub-atomic particle knowledge, the SM also predicted the existence of many other particles, and many details concerning the characteristics of these particles and their allowed interactions. While most of the predicted particles have been confirmed in the decades since (i.e. the top quark, the Higgs boson), certain *characteristics* of these particles have been counter to the SM prediction (such as neutrino oscillations), and some theorized particles (i.e. dark matter) aren't described by the SM at all. These facts have significantly driven the rise of alternate theories, one of which will be very briefly discussed in a later chapter.

The SM organizes matter, and even the interactions among its constituents, into fundamental particles: six *quarks* and six *leptons*, organized into three "families," comprising all matter, the families being the columns in Figure 1.1; three fundamental forces (gravity is not described by the SM), mediated by fundamental, "force-carrying" particles. In Quantum Mechanics, particles with half-integer spin (i.e. $s = 1/2, 3/2, 5/2,$ etc.) obey Fermi-Dirac statistics, and are named *fermions*. Similarly, particles with integer spin (i.e. $s = 0, 1, 2,$ etc.) obey Bose-Einstein statistics, and are named *bosons* (Griffiths, 2008). Further, fermions obey the Pauli Exclusion principle, which forbids any two of them from occupying the same quantum state. Meanwhile, bosons do not obey Pauli's Exclusion principle, and can freely occupy the same state. It turns out that *all* fundamental matter particles (and very nearly all stable composite matter particles) are fermions, while *all* force mediators are bosons. This separation of matter and the force mediators by the Pauli Exclusion principle is ultimately the explanation of our typical observations that
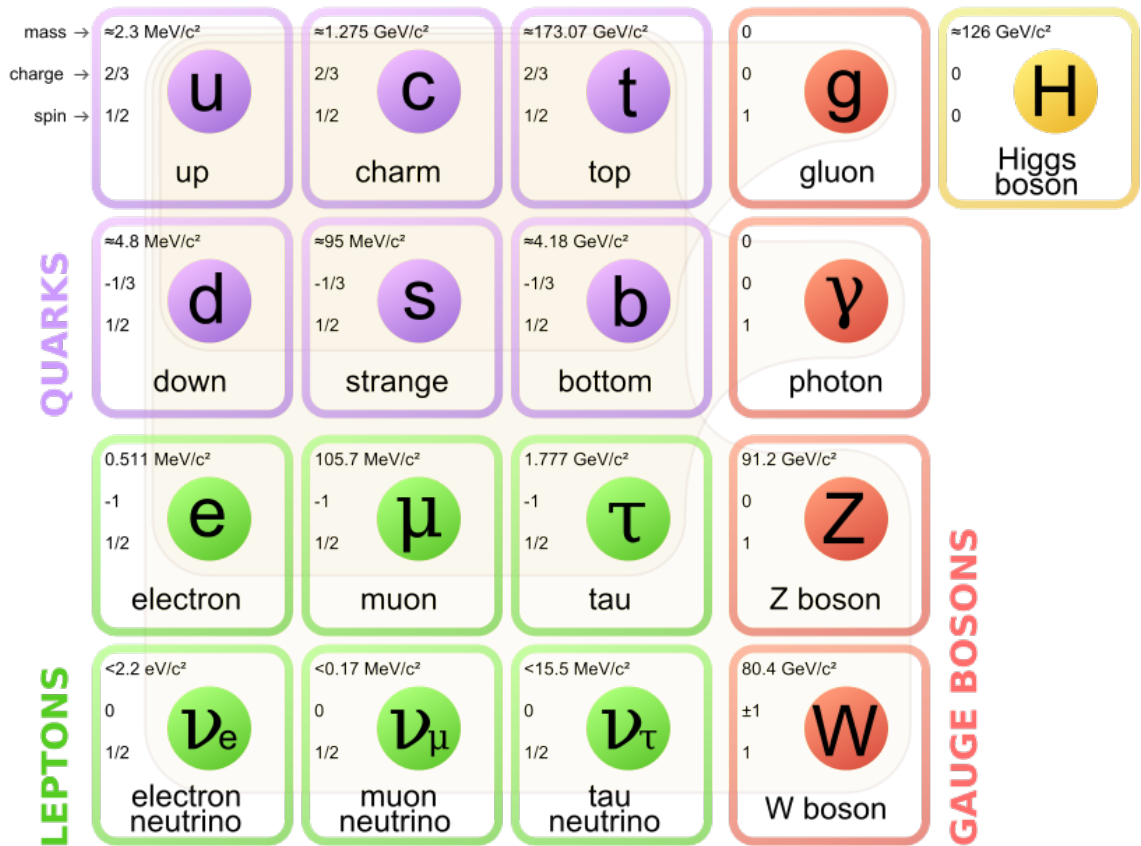
Figure 1.1: The Elementary Particles of the Standard Model. (Picture from `https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Par ticles.svg`

matter objects do not "pass through" one another, as Pauli's principle forbids fermions from occupying the same quantum state. Figure 1.1 displays the fundamental particles and their primary characteristics/quantum numbers.

The three fundamental forces in the SM are: the electromagnetic force, the strong force, and the weak force. These are mediated by *photons*, *gluons*, and the $W\pm$ and Z bosons, respectively. The familiar electromagnetic force occurs between any and all particles possessing electric charge. The strong force occurs between all particles possessing *color*, or *color charge*, which is a quantum property of quarks (and also the mediating gluons, unlike the photon and electric charge). The weak is most commonly

associated with flavor-changing decay processes, where *flavor* denotes the specific quarks and leptons in the interaction. Flavor will be discussed in a little more detail in the following section.

Although not actually a "force", a mention of the Higgs mechanism is appropriate here. The Higgs boson mediates the Higgs mechanism, analagous to how the photon mediates the electromagnetic force. The Higgs mechanism itself is the process by which all massive particles obtain their certain masses. Particles which do not have mass do not interact with the Higgs mechanism.

To finish our SM-basics discussion, particle-antiparticle annihilation must be mentioned. When a fundamental particle and its corresponding antiparticle collide, they can annihilate [1] in such a way that their mass energies are converted into "pure energy". This term is an unfortunately popular one, and the more formal meaning is that the mass energies of the two particles will rearrange into a "virtual" boson, which typically converts into a particle-antiparticle pair. This could be the same particle-antiparticle pair that went into the reaction, or any number of other particle-antiparticle pairs, so long as the quantum numbers before and after the reaction are conserved. Figure 1.2 shows a typical matter/anti-matter annihilation process.

1.2 Charged Lepton Flavor Violation

Despite being our current basis for understanding the sub-atomic world, there is no doubt in the scientific community that the SM is *not* the end of the story. One discovery counter to it, for example, was neutrino oscillations (Fukuda, Y. et. al., 1998). This is not just a flavor violating process, which is allowed but moderately suppressed in the SM, it is a flavor violating *conversion*, where only one particle is on either side of the interaction.

---

[1]Other annihilation processes can yield real bosons instead of virtual, but they require other constraints such as the presence of other matter, and are not as important to this discussion.
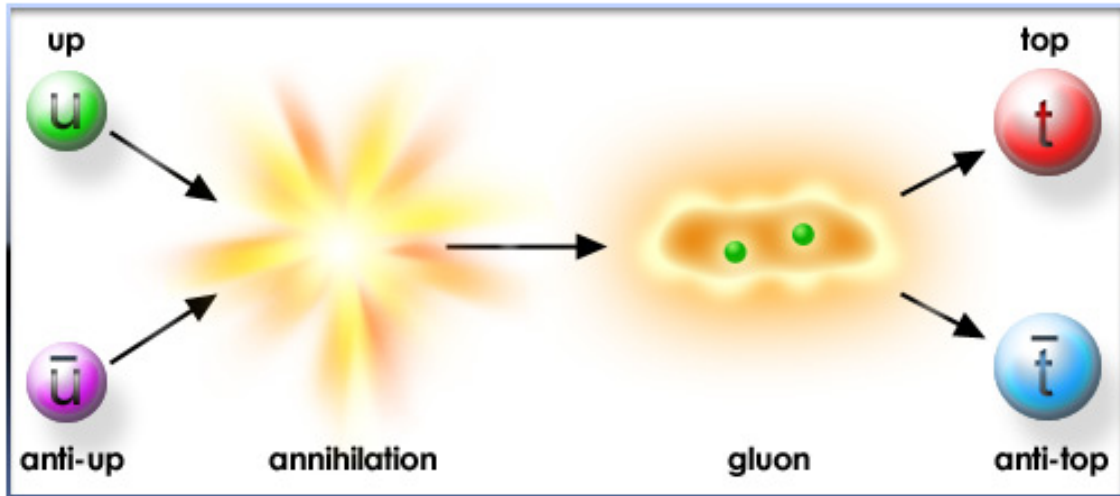
Figure 1.2: The above picture represents a matter/anti-matter annihilation process. (Picture from "The Particle Adventure" (Particle Data Group))

In typical HEP interactions, particle "flavor," which means the type of fundamental particle, is a conserved quantity. While quark flavor violation has been observed in weak interactions since the early days of HEP, *lepton* flavor violation was only first confirmed with neutrino oscillations by the Super-K experiment in Japan in 1998. Since neutrinos are leptons, this is "lepton flavor violation" with respect to the SM.

However, *charged* lepton flavor violation (CLFV) has yet to be observed, though not for lack of trying (Wintz, 1998). In many beyond the SM scenarios, like super-symmetry theories, CLFV occurs in some processes at rates which are far more reachable than the basically null rate we would expect (Ilakovac et al., 2013).

1.3 The Mu2e Collaboration

Mu2e is a planned flavor-physics/intensity-frontier experiment at the Fermi National Accelerator Laboratory, which will use the Main Injector. "Flavor-physics" is a term referring to particle "flavor", defined earlier. "Intensity-frontier" comes from a

convenient characterization of the current puzzling questions of physics, as viewed from a particle physics perspective, into three "frontiers". The "Energy Frontier" is explored by the highest of high energy HEP experiments, like those at CERN. It pushes the boundaries of the SM by searching for new particles which may fit into the theory (if they exist), but have not yet been reachable with the energies of previous HEP experiments. The "Cosmic Frontier", though gravity is technically not included in the SM, is concerned with the mysterious dissonance of gravity with the other, smaller-scale forces. Normally, it is more associated with astro-physical experiments. Finally, the "Intensity Frontier" deals with the detailed study of SM particles, discovering their mass relationships and quantum numbers, and are usually the more finely tuned, precise experiments, rather than the extremely high energy ones.

Mu2e specifically is aimed at measuring CLFV in the process of a muon directly converting to an electron with no accompanying neutrinos (hence the name), and with a higher precision than previous measurements (The Mu2e Collaboration, 2015). Even in alternate theories, these processes are very suppressed. Thus, the experiment's entire design is based upon high precision for this exclusive process, unlike previous intensity-frontier, multi-purpose experiments. The experimental design will be elaborated on in a later chapter.

At this point, the experiment is in its design stages, and initial construction has been going for just short of a year. The collaboration already consists of numerous universities and laboratories, both national and international (The Mu2e Collaboration, 2016). Until construction is complete, work is spent on simulating all the physical backgrounds necessary to understand our future data, as well as developing and optimizing the advanced equipment which will be paramount to this high-precision experiment. After construction, work will shift to different simulations to support the interpretation and management of data, as well as actual data management.

1.4   Radiative Pion Capture

The biggest *practical* implication of seeking to measure an ultra-rare process with high precision is that the collaboration will spend much more of its time studying processes that could *fake* the signal we're interested in than studying the signal itself. While there are many sources of background (fake signal) to the Mu2e process, the focus of this study is on one type: the radiative capture of pions by nuclei. This section will serve as a very brief overview of this process, a more thorough treatment of it will be presented at the end of Ch. 2.

Due to the nature of the muon-electron conversion, the electron will have a specific value of momentum (approximately $105 MeV$ in the lab frame). So, any process that can yield an electron of $\sim 105 MeV$ has the potential to fake our signal. Radiative pion capture (RPC) is such a process at the Mu2e experiment. This can occur when pions interact with the nuclei of a material. The "capture" is thus a nuclear one. This interaction is similar to the well-known nuclear phenomenon of electron-capture by a proton in the nucleus of an atom. Due to the pion's higher mass, the photon that is yielded in the RPC process can be of higher energy than the photon yielded in electron-capture.

The possibly high-energy photon emitted by an RPC process can, like all photons, go through a process called *pair-production*, where the photon spontaneously converts to an electron-positron pair. The electron in this pair can, in turn, have a momentum of $\sim 105 MeV$. Thus, when RPC occurs near the detector systems at Mu2e, there is a potential for background.

1.5   Thesis Overview

This section is a brief overview of this thesis:

- Ch. 1 Introduction - background information on SM basics, the Mu2e

Collaboration, CLFV, and RPC; also the thesis overview

- Ch. 2 Theory - background information on the theory of supersymmetry (SUSY), as well as the physics of the radiative pion capture process as a Mu2e background

- Ch. 3 The Mu2e Experiment - walkthrough of the experiment, with overviews of each detector system, and information on detector components

- Ch. 4 Software & Simulation - overview of Fermilab/Mu2e software and detailed walkthrough of the RPC background simulation and analysis techniques

- Ch. 5 Results - presentation of the results of the background study

CHAPTER 2

THEORY

2.1 Supersymmetry

This section is meant to give a *very* cursory treatment of SUSY and some of its different forms. While there are numerous books and papers giving more thorough developments of SUSY, this section is aimed at gleaning the basic information, emphasizing motivations for the theory and the testable implications which experiments (like Mu2e) have explored or plan to explore. Much of this section will be based on a very nice pedagogical paper on supersymmetry theory (Martin, 2016), with the rest coming from a phenomenological paper discussing processes like $\mu^- \rightarrow e^-$ (Ilakovac et al., 2013).

