

Distribution Agreement

In presenting this thesis as a partial fulfillment of the requirements for a degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis in whole or in part in all forms of media, now or hereafter now, including display on the World Wide Web. I understand that I may select some access restrictions as part of the online submission of this thesis. I retain all ownership rights to the copyright of the thesis. I also retain the right to use in future works (such as articles or books) all or part of this thesis.

Madeleine Bloomer

April 8, 2025

Analysis of Potential Target Materials for Production of High Intensity Secondary Beams

by

Madeleine Bloomer

Erin Bonning
Adviser

Department of Physics

Erin Bonning
Adviser

Justin Burton
Committee Member

Simbarashe Nkomo
Committee Member

2025

Analysis of Potential Target Materials for Production of High Intensity Secondary Beams

By

Madeleine Bloomer

Erin Bonning

Adviser

An abstract of
a thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Science with Honors

Department of Physics

2025

Abstract

Analysis of Potential Target Materials for Production of High Intensity Secondary Beams By Madeleine Bloomer

High intensity secondary beams are becoming increasingly important for the field of high energy physics (HEP), necessitating advancements in production targetry. Here, we consider some of the crucial factors for choosing a target material, including secondary production, radiation damage, and energy deposition, in the context of designing a production target for the Mu2e experiment. We use Geant4 and G4Beamline simulations to study secondary production across various materials and FLUKA simulations to study radiation damage and energy deposition across various materials. We find that number of nuclear interaction lengths may be a proxy for secondary production and that number of radiation lengths may be a proxy for damage to the target material. Thus, we propose that considering radiation length and nuclear interaction length is a valuable starting point for choosing production target materials, both for Mu2e and future accelerator contexts.

Analysis of Potential Target Materials for Production of High Intensity Secondary Beams

By

Madeleine Bloomer

Erin Bonning

Adviser

A thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Science with Honors

Department of Physics

2025

Acknowledgements

I would like to begin by extending my gratitude to my adviser, Dr. Erin Bonning, both for giving me the opportunity to continue with this work and for providing me with my first branch of support as I transitioned to a new campus and a new field. I never anticipated when I first stepped foot in the Emory physics department that I would be submitting a thesis at all, much less one on high energy physics, so I am endlessly grateful to her for allowing that to happen. Additional thanks to Dr. Alissa Bans for facilitating the process between Emory and Fermilab and to Drs. Justin Burton and Simbarashe Nkomo for contributing their time to serving on my committee.

Furthermore, I cannot understate my appreciation for my mentors from Fermilab, Drs. Kevin Lynch and Michael Hedges. Thank you for introducing me to this little corner of high energy physics, for spending countless hours helping me understand physics that I was certainly lacking the background for, for letting me sit in on meetings with interesting people, for always being excited to see a new plot, and for so much more. Most of all, thank you for believing in me through all of this, from my first weeks of my internship through the grad school application process and now as I finish writing my thesis and approach my second summer at Fermilab. I would also like to extend my thanks to everyone else in the Target Systems Department who have continually answered my questions, taught me to use new software, and explained engineering parameters. I appreciate all of your support.

I would be remiss not to mention all of the people at Oxford College who have supported me throughout my time there and beyond. In particular, thank you to Dr. Alfred Farris for spending hours every week encouraging my interest in physics and for being the first person to try to teach me E&M—believe it or not, I'm actually choosing to continue in a field where that's particularly important, so hopefully I know something. Additionally, thank you to Professor Brenda Harmon for her unconditional support from my very first research experience to my last experiences at Emory and for showing me the value of good pedagogy and mentorship. You have taught me so much about being a teacher, and I strive to be like you every day.

Finally, I have to thank my parents for being my built-in support from the very beginning. Thank you for never leaving me with any doubt that I could be anything I wanted to, from my blimp pilot dreams at three years old to my current aspirations of pursuing accelerator physics. I am so lucky to have had that confidence instilled in me so young.