2.1.1 The Hierarchy Problem

To begin our brief discussion of SUSY, one might ask "Why was it ever proposed in the first place?" To answer this question, we must examine the common "problem" with the SM that is referred to as the "hierarchy problem". The hierarchy problem takes its name from the fact that, in the SM, gravity appears to be so much weaker than the other forces (recall that gravity is not even properly represented in the SM). In the *hierarchy* of fundamental forces, gravity stands very far apart, and yet, the SM has no explanation for *why* this is (or in other words, the SM did not predict that this necessarily *had* to be the case).

To be a little more explicit, the energy of the Higgs field[1] was an undetermined parameter in the SM before the measurement of the $W\pm$ and Z boson masses. The measurement of these values at ~ 100 GeV fixed the non-zero Higgs value (the trivial solution of 0 GeV was always possible) at ~ 250 GeV. One might be thinking "So, what? What's wrong with ~ 250 GeV?" To answer this, let's take a cursory look at how gravity might fit into the SM framework that is already established.

From the implications of Quantum Mechanics and Quantum Field Theory, all properties must be quantized at some level, such that there is a smallest possible unit of length and a smallest possible unit of time. These are customarily called the "Planck length" and "Planck time", respectively, acknowledgements of Max Planck who is the grandfather of quantum mechanics. Now, maybe a little counter intuitively, the Planck *mass* is not defined to be the smallest possible mass that can exist, but instead defined to be the mass of the smallest possible black hole, where the black hole's Schwarzschild radius is set to the Planck length. The Planck mass, then, in this characterization, helps to roughly represent gravity's place in the SM theory. Using this definition, the Planck mass is $\sim 1,000,000,000,000,000,000$ GeV!

Ok, so still, where is this "problem"? The problem lies in the fact that, due to the SM theory, the stable Higgs field energy solutions should only be the trivial solution of 0, or the Planck mass energy, if gravity is even close to properly represented in the theory. However, due to the measurement of the $W\pm$ and Z boson masses, the non-zero solution of the Higgs field energy is definitely ~ 250 GeV! This is incredibly small compared to the Planck mass, but decidedly not the trivial solution of 0 either. The expression of this huge dissonance in the grand scheme of things is that gravity is not just the weakest force; it is shockingly, disturbingly weak in the hierarchy of fundamental forces. Thus, we have

---

[1]It's important to note here that the hierarchy problem is often characterized as a problem with the Higgs *boson* mass value, but this is not quite correct. The problem is with the Higgs *field* energy value, which was fixed in the SM after the measurement of the weak force bosons.

the hierarchy problem.

Now, the way that the Higgs field energy value is calculated in Quantum Field Theory is with the standard quantum corrections which come simply from the existence of the particles in the SM. These corrections lead us to the belief that only the Planck mass energy and the trivial solutions are stable solutions. However, it was noticed at some point that, if there were partners to every particle in the SM, by some natural symmetry between fermions and bosons, these quantum corrections cancel *exactly* to yield a Higgs field energy value close to the one we've observed. It is important to note that, it doesn't simply become *possible* to cancel the quantum corrections and achieve something close to our measured ~ 250 GeV, but completely *necessary*, once one has supposed that these "super-partners" exist. This is the basis of SUSY, and how it addresses the hierarchy problem.

2.1.2   Minimal Supersymmetric Standard Model

The simplest form of SUSY is the Minimal Supersymmetric Standard Model (MSSM). From its name, one can guess that it is only a "slight" extension of the basic SM theory, acknowledging that the SM is quite powerful in its own right. An interesting point about SUSY in general is that it must be a *broken* symmetry if it is physically realistic. If it were rigid, the supersymmetric partners would have identical masses to the basic SM particles they're partnered to, and the SUSY partners would have been found by experiments long ago. In the MSSM, there are two primary features that make it "minimal" in its supersymmetric nature: (1) there is only *one* family of supersymmetric partners to the SM particles[2], and (2) the supersymmetry is only "softly" broken.

The first approach to testing a theory like the MSSM is to try to measure the

---

[2]Other more exotic extensions of SUSY postulate more families of super partners, and so more than one "supersymmetry". These are not yet phenomenologically valuable, however, as all of their implications are quite untestable with modern science.

existence of those particles proposed to be super partners. However, the Energy Frontier sector has not found any unambiguous signs of new physics in this regard, suggesting that, if they still exist, the super partners are still outside the energy scales available at experiments. But this is far from a dirge for SUSY, or even MSSM. Recall the effect that the simple *existence* of SUSY particles had on the calculation of the Higgs field energy.

This phenomenon is rampant in MSSM, where the simple existence of the SUSY particles alters the the characteristics and behaviors of the normal SM particles, either directly or indirectly, and in varying degrees. As mentioned briefly in Ch. 1, MSSM predicts that the muon, already thought of as a "cousin" to the electron, is actually a much "closer" cousin than in the normal SM. This tighter coupling of the muon to the electron implies that the coherent $\mu^- \rightarrow e^-$ process is much less rare than in the SM case. There are a multitude of other processes (Ilakovac et al., 2013) which, if observed with the currently available sensitivities, would be unambiguous signs of physics beyond the SM, and would heavily support the MSSM theory.

2.2   Radiative Pion Capture (and Internal Conversion)

The process associated with this background study is the radiative, *nuclear* capture of a π^- by an Aluminum nucleus. The naive explanation is in the quark content of the particles involved: protons in the nucleus are composed of two up quarks and a down quark, while the π^- is composed of an anti-up quark and a down quark. When the π^- is captured, one of the up quarks of one of the protons in the nucleus annihilates with the anti-up from the π^-. The remaining quarks, one up and two downs, form a neutron. The annihilation process yields a photon. Quantum mechanically, there are two independent processes allowable for the annihilation: one which yields a real photon, and one which is

12

"virtual"[3]. The virtual particle, as such, almost instantaneously (exactly instantaneously, as far as our sensitivities are concerned) converts into an e^+e^- pair, while the real photon can propagate away from the capture point, and *may* pair-produce elsewhere, in an interaction with other material.

The real-photon RPC process is referred to as an *external* conversion, and the virtual-photon RPC process is referred to as *internal*. The external process goes like

$$\pi^- + (A, Z) \rightarrow \gamma + (A, Z-1)$$
$$\hookrightarrow e^+e^-$$
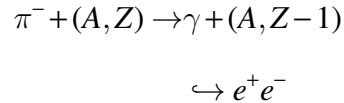
whereas the internal conversion process goes like

$$\pi^- + (A, Z) \rightarrow e^+e^- + (A, Z-1)$$
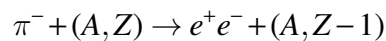
It is worth noting here that the RPC process has been studied before with a Magnesium target (Bistirlich et al., 1972), and the energy spectrum of the photons followed the Bistirlich distribution. This same spectrum is assumed to be correct for the Aluminum target as well.

---

[3]The virtual particle concept is an artifact of quantum field theory. If all interactions between matter particles are to be represented as the exchange of *other* particles, there becomes a need to define "virtual" particles which mediate those interactions. It would not serve this background study to delve any deeper than that. Suffice it to say that the *effects* of virtual particles are quite observable.

CHAPTER 3

THE MU2E EXPERIMENT

3.1 Outline

Due to conservation of momentum, the direct $\mu^- \to e^-$ conversion process, if physically realistic, could only occur in the presence of an atomic nucleus. Experimentally, this means that muons must be atomically captured in a stopping target (ST) of a certain material (currently planned to be Aluminum), and then monitored for the conversion signal. This signal is an electron of a very well-defined, specific value of momentum ($\sim$ 105 MeV/c). The mono-chromatic nature of the signal-electron's momentum is completely due to the fact that in a conversion process, the single daughter must have a kinetic energy exactly equivalent to the difference in the mother and daughter's mass energies (the excess energy has no where else to go). Were it not for this mono-chromatic signal-electron, there would be no hope to observe this ultra-rare process with modern technology.

The way Mu2e produces a beam of muons, to stop upon a ST, is by colliding protons with yet *another* target, called the production target (PT), in such a way as to produce prodigious amounts of pions (The Mu2e Collaboration, 2015). These pions are "swept" down a beam pipe by solenoidal magnetic fields and allowed to decay during their flight towards the stopping target. Pion decays *very* often include muon daughters, and so we achieve a beam of muons which collide with the stopping target. Finally, after a significant number of muons have been stopped (and counted), we watch for the

conversion signal with the detector system.

3.2   Where does Radiative Pion Capture Fit?

Ideally, *all* the pions would decay in flight to muons, and *all* the muons would be stopped by the ST. Realistically, most of the muons are *not* stopped, and some of the pions don't even decay before getting all the way down the beam pipe to the ST. These pions that reach the ST can have an interaction with the ST nuclei that can in turn yield a high energy photon. This high energy photon can then go through a process called *pair-production*, where the photon spontaneously converts to an electron-positron pair. The electron from this pair can have a momentum which is near the signal value, thus making the RPC process a possible background to Mu2e.

3.3   Protons to Stopped Muons

This section describes the parts of the accelerator complex at Fermilab that are relevant to the Mu2e experiment. Figure 3.1 is a conceptual diagram of part of the complex that has yet to be decommissioned (like the TeVatron). Much of the diagram pertains to Mu2e and the other muon experiment, Muon g-2, and will serve as a good visual reference for the remainder of this section. All figures in this section by courtesy of Mu2e (The Mu2e Collaboration, 2015).

3.3.1   Proton Pulse

The Mu2e experiment begins with a pulsed, ~ 8 GeV beam of protons. The relatively low energy proton beam is for the purpose of achieving, eventually, a relatively low energy *muon* beam, which optimizes the yield of stopped muons in the ST. The *pulse* is a clever design for background prevention: by pulsing the beam, the entire experiment
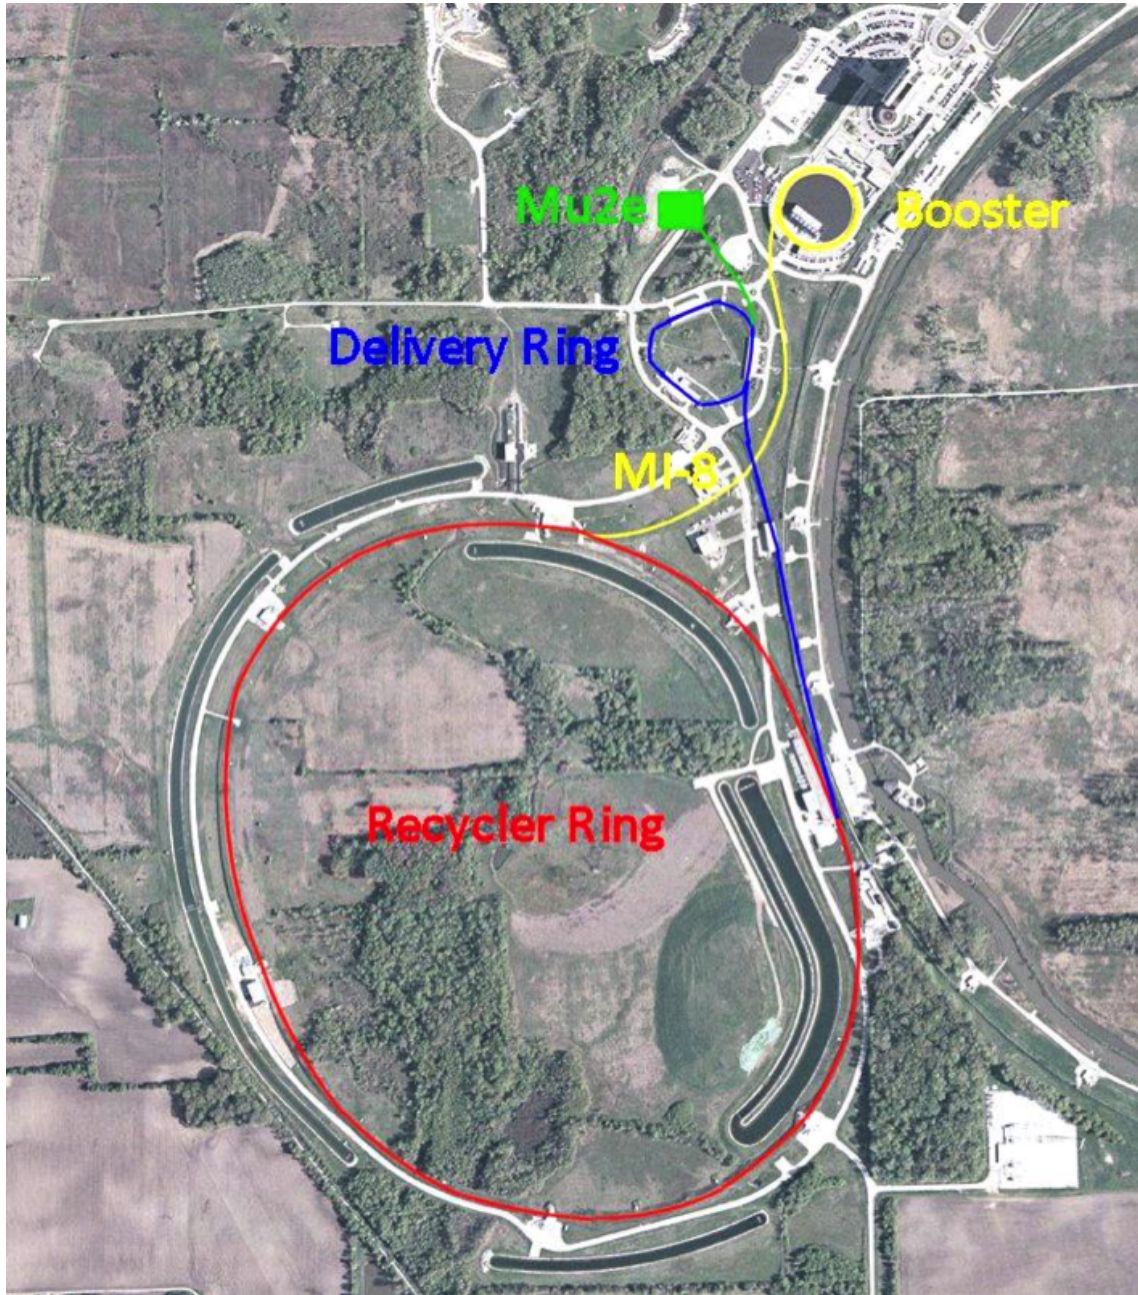
Figure 3.1: The above figure shows the Proton Beamline layout. Note the Main Injector is below the Recycler (into the page, underground).
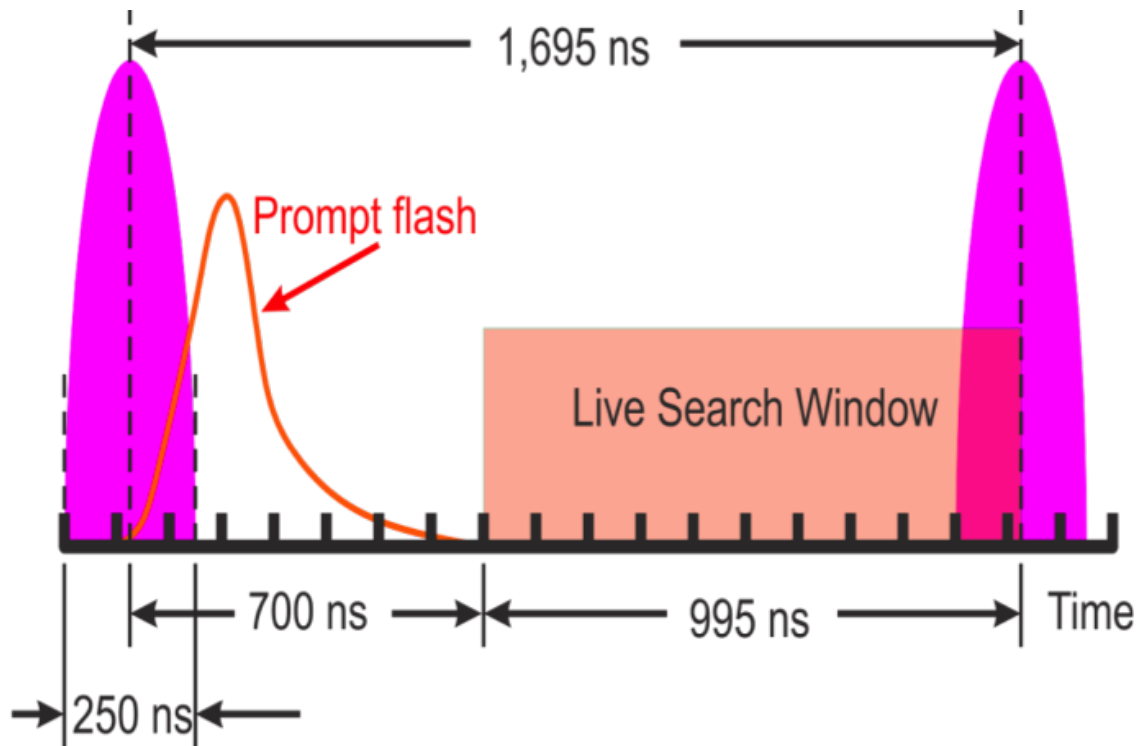
Figure 3.2: This figure shows the pulse shape for protons on the production target. Also shown is the time-distribution of primary backgrounds, referred to here as "prompt flash."