Contents

1 Introduction	1
1.1 The Standard Model	1
1.2 Charged Lepton Flavor Violation and $\text{Mu}2e$	3
1.3 Production Targetry	8
2 Methods	16
2.1 Software	16
2.2 Muon Production	17
2.3 Material Damage	19
3 Muon Production	21
3.1 Results	21
3.2 Discussion	27
4 Material Damage	31
4.1 Results	31
4.2 Discussion	35
5 A General Framework	39
6 Conclusions	42
Appendix A	44
Appendix B	47
Appendix C	50
References	58

List of Figures

1.1 The Standard Model of particle physics.	2
1.2 A summary of the frontiers of particle physics.	3
1.3 A Feynman diagram of typical muon decay into an electron and two neutrinos.	4
1.4 A Feynman diagram of the “allowed” flavor violating muon-to-electron conversion via neutrino oscillation.	5
1.5 Flavor violating muon to electron decay in the presence of a nucleus via a virtual photon.	6
1.6 A more general diagram of flavor violating muon-to-electron conversion in the presence of a nucleus. The dot in the center could represent any number of mechanisms, and Mu2e does not rely on understanding what that dot might represent.	7
1.7 Mu2e experimental setup.	7
1.8 Proposed tungsten target design for Mu2e.	10
1.9 Pion and muon production study from IMCC.	11
1.10 Bicycle wheel mount for the production target.	14
3.1 Total muon production per POT versus Z	22
3.2 Total muon production per POT versus density.	22
3.3 Low momentum muon production per POT versus density.	22
3.4 Muon stops per POT versus density.	23
3.5 Muons per POT versus number of nuclear interaction lengths.	24
3.6 Low momentum muons per POT versus number of nuclear interaction lengths.	24
3.7 Muon stops per POT versus number of nuclear interaction lengths.	24
3.8 Low momentum muons per POT versus number of nuclear interaction lengths for graphite, molybdenum, and tungsten cylinders of varying length.	25
3.9 Total muon production per POT versus target volume for graphite, molybdenum, and tungsten cylinders of varying volume.	26
3.10 Position of pion production for tungsten, Inconel, molybdenum, and graphite targets with the Hayman 2 target geometry.	27
4.1 Peak DPA per year versus density.	32

4.2 Peak DPA per year versus Z^2	32
4.3 Peak DPA per year except DPA caused by heavy ions versus Z	33
4.4 DPA per year by particle for a tungsten cylinder.	34
4.5 DPA per year by particle for a graphite cylinder.	34
4.6 Approximate target temperature versus number of radiation lengths.	35
5.1 Nuclear interaction length versus radiation length for a variety of materials.	40
5.2 Nuclear interaction length versus radiation length for a variety of materials (on broader axes).	40

Chapter 1

Introduction

1.1 The Standard Model

The Standard Model of particle physics describes our current understanding of the most fundamental particles and forces which make up the universe. These fundamental particles are shown in Fig. 1.1, where they are separated into two main categories: the matter particles and the force carriers. The matter particles are further broken down into quarks and leptons, each of which have three generations. The least massive of these generations are the ones which we interact with most, namely the up and down quarks (which make up protons and neutrons) and the electron. The more massive generations are unstable because it is favorable for them to decay into their less massive counterparts, so to study them, we have to collide the more stable particles we have easier access to (i.e., the less massive ones) at very high energies.