gains a swelling, Gaussian time-distribution, shown by Figure 3.2. From the figure, it is fairly obvious that there is, with this time-distribution, an optimal window of time for all of the detector components to "open their eyes" and begin collection. By restricting detector component collection to *only* that optimal window, the collected data is optimized for as many conversion electrons as possible relative to the amount of background collected.

The proton beam begins at Fermilab's Ion Source. The protons receive their initial kick from the Linac, and then are fed immediately to the Booster Ring. The Booster is almost entirely responsible for the beam's kinetic energy, and feeds two batches of ~ 8 GeV protons to the Recycler Ring. The Recycler's primary purpose is to divide the batches of protons which enter it into well-defined bunches. The Recycler uses a

radio-frequency sequence to coalesce each proton batch into four bunches. These are transferred to the Main Injector below the Recycler Ring, which then synchronously transfers the bunches to the Delivery Ring. Finally, the Delivery Ring sends the bunches, containing $\sim 3 \times 10^7$ protons, to the PT.

3.3.2 Production Target & Solenoid

The entire Mu2e experiment is surrounded by a system of three solenoids (see Figure 3.3) which serve the purpose of transporting an optimal amount of muons (of appropriate momenta) to the ST and reducing the amount of other particles transported. The first of these solenoids is the Production Solenoid (PS), and it surrounds the PT area. It ranges in strength from 4.6 T at the upstream end to 2.5 T at the downstream end. This field gradient is the first significant filter of background after the proton beam pulse, as it tends to dump positively charged particles away from the Transport Solenoid (TS) and route negatively charged particles towards the TS.

Most important for Mu2e amongst the particles created when the proton pulse hits the PT are negative pions. These are swept towards the TS where they decay in flight to muons. The PT's geometry and the geometry of the PS area are optimized to reduce reabsorption of the pions by the PT and maximize the amount of pions transmitted to the TS.

The extinction monitor is also located in the PS area. This monitors "beam extinction," referring to the beam pulses coming into the PS area from the Delivery Ring. The beam pulses, or bunches of protons, are not perfectly tuned and there can be protons which arrive at the PT outside of the pulse time-window. Any of the protons-on-target (POTs) which are outside of the pulse window are referred to as out-of-time (OOT). The extinction monitor watches for when these OOT POTs are at a minimum, or in other words, when the "extinction" is at a maximum. This is the primary trigger for the other

detector components of the experiment.

3.3.3   Transport Solenoid & Collimator

The aforementioned TS holds the primary design feature for getting stopped muons in the ST. As can be seen from Figure 3.3, the TS's geometry is arranged in an S-shape. The magnetic field for this region is graded as well, albeit more gradually: 2.5 T at the upstream end, and 2.0 T at the downstream end. The field gradient in the TS helps to continue the primary function of the *PS*, but with less gusto. The *primary* function of the TS region is driven by it's S-shape curvature: charged particles with momenta above a certain threshold, which is determined by the TS curvature, will tend to crash into the walls by not being deflected quite enough in their trajectories to stay on track. Meanwhile, charged particles with momenta below a certain threshold will tend to crash into the walls by being deflected *too much* to stay on track.

Further, the S-shape of the TS helps to eliminate backgrounds from neutral particles created in the PT collision. Neutral particles are not deflected in the solenoid fields, and so have straight trajectories. The S-shape 90° bends, combined with a series of absorbers, prevent any of these "line-of-sight" trajectories from reaching the DS area downstream.

The primary collimator located in the center of the central straight section of the TS serves the important function of further filtering the beam based on sign-of-charge and momentum. Particles will tend to separate very nicely based on charge in the 90° bend immediately before the collimator. This allows us to block the postively charged particles preferentially at the central collimator. Further, the beam pulse will gain a systematic momentum distribution in the bend: particles with higher momentum will tend to drop vertically lower, and lower momentum particles will tend to drift higher. Since the beam pulse is distributed smoothly based on momenta as it arrives at the central collimator, it is

also possible to select an optimal window of momenta which is transmitted through to the downstream TS and DS areas. This momentum selection is paramount to optimizing the stopped muon yield.

3.3.4   Detector Solenoid & Stopping Target

The DS begins with a magnetic field of 2.0 T, coming from the TS area, and sharply decreases to 1.0 T. The ST is situated toward the upstream end of this gradient. After the graded area, the field is kept at a uniform 1.0 T for the remainder of the DS area. This field gradient serves a primary purpose of directing conversion electrons toward the detector components, even when they are emitted in the upstream direction. The uniform field area is where nearly all of the detector subsystems reside. Signal electrons will follow well-defined, helical trajectories in this uniform magnetic field, which facilitates optimal design of the detector components to take advantage of this.

The ST, while originally planned to consist of 17 circular foils arranged in a conical shape, is now planned to consist of 34 foils with no conical taper to the foil radii (in other words, 34 circular foils of equal radius arranged in a cylindrical shape). Each foil is planned to be 75 mm in radius, 0.1 mm thick, and spaced 24.24 mm apart, center to center. The circular faces of the foils are perpendicular to the beamline. The geometry of the ST is designed to optimize stopped muon yield, as well as minimize reabsorption of signal electrons and photons by other foils.

3.4   The Detector System

This section, and Figure 3.3, focus more closely on the Mu2e experiment itself, and will serve as a good visual aid for the following detecor component subsections. All figures in this section by courtesy of Mu2e (The Mu2e Collaboration, 2015).
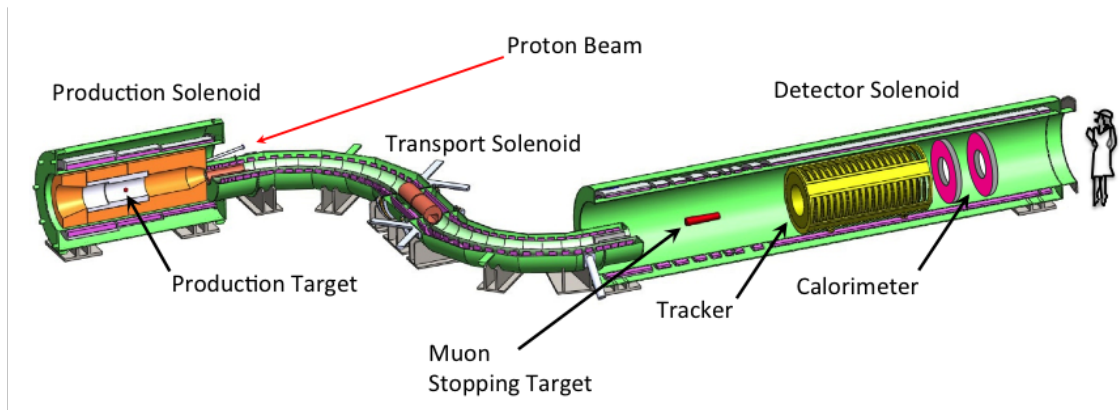
Figure 3.3: The above figure shows the Mu2e detector system. This system is to be housed in a building located at the block labeled "Mu2e" in Figure 3.1.

3.4.1   Stopping Target Monitor

The ST Monitor is a very straightforward detector component: its function is to count the number of muons stopped in the ST. One might think that the most straightforward way to do this would be to measure the photon spectrum which is produced in the atomic capture of the muons, thus measuring the amount of muons stopped in the ST. However, due to the timing considerations, this is a problematic approach, because, while the photon spectrum from the muon captures is rich, it is coincident with most of the other background processes coming from the "beam flash".

To aleviate the high rate and radiation problems of directly measuring the muon-capture-photon spectrum, an alternative approach was developed: detecting photons coming from the decays of radioactive nuclei which are produced in *nuclear* muon capture, and are delayed relative to the atomic muon capture photons[1]. Energy and intensity information on *this* photon spectrum is highly dependent on the ST material, and

---

[1]Nuclear capture refers to an interaction with a nucleus, whereas atomic capture is an interaction with an atom as a whole. In the latter case, one can often think of "muonic" atoms, for example, where the atomic capture results in a muon taking the place of an electron in an orbital. For nuclear captures, such models aren't readily available. For muons, atomic capture occurs much more frequently than nuclear.

available in the literature.

3.4.2 Tracker

The first detector component which the conversion electrons will interact with is the tracker. The term "tracker" has become a popular catch-all phrase in HEP, and for good reason. With very few exceptions, nearly all HEP experiments are driven by track reconstruction, at some level, and this is only possible with some sort of particle trajectory information, usually in a magnetic field. Most tracker devices measure the position of a charged particle (with minimal deflection) by measuring the ionization the charged particle leaves when passing through some medium. However, not all trackers achieve this with the same design. The Silicon Vertex Tracker (SVT) at the BaBar experiment, for example, is composed of silicon strips. Charged particles passing through the thin semiconductor material ionize nearby atoms, and the ions are accelerated to electrical readouts by an applied voltage. The ion drift times give timing information of the passing particle, while the locations of ionizations give corresponding position information.

The Mu2e tracker, shown in Figures 3.4 & 3.5, is composed of straw drift tubes, which are 5 mm thick, metalized Mylar® tubes, with a 25 μm thick sense wire concentrically inside. The tube walls are 15 μm thick, and the inner chamber around the sense wire is filled with Argon. As charged particles pass through the tubes, ionization occurs in the Argon gas, and the ions drift to the sense wire. Similarly to the previously described SVT, the straw ionized gives position information of the passing charged particle. These tubes are arranged in bunches (shown by the red and blue in Figure 3.5) and held in place by the straw assemblies (one of which is shown in Figure 3.5), which are connected to the electrical readouts. These assemblies are then connected concentrically to form the tracker as shown in Figure 3.4.

The key physical feature of the tracker is its radius. Since the signal electron from

22
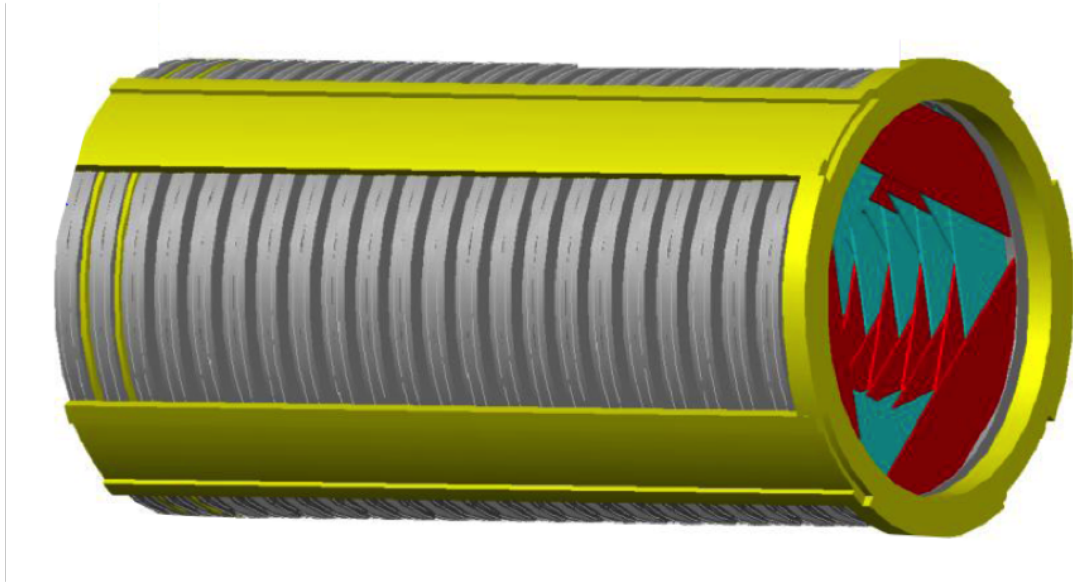
Figure 3.4: The above figure shows the tracker, with all of its straw tube assemblies arranged.

coherent $\mu^- \to e^-$ conversion has a single value of momentum, the radii of the helical

trajectories of signal electrons in the uniform magnetic field of the downstream DS area

will all be constrained. [2] There is still some ambiguity from the fact that not all signal

electrons will have the same initial positions, but the constraint on their trajectory radii

allows for an optimal tracker radius. This is a very effective method for removing

background coming from charged particles of too low momenta; they're never measured

in the first place!

---

[2] The radius of a charged particle's helical path in a magnetic field only depends on its *transverse* momentum, strictly speaking, and not the magnitude. However, so much of the background is of lower momenta than the signal, that even when most of those particles have nearly *all* their momenta in the transverse direction, it is not enough for their radii to reach the tracker (again, for most of them). In fact, even some *signal* tracks will be missed due to this design, but the background rejection advantage is so great that it still becomes optimal to apply this constraint to the tracker radius.

Figure 3.5: The above figure shows a single straw assembly. The blue and red regions represent the straw tube bunches. The central yellow circle represents the stopping target, and the other circles represent electron trajectories in the transverse plane. The green circle is a signal electron.

3.4.3 Calorimeter

The calorimeter serves the primary function of supporting the tracker. The Mu2e Online Software will perform track reconstruction on "hits" in the tracker, or measured charged-particle locations. However, the tracker will undergo high rate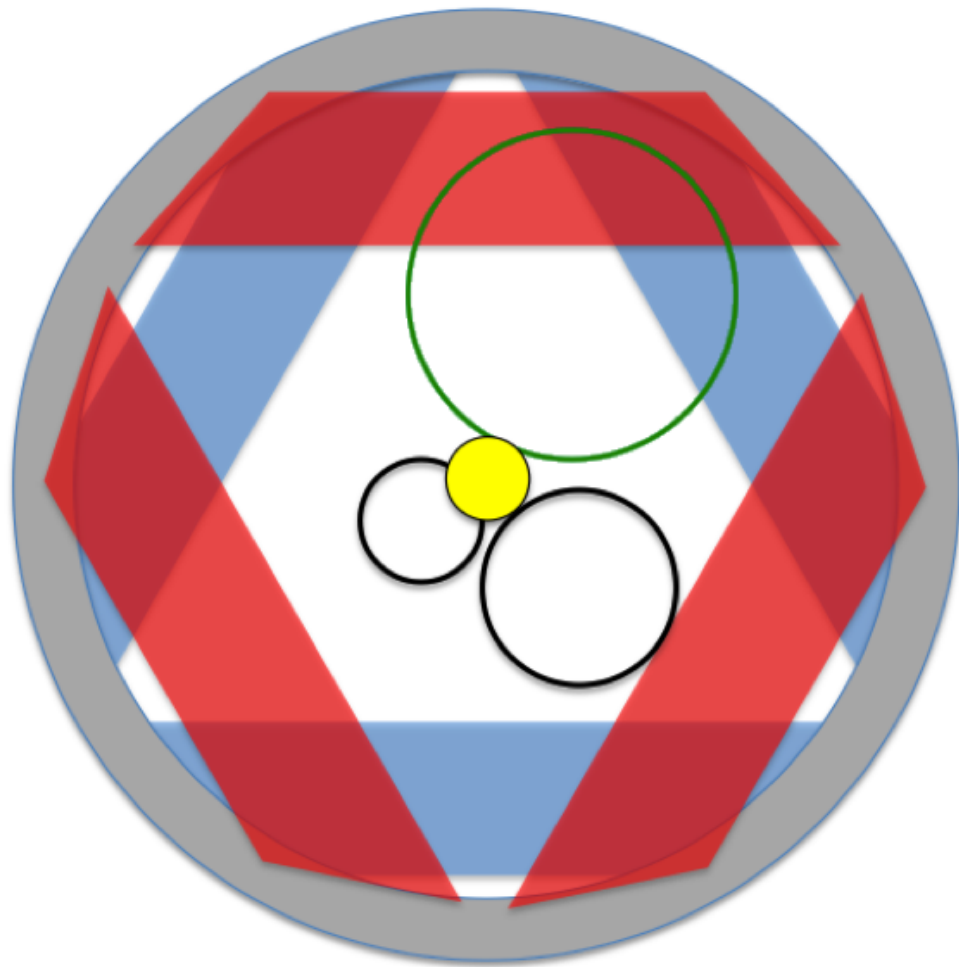s of charged-particle flux for much of the experiment's lifetime, with small timing separations for the hits. This facilitates the possibility of mis-reconstruction of tracks. When the tracks realistically have energies already close to the signal energy, this makes it very likely that a track will be mis-reconstructed to appear consistent with the $\mu^- \to e^-$ process. The calorimeter helps alleviate this issue by giving an independent measurement of charged-particle momenta and energy. Basically, reconstructed tracks are rejected as background when they cannot be sufficiently associated with a corresponding hit in the calorimeter.

The Mu2e calorimeter is of the "total absorption" variety. This means that, opposite the tracker, the calorimeter is designed to interact strongly with charged particles, completely stopping them and absorbing their energy. For this reason, the calorimeter is to be made of scintillating crystals, where the photon shower initiated by the charged particles is collected by avalanche photo-diodes (APD's) attached to each crystal. Initially, luttetium-yttrium oxyorthosilicate (LYSO) was considered for the scintillating crystal, but was found to be cost prohibitive. Among the several other materials being considered, barium fluoride (BaF_2) is currently performing the best in simulations.

As shown in Figure 3.6, the calorimeter is comprised of two annular disks, each of which holds an array of 1860 crystals. The yellow hexagons are the BaF_2 crystals, which are 33 mm per hexagonal side, and 200 mm in depth. On the downstream side of the disks (not shown in the figure), the APD's and electronics readouts are attached. As described before, charged particles will spiral about the azimuth as they travel towards the
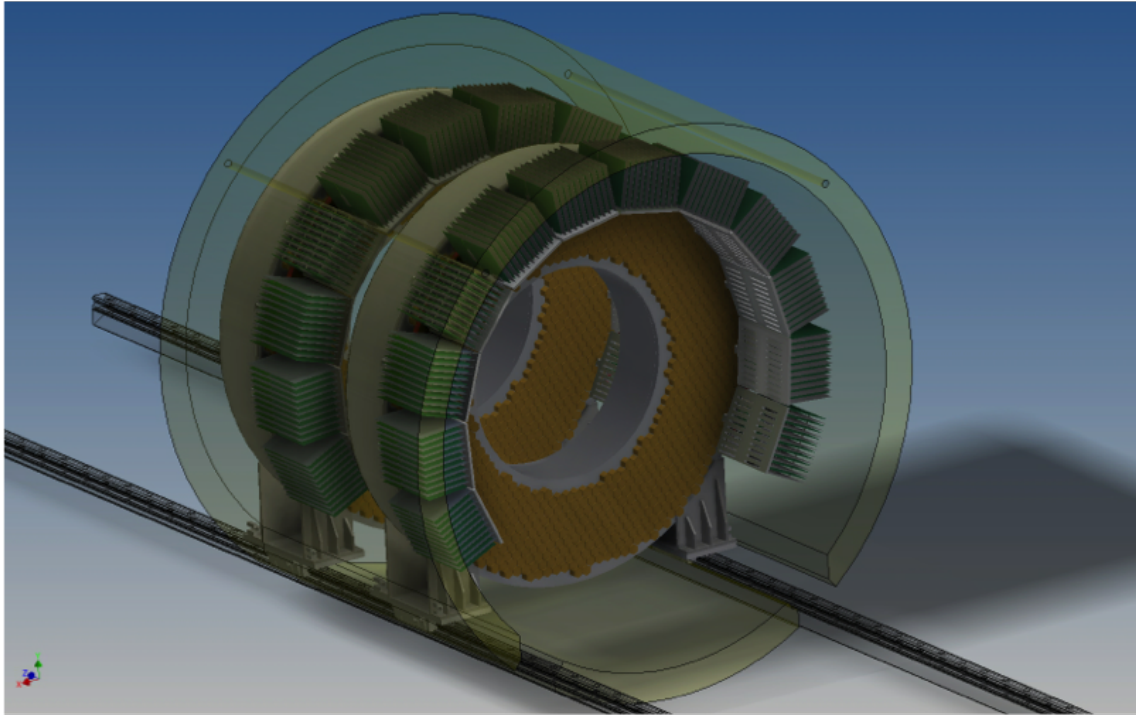
Figure 3.6: The Mu2e calorimeter.

calorimeter, causing photo-showers in the calorimeter crystals as they go (until they are stopped). The precise timing resolution of the APD's allows for very accurate correlation with reconstructed tracks, facilitating the calorimeter's primary function.

3.4.4   Cosmic Ray Veto

The cosmic ray veto (CRV) is a unique detector component, largely because it is designed to address a unique source of background: cosmic rays. Cosmic rays are particles coming from interactions elsewhere in the cosmos, outside Earth's atmosphere (or are at least direct by-products of *those* particles' interactions/decays in Earth's atmosphere). Since these particles can come from a wide variety of processes having a wide variety of initial conditions, they can also have a wide range of momenta and energy. Typically, cosmic rays which survive Earth's atmosphere all the way to particle detectors

are muons. Being charged, and possibly having energies near the signal energy, they present a significant possibility of creating charged-particle tracks in the tracker and corresponding hits in the calorimeter which will be consistent with the $\mu^- \to e^-$ signal.

The first line of defense to this background is actually not the CRV, but the various forms of passive shielding surrounding the experiment, such as the overburden above and to the sides of the detector hall and the concrete shielding around the DS area. However, cosmic rays can still be quite penetrating. A further step taken to reject cosmic ray backgrounds is by particle identification with the tracker and calorimeter (cosmic rays which obviously didn't come from the ST by their tracks can be rejected, for example). However, there is still the possibility that a cosmic ray will penetrate the passive shielding *and* initiate a particle with ~ 105 MeV/c which appears to emanate from the ST. This specificity may seem unlikely, but one must keep in mind the rarity of the $\mu^- \to e^-$ process being sought.