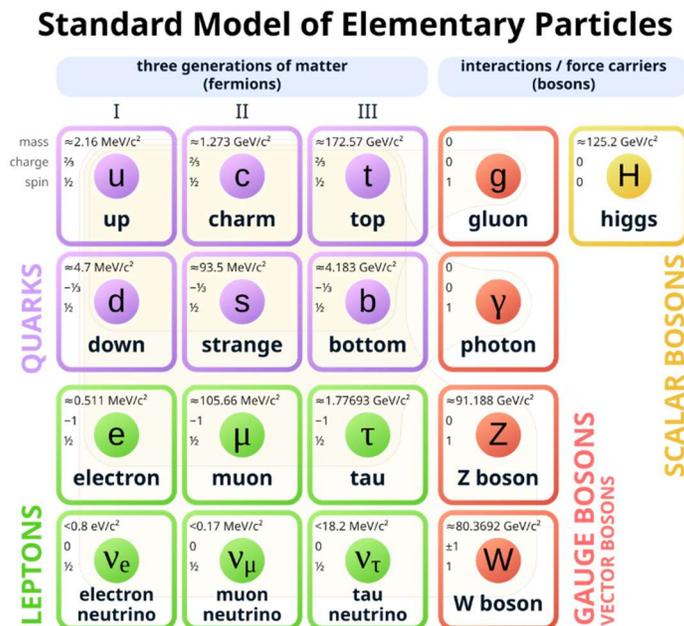


Figure 1.1: The Standard Model of particle physics [1].

The high energy physics (HEP) community breaks the field into categories called frontiers. The three overarching frontiers, which are summarized in Fig. 1.2, are the energy frontier, the intensity frontier, and the cosmic frontier. The cosmic frontier explores the large-scale structure of the universe through particle physics observations, so the high energy accelerator experiments (i.e., those experiments where we collide particles at high energies) generally fall into the other two frontiers. The energy frontier aims to discover new physics through the highest energy experiments and is the most well-known, as it houses the Large Hadron Collider (LHC). The intensity frontier, which, as the name suggests, employs high intensity particle beams, is further broken down into two subcategories: precision measurements and rare processes. Precision measurement experiments aim to measure physical observables to higher precision; for instance, the recent Muon g-2 experiment at Fermilab made the most precise measurement yet of the muon's magnetic moment. Rare

processes experiments aim to observe specific theoretical processes which could provide evidence of new physics.

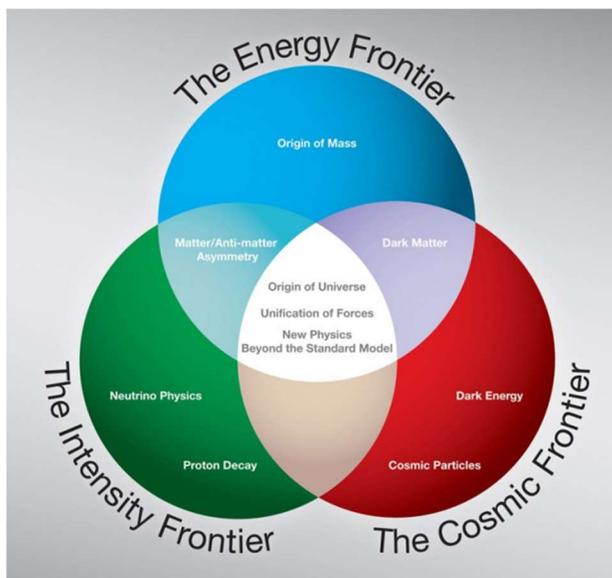


Figure 1.2: A summary of the frontiers of particle physics [2].

1.2 Charged Lepton Flavor Violation and $\text{Mu}2e$

As of today, no experiment has ever violated the Standard Model, but there are plenty of ideas about where to look to find physics beyond the Standard Model (BSM). One of these ideas, charged lepton flavor violation (CLFV), is currently one of the major focuses of the rare processes frontier. Each of the generations of leptons shown in Fig. 1.1 has its own “flavor”; e.g., the electron and electron neutrino both have electron flavor. By the Standard Model, this lepton flavor is generally conserved, just as properties like charge are conserved. Take, for example, the process by which muons decay into electrons, shown in Fig. 1.3. The original muon has muon flavor, and it decays into a muon neutrino, an electron, and an electron antineutrino by way of a W boson. The muon neutrino preserves the original muon flavor,

and the electron and electron antineutrino add up to zero lepton flavor (lepton antiparticles have opposite lepton flavor to their counterparts), so lepton flavor is conserved.