In order to reject cosmic ray processes further than with tracker/calorimeter particle identification, the CRV is arranged around the DS area, and even the downstream half of the TS, as in Figure 3.7. The enclosure is on the top and sides, and is composed of four layers of long scintillating bars. Silicon Photo-Multipliers (SiPM's) are attached to the ends of the bars, with wave-shifting fibers run through the bars to optimize the SiPM readouts. The bars are then attached in a staggered configuration, as shown in Figure 3.8 which shows one "module", with Aluminum absorber layers between them. The staggered configuration serves to prevent line-of-sight cosmic rays from penetrating undetected. When the CRV detects a cosmic ray of a certain energy, this is used to trigger a "veto" for the experiment, where the Online Software flags the data as having cosmic background.
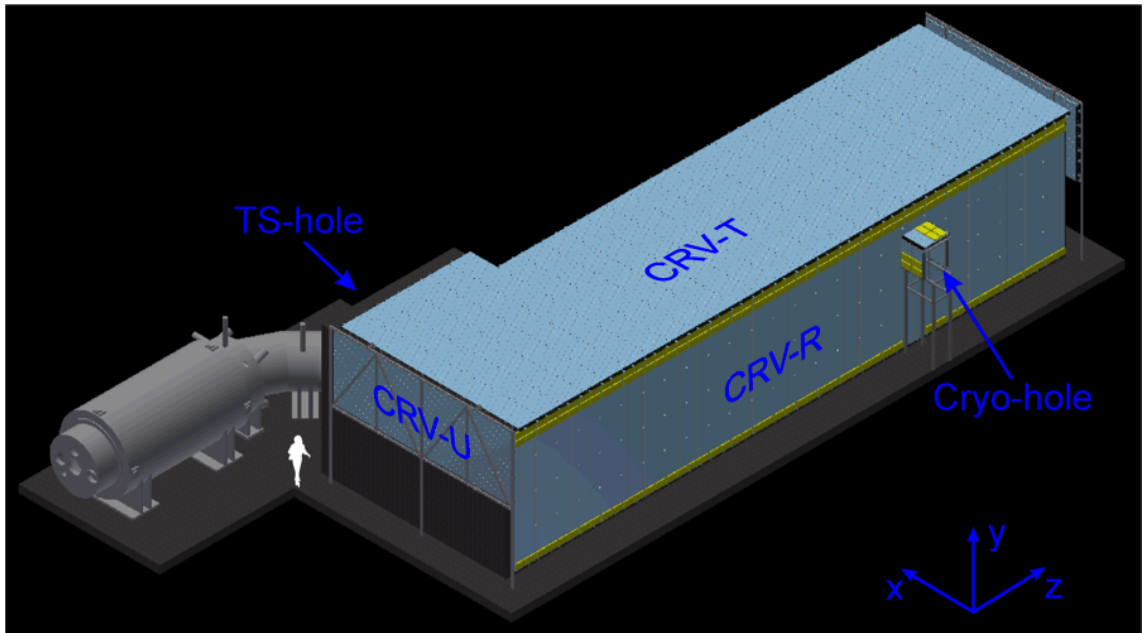
Figure 3.7: The above figure shows the full cosmic-ray-veto subsystem.
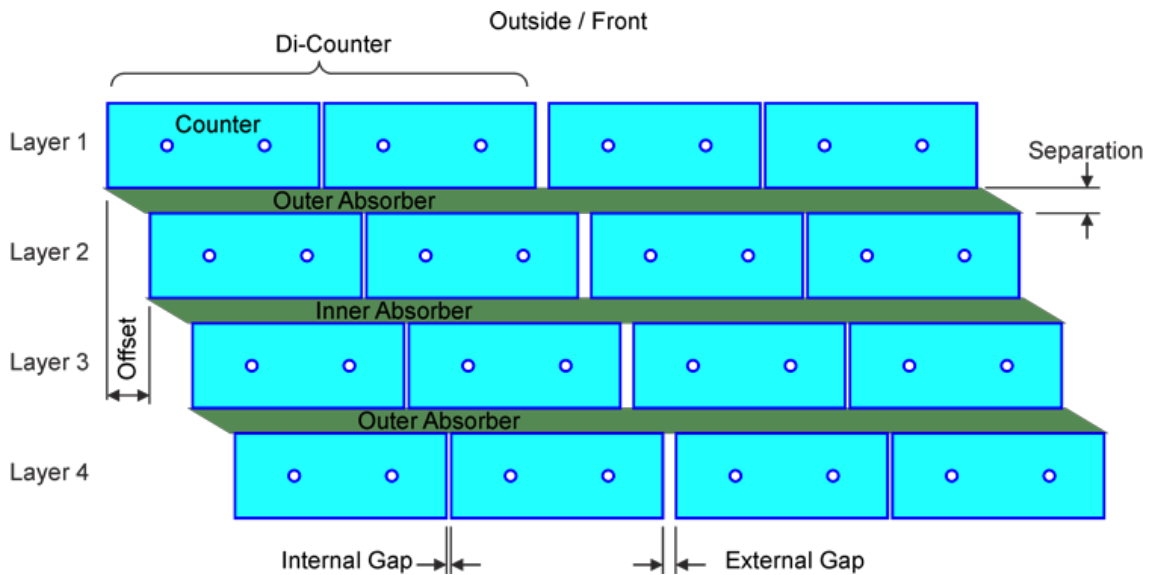
Figure 3.8: The above figure shows a downstream-facing cross-section of a single cosmic-ray-veto module.

CHAPTER 4

SOFTWARE & SIMULATION

4.1   Software Overview - Frameworks & Packages

In order to optimize and understand the experiment we mean to undertake, we simulate the processes involved, at as many levels as possible, including our detector components' defects. Further, HEP experiments have been undertaken for nearly a century and a lot of groundwork has been laid in the computer sciences realm as well as physics. One of the most important foundations laid (especially for this study) are the various software frameworks and packages which have been developed, often times exclusively, for HEP research.

At Mu2e, there are several broad areas of software:

- The "official" Offline software package, which is actually called "Offline". It carries out the detailed physics simulations for the experiment. It also includes software for analysis of both simulated and real data.

- G4beamline. It is used mainly for quicker, coarser studies, aimed primarily at preliminary construction, both of the detector building and components.

- MARS. It is used primarily for simulation of radiation fluxes and dosages; it is used to study average doses from radiation types of interest, largely for safety and health-related designs, both for equipment and people.

This particular background study is based on the Mu2e Offline package, which is further based on Fermilab's broader *art* software framework. The remainder of this section covers both.

4.1.1 Fermilab's *art* Software Framework

Fermilab's intensity-frontier software framework, the *art* [1] Event Processing Framework, has been designed for use in a relatively broad range of HEP-related applications (Kutschke et al., 2015):

- high-level software triggers

- online data monitoring

- calibration

- reconstruction

- analysis

- simulation

and has not been designed for use in data-acquisition interfaces, for direct hardware applications. Mu2e's Offline package currently uses *art* for simulation and analysis; as the Offline software evolves, it will utilize all of the functionality of *art*.

Its developers intended *art* to be a standard software platform for many different experiments, to address the common problem in HEP of each experiment's software infrastructure being too tightly coupled to its own specific code to be of much use to other experiments. This is obviously inefficient in a collaborative community, duplicating a lot of effort, and making it difficult for experiments in the community to compare results in a

---

[1] *art* is always lower-case, always italicized, and not an acronym.

standard fashion. In light of this goal, *art* was designed with clear boundaries between itself as a framework and the specific user code, where user code refers to code written by experiments with the intent to be used with *art*.

As may be no surprise to readers from the HEP community, *art* is based on the C++ programming language. The way users (users meaning experiments and people working on them) interact with *art* is through special C++ classes called *modules*. The concept of a module is widely used in programming, and here it simply has some specific constraints defined by the framework. The code of an experiment, which typically contains a great many C++ classes, is organized into these *art* modules, with the more basic, non-module classes being used within the modules. These modules are then configured in text files in the Fermilab Hierarchical Configuration Language (FHiCL, pronounced "fickle"), which end in the file extension `.fcl`. The most basic command for running *art* from a terminal looks something like "`art -c filename.fcl`", where the `.fcl` file basically acts as an instruction list for *art*, telling it which modules to use, in what order, and with what parameters. Figure 4.1 is a simple example of how a `.fcl` file configures an *art* job.

In the vein of standardization, *art*'s most basic unit of information is the *event*, in the software sense of the word. For triggered experiments like Mu2e, an event is all information associated with a single trigger. The next unit up in the information/bookkeeping hierarchy is the *subRun*, and then the *run*, where runs contain subRuns which contain events. Further than that, where subRuns and runs end and begin is completely determined by each experiment.

It was mentioned earlier that *art* modules are special C++ classes; what makes them special are the rules defined for modules by *art*. Every module in a `.fcl` file that is executed as an *art* job must provide code that is called once for each event. Further, any module *may* also provide code that is called at the start and end of the *art* job, the start and end of each run, and at the start and end of each subRun. Modules in *art* must also inherit,

```
#
#
# $Id: hello.fcl,v 1.1 2011/05/18 21:08:02 kutschke Exp $
# $Author: kutschke $
# $Date: 2011/05/18 21:08:02 $

#include "minimalMessageService.fcl"

# Give this job a name.
process : HelloWorld

# Start form an empty source
source : {
  module_type : EmptyEvent
  maxEvents : 3
}

services : {
  message : @local::default_message
}

physics :{
  analyzers: {
    hello: {
       module_type : HelloWorld
    }
  }

  p1 : [ ]
  e1 : [hello]

  trigger_paths  : [p1]
  end_paths      : [e1]

}
```

```
-uu-:---F1  hello.fcl      All L1    (Fundamental)-----------------------------------------
```

Figure 4.1: The above figure shows a trivial example of how a `.fcl` file configures an *art* job.

in the C++ sense, from one of several *art* base classes, overriding one or more of the pure-virtual member functions. Lastly, *art* modules must be one of the following types, defined by how they interact with events:

- *analyzer module* - may inspect information found in the event but may not add new information to the event.

- *producer module* - may inspect information found in the event and may add new information to the event.

- *filter module* - same functions as producer module, but may also tell *art* to skip processing of some, or all, modules for the current event; may also control which events are written to which output.

- *source module* - reads events, one at a time, from some source; *art* requires that every *art* job contain exactly one source module. A source is often a disk file but other options exist.

- *output module* - reads selected data products from memory and writes them to an output destination; an *art* job may contain zero or more output modules. An output destination is often a disk file but other options exist.

Where modules are the *art* C++ classes which execute tasks, *data products* are special C++ classes which are passive, typically being collections of information from modules in a simulation chain or actual data files from the data-acquisition system. In all cases, data products are almost entirely defined by the experiment.

There is a final special type of C++ class: *services*. Services in *art* are intended to support the management of information which is *not* represented in data products. Data products are passed from module to module in each event of an *art* job, but some

information is valid for bigger intervals like subRuns and runs, or even an entire job, like the geometry specifications, calibration information, and particle properties (for simulations). In *art*, services are the C++ classes which manage these types of information. Services are initialized at the beginning of an *art* job and may be called by any module in the event-loop.

To finish our *art* background discussion, there are some software packages used by experiments which are neither a part of *art* nor the user code; these are called *external products* within *art*. The FHiCL package used to execute *art* modules is one example. ROOT, a software framework widely used in HEP for large data-set histogram management, is also included in *art* as an external product. One last external product of note is Geant4 (G4), which is a powerful, broadly used software package in HEP that simulates particle interactions in realistic, macroscopic systems (like experiment control volumes).

4.1.2   Mu2e's Offline Package

The Mu2e Offline package, as mentioned before, is the workhorse of the experiment, and also the only software package developed and maintained entirely by Mu2e. The rest of this subsection will overview the rough divisions of code in Mu2e Offline, and call out some key directories that are of particular importance to this background study.

Mu2e Offline has a total of 70 directories which contain the code. Other than a few exceptions like the directories `lib` and `bin`, the basic structure of each directory is to contain a `src` subdirectory which has all of the `.cc` source files. Many also (a few *only*) contain an `inc` subdirectory, holding any associated `.hh` header files. Fewer, but still many others, also contain `fcl` subdirectories, holding the standard `.fcl` file(s) to execute for the related code, or `.fcl`'s to include in other `.fcl`'s for more general jobs.

34

The code in Mu2e Offline can be roughly grouped:

- "ExperimentComponentGeom" directories - these are directories purely containing code which models the geometry of the relative experiment component. They are linked to a common geometry text file which is included, through *art* services, in jobs.

- Utility/Helper/Service directories - these directories contain code which has a variety of miscellaneous supporting functionalities.

- Tracking directories - these directories contain code related to simulating the tracker and its response during collections.

- Reconstruction directories - these directories contain code which simulates the reconstruction of tracks using tracker hits, as well as calorimeter clusters, and also tracker-calorimeter track matching.

- Data Product directories - in the true *art* fashion, the various data products for Mu2e Offline are organized into their own directories. Notable data products are the `StepPointMC` class, which is a persistent data product used throughout a Monte Carlo (MC) simulation, and the `SimParticle` class, similar to the former. The difference is that `SimParticle` is used to represent an actual particle, with its various characteristics, while `StepPointMC` represents a more abstract "point on a track", which is also inside, or on the boundary of, a G4 volume within a MC simulation.

An important directory to call out separate from those listed above is `Mu2eG4`. This directory contains both the `G4_module` code and the `Mu2eG4_module` code, as well as many other modules and classes associated with G4. The G4 module is the basic module

35

used to implement the G4 external product, and the Mu2eG4 module is a re-working of it which was used in this background study. This directory contains all of the primary code used in the simulations for this study, handling all of the particle interactions and detector responses. Of all steps in the background study, only the event mixing, track reconstruction, and final analysis did not make use of G4 and the Mu2eG4 module.

One more important directory to this study is `JobConfig`. This is a very straightforward directory; it primarily contains `.fcl` file(s) for various background studies at Mu2e, and also the necessary text files that link in all of the Mu2e geometry code. The `.fcl`'s used in this background study are all contained in this directory.

4.2   Simulation & Analysis

This section will detail the simulation portion of the background study undertaken for this paper, and then the analysis. It will walk through the `.fcl`'s executed, using these as a guideline for pedagogically describing the simulation process. It will highlight important aspects and will discuss various intermediate quantities acquired during the simulation process. Finally, it will detail the event mixing and final analysis steps of the study.

4.2.1   Protons-On-Target to Stopped Pions

As described in Ch. 3, the muon beam essential to the experiment is created by colliding protons with the PT, which then gives us a beam of pions, which then decay somewhere in the TS area into muons (usually, and with other by-products as well). The RPC background occurs when pions from the pion beam *don't* decay, and instead make it all the way to the ST. So, during this study, POT's were generated and made to interact with the PT. The by-products of this interaction were propagated through the solenoid

36

system using G4, with pion decay disabled, and their proper times stored. The proper time information allows us to weight the sample of stopped pions by their survival probability in final analysis. They were then propagated through the ST and detector area, with information about stopped pions being stored: positions, times, and proper times. This sample of stopped pions, the final output of the first stage, was then used as input for later stages.

The first step in this part of the simulation was configured with `pions_g4s1.fcl`. This sub-stage of the simulation generated 10^{10} POT's, all with a single initial time value[2]. The by-products of the POT interactions with the PT were propagated through the experiment components, from the PT in the PS area up through the first half of the TS area to the central collimator. These particles were then "killed" upon entering the "TS3" volume of the TS area and their `SimParticle` and `StepPointMC` information was written out to disk.

Several preliminary cuts were applied in the earliest parts of this substage and kept throughout the first stage of the simulation, greatly improving the computation time of the simulation while having little to no effect on the RPC study:

- Neutrinos were cut from the simulation immediately upon their generation, if it occurred; neutrinos have no reasonable chance of affecting the rate at which pions stop in the ST.

- Electrons, positrons, photons, and neutrons initiated with kinetic energies below 100 MeV were immediately cut from the simulation; these are likewise extremely unlikely, due to their relatively low energies, to affect the rate of pions stopping in the ST.

- Any and all particles entering the "Hall Air" G4 volume (those leaving the solenoid

---

[2]This is obviously not representative of the POT pulse, but this is addressed in the final analysis stage.

system) were cut immediately upon entering; these are the least likely to contribute to changes in the stopped pion rate and likely save the most computation time.

The next substage is fairly straightforward in comparison, especially after all the preliminary cuts from `pions_g4s1.fcl`. This substage, configured with `pions_g4s2.fcl`, simply takes the information written out from the previous stage and creates "new" pions where the old ones were "killed", with the same kinematic information they ended with. These are then propagated up through the remaining TS area to the DS area, and "killed" again upon entering, kinematics written out exactly as before.

Finally, the last substage in stage one is configured with `pions_g4s3.fcl`, and brings the remaining pions (those which have not been lost to the solenoid walls) through to the ST. It also simulates their interactions with the ST, thus determining whether or not they stop. There is one final `.fcl` file that simply dumps the data from this substage's output into a ROOT *ntuple*[3]. This final ntuple contains the positions of all the stopped pions in the ST, their times, and their proper times for weighting against the pion survival probability in the analysis stage. Ultimately, the 10 billion generated POT's yielded 24,842,100 (unweighted) stopped pions.

At this point it is useful for the experiment as a whole, though not for this particular background study, to briefly diverge and examine the position distributions, times, and survival weighting of the captured pions in the ST. Figures 4.2 through 4.7 show various pion quantities for the captured pions. Note that the number of foils a pion penetrates before stopping, or in other words its z-direction[4] since z is the downstream direction, completely changes between the unweighted and weighted distributions. In the weighted distributions, pions with low survival probability (longer proper times), are removed by the weight factor, and the stopped-pion count becomes an increasing function

---

[3] An *ntuple* is a math term common in HEP. In simple terms, it is an ordered list of information.

[4] The z axis faces "downstream" parallel to the PS and DS solenoid axes at Mu2e. The y-axis is vertically upward, with the origin being the beam height. The x-direction is horizontal.

with z-position. This can be explained from the fact that pions with higher survival probability (shorter proper time) necessarily have higher momentum, since, from Special Relativity, momentum and proper time are correlated. Thus, the pions most likely to make it to the ST are those with higher momentum, and, having higher momentum, penetrate larger amounts of the ST material before being captured, if at all.

Another feature of note is contained in Figure 4.5. As described in Ch. 3, the central collimator of the TS acts as a momentum filter of sorts. Charged particles tend to separate vertically (along the y-axis) in the first 90° bend of the TS based on charge and momentum. Further than the very effective charge selection, this allows for shields to be placed which block particles outside a preferred range of momenta as well, simply by constraining the dimensions of the window through the shields. The end result of this selection at the collimator is that higher momentum particles tend to come through to the DS area with lower y-positions. This is represented in Figure 4.5 by the negatively skewed but otherwise Gaussian shape in stopped-pion y-positions.

Figure 4.2: The above figure shows the (unweighted) distribution of stopped-pion z-positions.

Figure 4.3: The above figure shows the (weighted) distribution of stopped-pion z-positions.

Figure 4.4: The above figure shows the (unweighted) distribution of stopped-pion x-positions.

Figure 4.5: The above figure shows the (unweighted) distribution of stopped-pion y-positions.
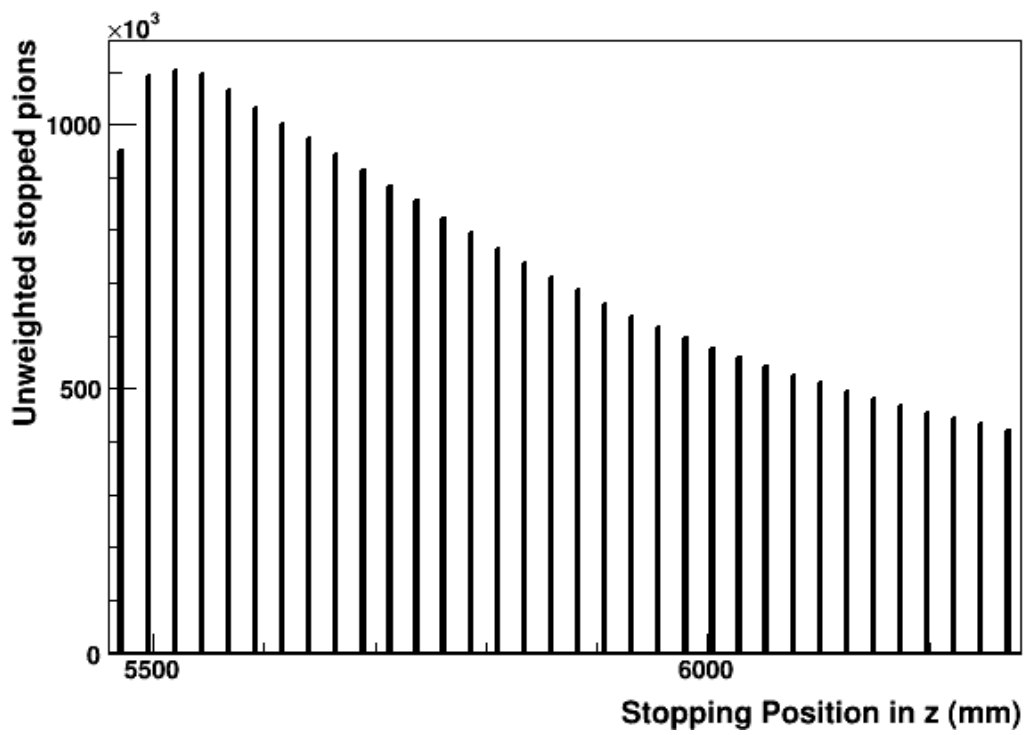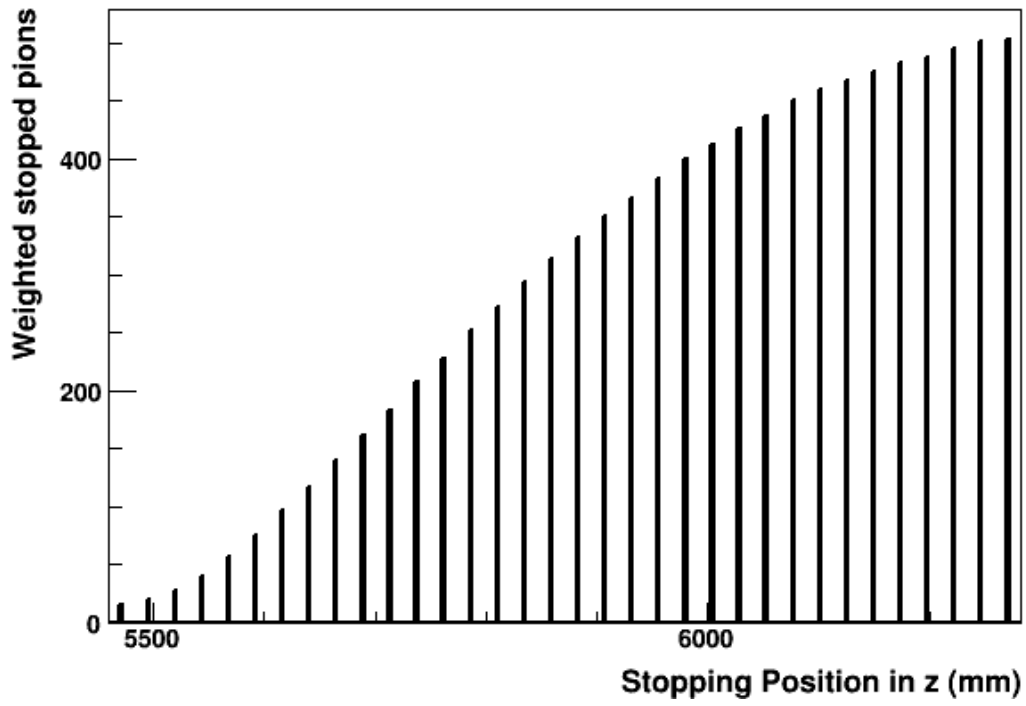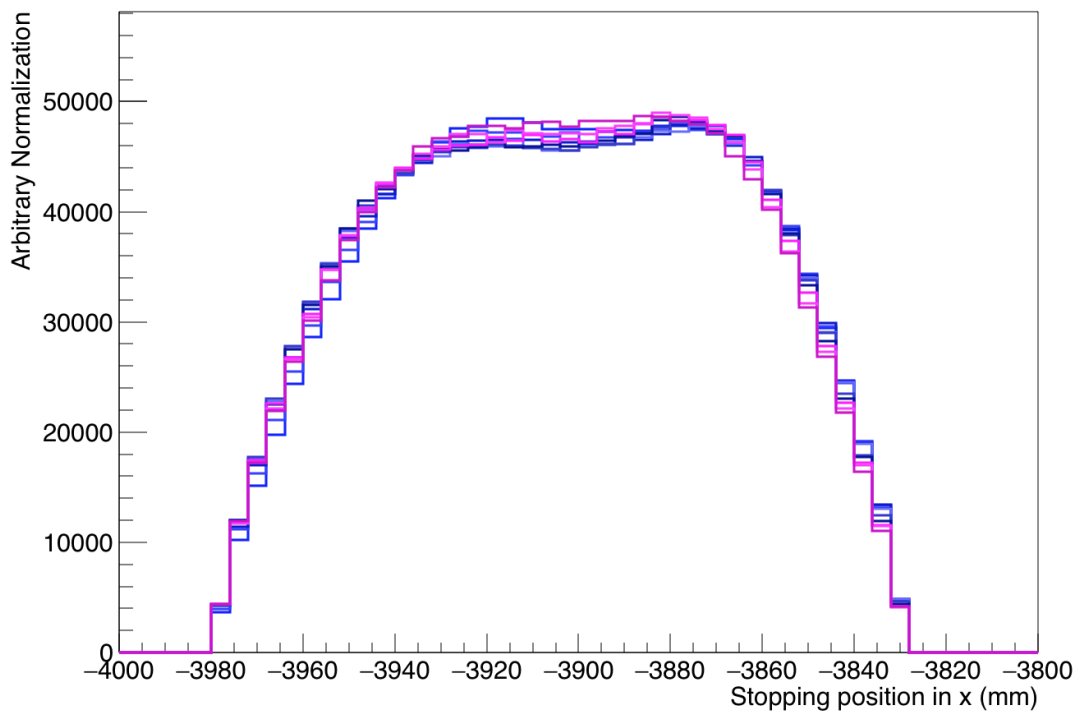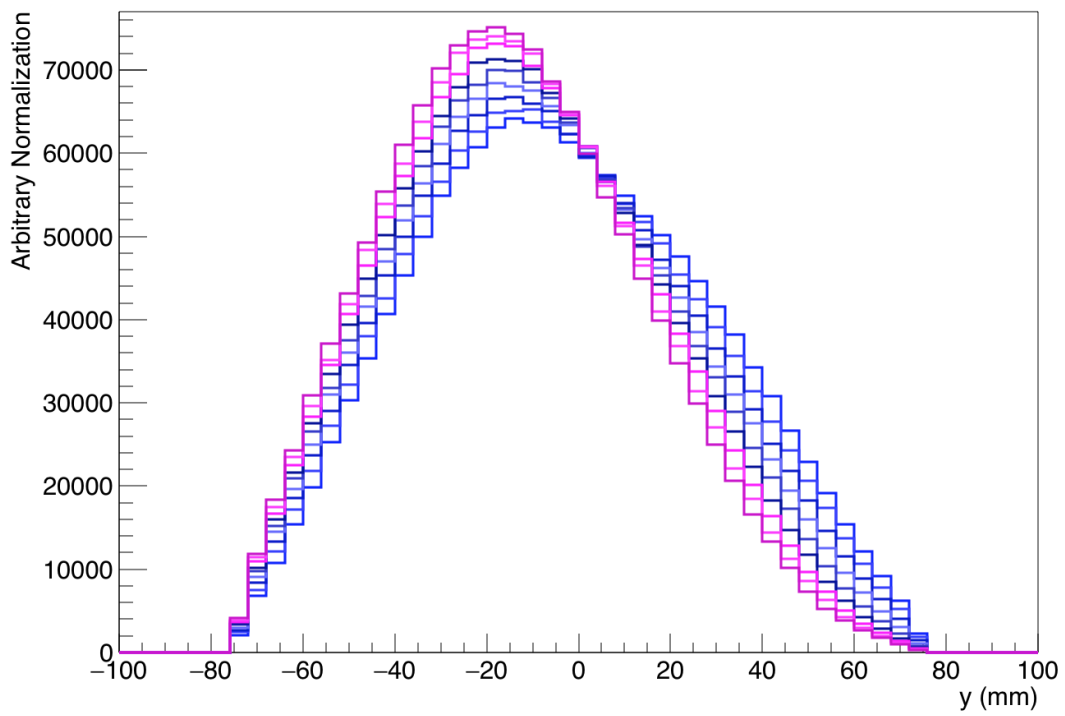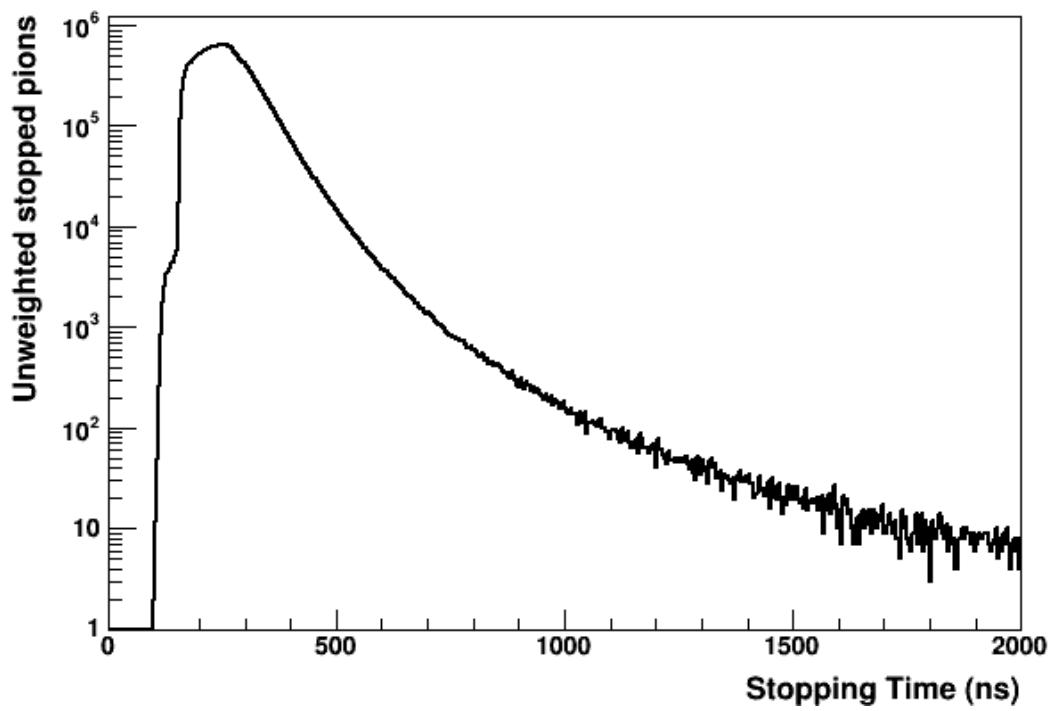
Figure 4.6: The above figure shows the (unweighted) distribution of stopped-pion times.
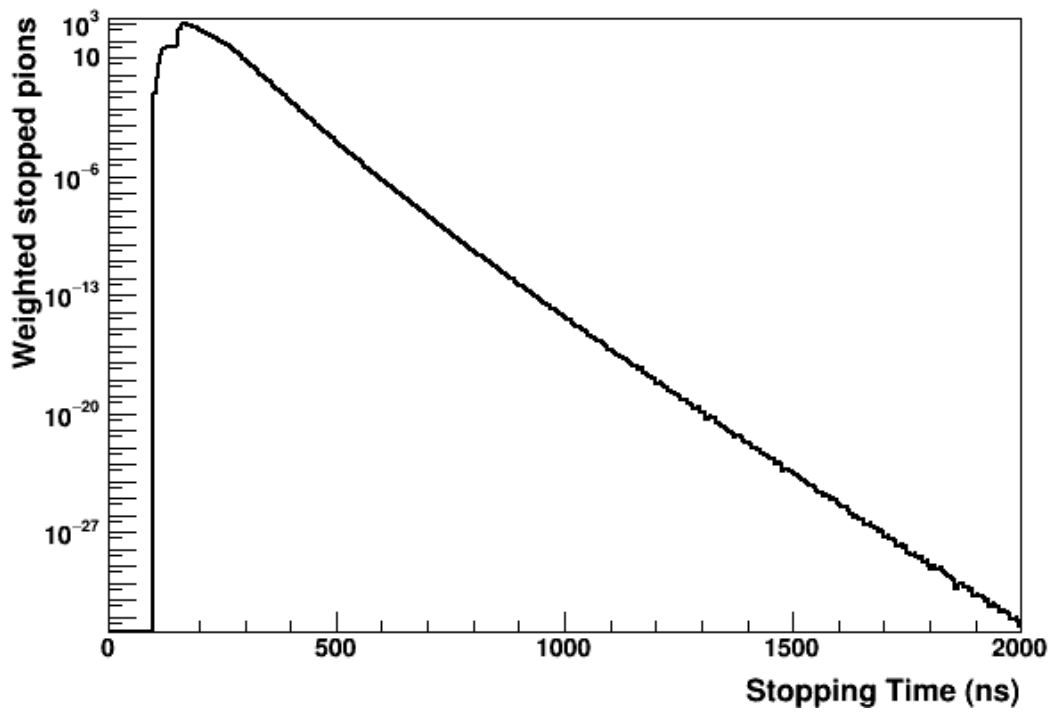
Figure 4.7: The above figure shows the (weighted) distribution of stopped-pion times.