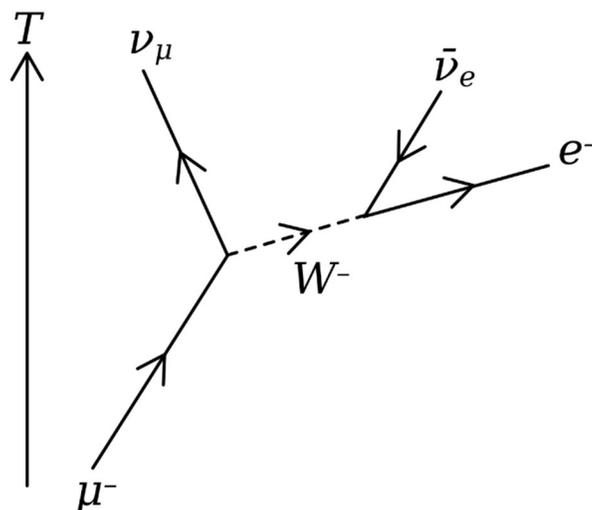


Figure 1.3: A Feynman diagram of typical muon decay into an electron and two neutrinos [3].

Lepton flavor conservation (and most of the Standard Model, for that matter) assumes that neutrinos have no mass. However, it has been discovered that neutrinos “oscillate,” or change flavor, which means that they must have mass. This has led to an interest in the possibility of charged leptons violating flavor conservation. There is an “allowed” way for this to happen via neutrino oscillation—in fact it *must* happen—which is shown for a muon to electron conversion in Fig. 1.4. It is possible to calculate the probability of this occurring, though, and that turns out to be on the order of 10^{-54} —in other words, we will never observe this. Therefore, if we were to observe a charged lepton decay which does not conserve flavor, it would almost certainly be evidence of BSM physics, which is why CLFV is of such interest.

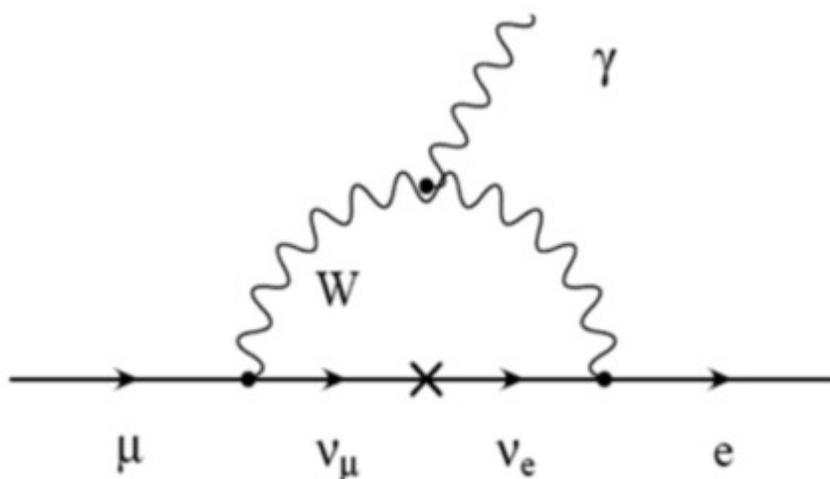


Figure 1.4: A Feynman diagram of the “allowed” flavor violating muon-to-electron conversion via neutrino oscillation [4].

There are multiple experiments which aim to observe CLFV, including the Mu2e experiment at Fermilab. Mu2e is specifically looking for flavor violating muon-to-electron conversion in the presence of a nucleus, which is experimentally favorable for a few reasons. For one, it has a clear experimental signature. To detect a muon decay with the products shown in Fig. 1.4—i.e., a muon decaying into an electron and a photon, though not necessarily via the mechanism shown in Fig. 1.4—you would have to detect both the electron and the photon as well as trace both particles back to the original muon (this is actually the premise of the Mu to E Gamma, or MEG, experiment at the Paul Scherrer Institute). However, if you had a muon decay this way in the presence of a nucleus, as shown in Fig. 1.5, the photon becomes a virtual particle (a complicated concept, but for the scope of this work, it’s represented by an internal line in a Feynman diagram, and we don’t have to consider it in the products of the conversion for things like energy conservation). Additionally, we can consider the muon to be at rest when it decays because it has been captured by the nucleus, so the muon’s energy at the time of decay is simply its rest mass energy, and we can consider