4.2.2 Stopped Pions to Pair Production to Tracker Hits

The stopped pion information coming from the end of stage 1 was used as the input for stage 2, which simulates the RPC process and the photon which is produced. As described in Ch. 2, the *external* conversion yields a real photon, which means that it will travel away from the capture point, possibly pair producing elsewhere in the presence of other material. The *internal* conversion process yields a virtual photon, which experimentally means that it will almost instantaneously pair produce, essentially *at* the capture point. For both types of conversion, the photon energy spectrum was assumed to be the Bistirlich distribution. Each conversion type was simulated, and the simulation carried out through the pair production process and propagation of the e^+e^- pair, all the way out to hits in the tracker and calorimeter. These hits, as well as the kinematic information of the particles making them, were written to the output file. The two files used to configure these jobs were `pions_g4s4_RPC.fcl` and `pions_g4s4_IntConv.fcl`, associated with the external and internal conversions, respectively. It's worth noting that this is final stage which makes use of G4 and the `Mu2eG4` module.

The only preliminary cut applied at this stage was the range of the Bistirlich energy distribution that was sampled; photons were generated with energies between 100 MeV and the tail of the spectrum, 140 MeV. Photons below 100 MeV are highly unlikely to yield electrons in the signal region. This energy constraint could have been applied to the e^+ and e^- as well, but this was neglected to ensure simplicity of the simulation.

4.2.3 Reconstruction & Analysis

The final stage of the simulation takes the output file from stage 2, containing the tracker and calorimeter hits and the corresponding particles' kinematic information, as input. This stage performs the reconstruction of tracks based on these hits, using the

reconstruction software available in Mu2e Offline. Recall that the proper time information for these particles and their reconstruction tracks has been kept throughout the simulation. This means that the output of the final stage contains both the weighted and unweighted number of reconstructed tracks, where the weight factor is the survival probability of the stopped-pion from which these originated. For the purposes of standardization across the various background studies at Mu2e, this stage also then applies the standard set of cuts for Mu2e, rather than this being separated into another stage.

These cuts are as follows, applied in the order presented:

- Track status - simple cut based on whether the reconstruction software could even successfully fit a track to the hits or not.

- Track quality constrained to be between 0.4 and 1.3 - the reconstruction software available at Mu2e also assigns a track quality measure to each track, which characterizes the confidence in the reconstruction; both very high and very low values of track quality are bad

- pitch (polar angle) constrained between 45° and 60° - this reduces background due to high energy electrons from other processes, like muon and pion decay-in-flight, or that simply come from the beam somewhere in the TS area

- Minimum track transverse radius constrained between -80 mm and 105 mm - this cut basically ensures that the track actually comes from somewhere in the ST; any that don't are unambiguosly not signal

- Maximum track transverse radius constrained between 450 mm and 650 mm - this requires that the particle's track actually intersects the tracker straws

- Track initial time constrained between some low value and 1695 ns - the initial time of the first hit in the track; the low value is chosen to optimally reduce the RPC

background while accepting as much signal as possible. This cut actually determines the "livegate" of the experiment, when the detector components are collecting data

- Tracks must match calorimeter hits - this simply rejects tracks which do not have corresponding calorimeter hits; it helps reduce mis-reconstruction backgrounds

- Tracker-Calorimeter χ^2 constrained to be less than 100 - this ensures the calorimeter matching with the tracker is a physically realistic one

- The energy deposited in the calorimeter constrained to be between 10 MeV and 120 MeV - mostly a redundant cut this far down the cut-flow and after all the preliminary cuts

- A particle-identification (PId) algorithm is applied - this algorithm combines information from the tracker and the calorimeter to accomplish the subtle task of separating particles like electrons from muons, which look very much the same to most of the detector components

- Momentum of the track constrained between 103.85 MeV and 105.5 MeV - tracks outside this window are unambiguosly *not* signal tracks

In stage 1, the POT's were generated all with a single initial time value. It was mentioned that this was addressed in the analysis stage. When the final stage job is run, it shifts the initial times of every particle, and therefore the tracker/calorimeter hits associated with them, by a random amount. This random amount is sampled from a function representing the POT pulse distribution in time, using the initial time as the midpoint. The size of the time shift is then the absolute value of the difference between the original initial time and the randomly generated time.

48

It was mentioned in Ch. 3 that the beam pulses are not perfectly tuned, and that there is a non-zero number of POT's which arrive at the PT outside of the nominal pulse window. These POT's are classified as out-of-time (OOT), as opposed to in-time. In order to represent this effect in this background study, the final stage `.fcl`'s again shift the initial times of every particle and their corresponding tracker/calorimeter hits, only with a different probability distribution function. In this OOT case, each particle had a new initial time generated according to a uniform random distribution, where the range was the full amount of time between pulses, excluding the pulse window area. The OOT contribution to background is then modulated by something called the "extinction factor", which is a representation of how well we expect to be able to extinguish excess beam outside the nominal pulse window. This factor is determined by other studies, and a conservative nominal value was used in this study.

Due to the fact that the external and internal conversions are modeled separately, and that both in-time and OOT studies were done, there were four basic configurations for this final stage: `dra_pure_pions_RPC.fcl` & `dra_pure_pions_RPC_oot.fcl`, and `dra_pure_pions_IntConv.fcl` & `dra_pure_pions_IntConv_oot.fcl`. The outputs of these jobs were more ROOT ntpules. The final step in the full background study was to use some basic scripting to take the (weighted) tracks which were still accepted after the standard cut-set from each job's output and apply the appropriate normalization factors. These normalize the background contributions to the full run-time of the experiment.

CHAPTER 5

RESULTS

5.1 Cut-flow Histograms

Figures 5.1 through 5.4 show the (weighted) number of tracks rejected by each cut. The cuts were applied successively, left to right. The last column shows the accepted tracks.

Figure 5.1: The above figure shows the cut-flow histogram for the in-time, external conversion study. The cuts were applied successively, left to right.

Figure 5.2: The above figure shows the cut-flow histogram for the out-of-time, external conversion study. The cuts were applied successively, left to right.
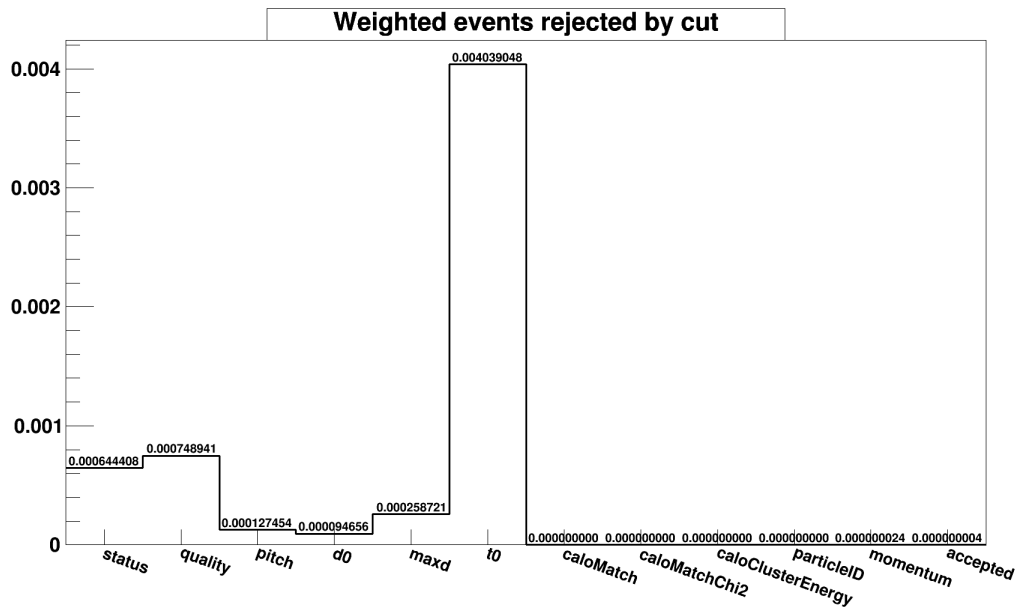
Figure 5.3: The above figure shows the cut-flow histogram for the in-time, internal conversion study. The cuts were applied successively, left to right.
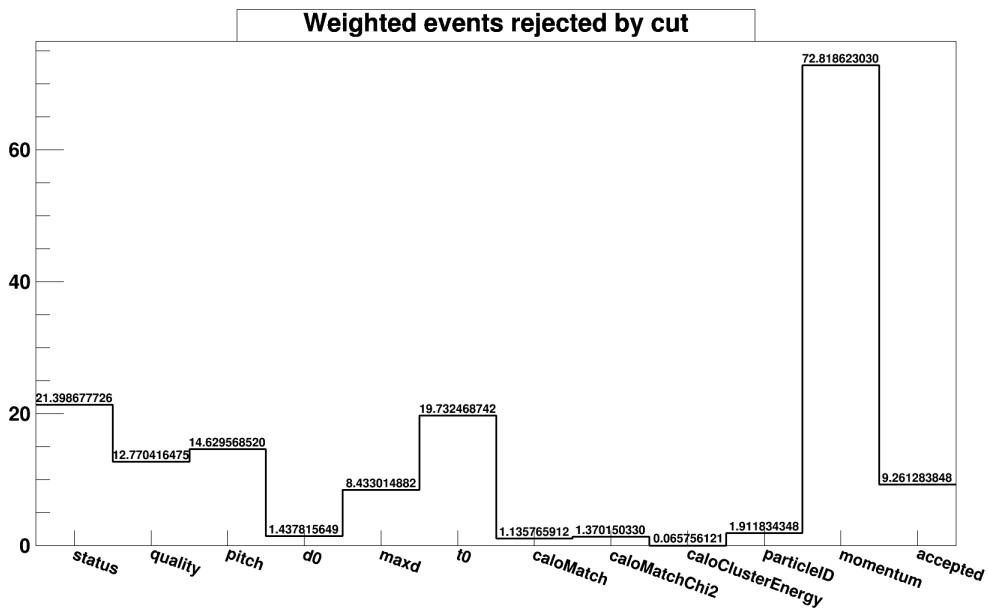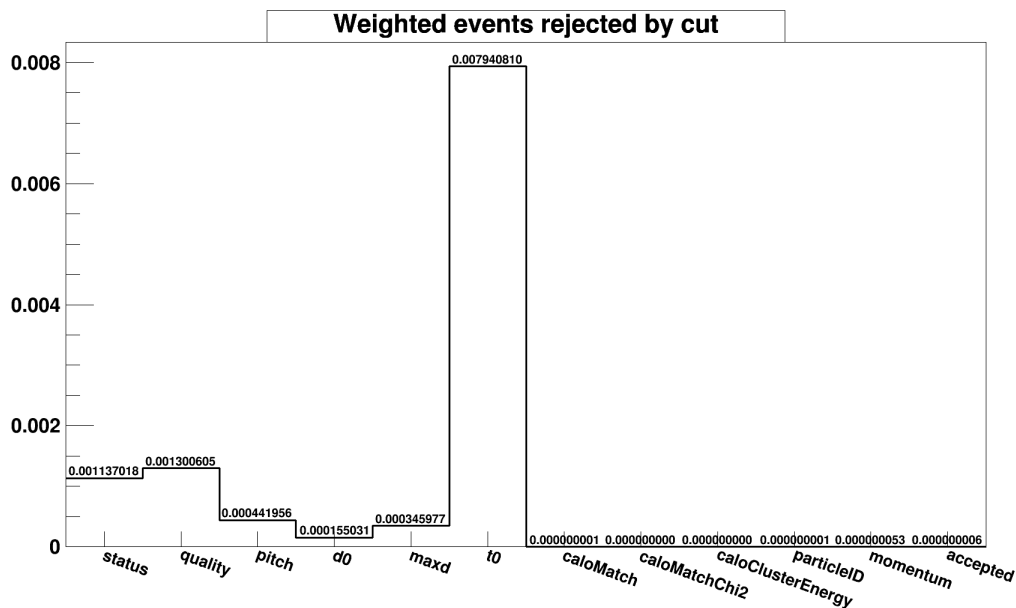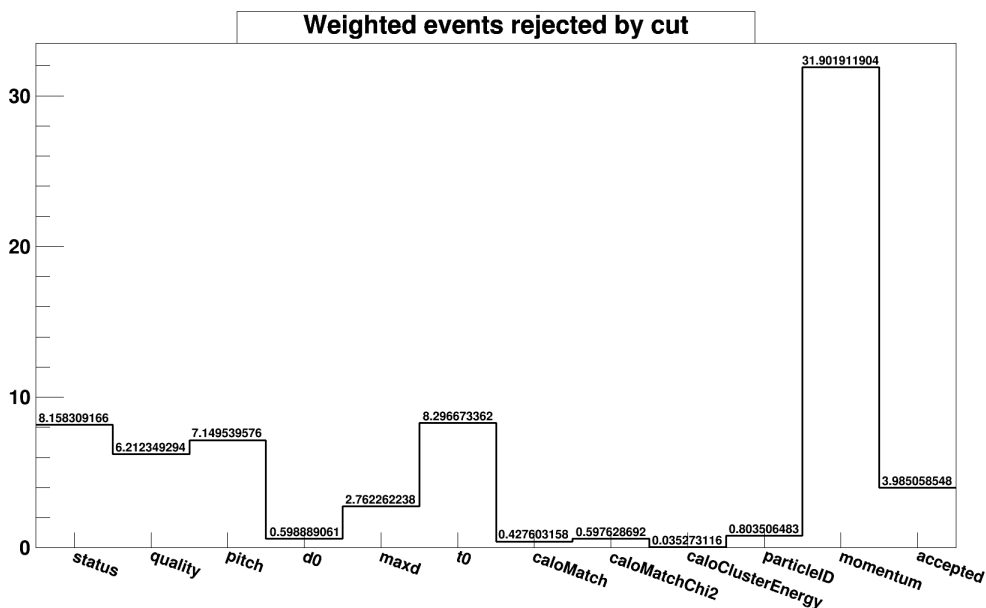

Figure 5.4: The above figure shows the cut-flow histogram for the out-of-time, internal conversion study. The cuts were applied successively, left to right.

5.2 Normalization

The number of accepted tracks at the end of each simulation is not the end of the background study. As detailed in Ch. 4, the simulation was broken into stages and substages, and the results must be normalized down the simulation chain to match them to the original number of POT's. Further, the number of POT's must also be normalized to the number we expect for the lifetime of the Mu2e experiment. The correct normalizations are given by the following equations

$$N^{e^-}_{external} = N_{POT} \cdot \xi \cdot p^{\pi^-(\tau=\infty)}_{stop} \cdot f_{RPC} \cdot F_\gamma(E_1, E_2) \cdot [\frac{1}{n_{sim}} \sum_{i=1}^{n^{accepted}_{sim}} w_i]$$

$$N^{e^-}_{internal} = N_{POT} \cdot \xi \cdot p^{\pi^-(\tau=\infty)}_{stop} \cdot f_{RPC} \cdot F_{\gamma^*}(E_1, E_2) \cdot \rho_{\gamma^* \to e^+ e^-} \cdot [\frac{1}{n_{sim}} \sum_{i=1}^{n^{accepted}_{sim}} w_i]$$

where the terms are defined as follows:

N_{POT} - the number of protons-on-target for the life of the experiment

ξ - the beam extinction factor (note, taken to be 1 for in-time calculations)

$p^{\pi^-(\tau=\infty)}_{stop}$ - probability of an infinite lifetime pion of stopping on the ST

f_{RPC} - probability that a stopped pion will go through an RPC process

$\rho_{\gamma^* \to e^+ e^-}$ - the internal conversion coefficient

$F_\gamma(E_1, E_2)$ - the probability function associated with the Bisirlich distribution on the photon energy spectrum

$F_{\gamma^*}(E_1, E_2)$ - same as above but for the internal conversion process

n_{sim} - the number of simulations in stage 2, where internal and external photons were simulated using the stopped-pion locations as input

$n_{sim}^{accepted}$ - the number of stage 2 simulation events which pass the standard cuts and are accepted as Mu2e signal

w_i - the weighting of the event based on the stopped-pion survival probability

5.3 Tables of Background Contribution

Due to this background's very prompt nature, it has a strong dependence on the livegate of the experiment. To address this, multiple initial time cuts were applied, as mentioned in Ch. 4. The following tables show the background contributions of the RPC process for various values for the initial time cut. All uncertainties in the tables are statistical only. For the OOT calculations, the extinction factor was taken to be $\xi = 8.2 \times 10^{-13}$. The final result is reported as a function of the extinction factor.

TABLE 5.1
The external conversion process contribution to background.

t0 cut		In-Time	OOT ($\times 10^{-5}$)
500 ns	242.319	242.319 ± 9.5235	$11.0602 \pm 1.40109\text{e-}1$
550 ns	13.262	13.262 ± 0.7067	$10.5977 \pm 1.37104\text{e-}1$
600 ns	1.04452	1.04432 ± 0.07421	$10.0917 \pm 1.33593\text{e-}1$
650 ns	0.082320	0.082128 ± 0.01098	$9.59646 \pm 1.30590\text{e-}1$
700 ns	0.0071056	0.006923 ± 0.000939	$9.12578 \pm 1.27143\text{e-}1$
750 ns	0.00065720	0.0004838 ± 0.0000988	$8.66986 \pm 1.23530\text{e-}1$

TABLE 5.2

The internal conversion process contribution to background.

t0 cut		In-Time	OOT ($\times 10^{-5}$)
500 ns	316.874	316.874 ± 9.9181	13.2397 ± 2.51965e-1
550 ns	20.5189	20.5188 ± 0.9490	12.7328 ± 2.47510e-1
600 ns	1.206262	1.20614 ± 0.09750	12.2000 ± 2.43416e-1
650 ns	0.1161766	0.11606 ± 0.01099	11.6596 ± 2.37965e-1
700 ns	0.00848715	0.008376 ± 0.001068	11.1151 ± 2.28075e-1
750 ns	0.00092142	0.0008160 ± 0.0001129	10.5419 ± 2.22299e-1

TABLE 5.3

The total background calculation, separated into in-time and out-of-time components.

t0 cut	In-Time	OOT
500 ns	559.193 ± 13.750	0.000243 ± 2.8838e-6
550 ns	33.7808 ± 1.1832	0.000233 ± 2.8303e-6
600 ns	2.25046 ± 0.12253	0.000223 ± 2.7774e-6
650 ns	0.19819 ± 0.01554	0.000213 ± 2.7151e-6
700 ns	0.01530 ± 0.00142	0.000202 ± 2.6118e-6
750 ns	0.00130 ± 0.00015	0.000192 ± 2.5438e-6

5.4 Systematic Uncertainties

The following are the systematic uncertainties which have been determined by the previous studies of this background and also by the original study of the RPC process using a Magnesium target (Bistirlich et al., 1972).

- f_{RPC} : 9.3%

- IntConv Coefficient : 5.5%

- Virtual Photon Spectrum : 30%

- POT Shape/Extinction : 10%

The internal conversion coefficient and virtual photon spectrum apply only to internal conversions. Thus, the combined systematic for externals is $\sim 13.7\%$; for internals, it is $\sim 33.4\%$

5.5 Final Result

Using the nominal initial time cut of 700 ns, the total background to the Mu2e process coming from RPC processes, normalized to the life of the experiment, is:

$$N_{RPC}^{e^-} = 0.01530 + 0.000202 \cdot \frac{\xi}{8.2 \times 10^{-13}}$$

REFERENCES

J. A. Bistirlich, K. M. Crowe, A. S. L. Parsons, P. Skarek, and P. Truöl. Photon spectra from radiative absorptions of pions in nuclei. *Physical Review C*, 5:1867–1883, Jun 1972.

M. Born. *Einstein's Theory of Relativity*. Dover Publications, 1962.

Fukuda, Y. et. al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81: 1562–1567, Aug 1998. doi: 10.1103/PhysRevLett.81.1562. URL `http://link.aps.org/doi/10.1103/PhysRevLett.81.1562`.

D. J. Griffiths. *Introduction to Elementary Particles*. Wiley-VCH, 2008. The primary source for HEP and SM background information.

A. Ilakovac, A. Pilaftsis, and L. Popov. Charged Lepton Flavor Violation in Supersymmetric Low-Scale Seesaw Models. *arXiv:1212.5939v4*, Feb. 2013.

R. Kutschke, M. Paterno, and M. Wang. *Intensity Frontier Common Offline Documentation: art Workbook and Users Guide*. Fermilab Scientific Computing Division, Aug. 2015. The primary source for the *art* background info in Ch. 4.

S. P. Martin. A Supersymmetry Primer. *arXiv:hep-ph/9709356v7*, Jan. 2016.

Particle Data Group. The Particle Adventure. URL `http://www.particleadventure.org/standard%5fmodel.html`.

The Mu2e Collaboration. Mu2e Technical Design Report. *arXiv:1501.05241v2*, Mar. 2015. The primary source for the information in Ch. 3.

The Mu2e Collaboration. Official Mu2e Collaboration List, 2016.

The Nobel Foundation. The Nobel Prize in Physics 1918.

L. L. Whyte. *Essay on Atomism*. Wesleyan University, 1961.

P. Wintz. Results of the SINDRUM-II experiment. *Conf. Proc.*, C980420:534–546, 1998.

CURRICULUM VITAE

NAME: Jacob Ross Colston

ADDRESS: Department of Physics

                   University of Louisville

                   Louisville, KY 40292

DOB: January 23, 1990

EDUCATION: B.S. Physics

                   University of Louisville

                   2014

WORK EXPERIENCE

University of Louisville

Graduate Teaching Assistant, 2014-2016

Dept. of Physics & Astronomy

University of Louisville

Graduate Research Assistant, 2014-2016

UofL HEP Group

University of Louisville, PRIMES Program

Undergraduate Teaching Assistant, 2013-2014

Dept. of Physics & Astronomy

University of Louisville

Undergraduate Research Assistant, 2014-2016

UofL HEP Group

Duke Energy

Mechanical Engineering Co Op, 2010-2012

Gallagher Station

Kristie Beaven (Engineering Manager)

TEACHING EXPERIENCE

University of Louisville

Intro to Modern Physics Laboratory (Physics 301)

Calculus Based Physics Laboratory I (Physics 295)

Elementary Physics Laboratory (Physics 111)