the nucleus at rest because it is in a lattice. As a result of this and conservation of energy, the electron produced from the decay gets nearly all of the muon's energy (minus a negligible amount which goes to the nucleus). Therefore, not only do we not have to worry about measuring the photon because it's virtual, we also don't have to worry about tracing the electrons anywhere because we know that they will have a characteristic energy equal to the rest mass energy of the muon. If we detect an electron with that energy, then we will have evidence of CLFV.

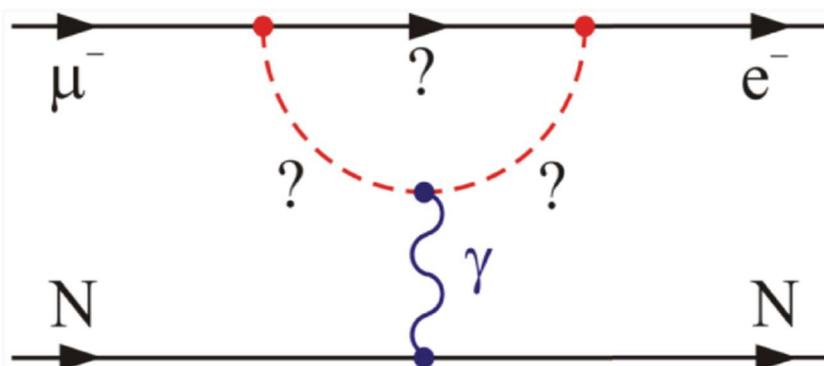


Figure 1.5: Flavor violating muon-to-electron conversion in the presence of a nucleus via a virtual photon [5].

Furthermore, looking for muon-to-electron conversion in the presence of a nucleus expands the possible mechanisms of CLFV we can detect. It is possible to detect that conversion with a photon shown in Fig. 1.5, but because we aren't concerned with detecting anything other than the electron (and we know the muon and nucleus are there), the mechanism that happens in between isn't necessary to understand to do the experiment. A more specific experiment like MEG puts more constraints on the mechanism which it can detect (e.g., there has to be a photon involved), whereas Mu2e can detect any mechanism of flavor violating muon to electron conversion which happens in the presence of a nucleus, which is outlined in Fig. 1.6.

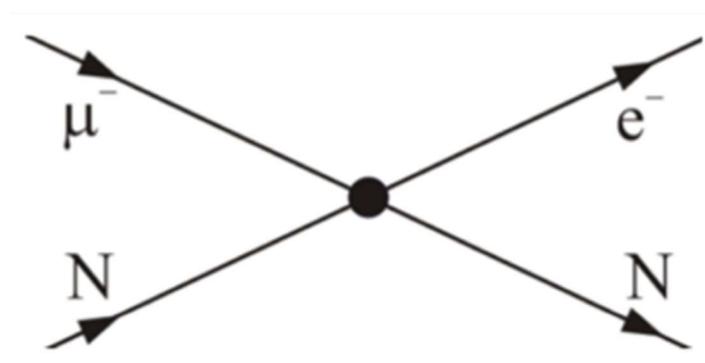


Figure 1.6: A more general diagram of flavor violating muon-to-electron conversion in the presence of a nucleus [5]. The dot in the center could represent any number of mechanisms, and Mu2e does not rely on understanding what that dot might represent.

Mu2e begins with an 8 GeV proton beam (made of the particles we can easily access, namely up and down quarks) which enters the production solenoid and hits a production target, as shown in Fig. 1.7. This collision with the production target is how we produce the muons we want to study. These muons are directed through the transport solenoid and stopped in the detector solenoid by the stopping target, which provides the nuclei for the conversion mechanism we are looking for. The tracker and calorimeter then trace and reconstruct the resulting events in the detector solenoid, allowing the muon to electron conversion to be detected if it occurs.

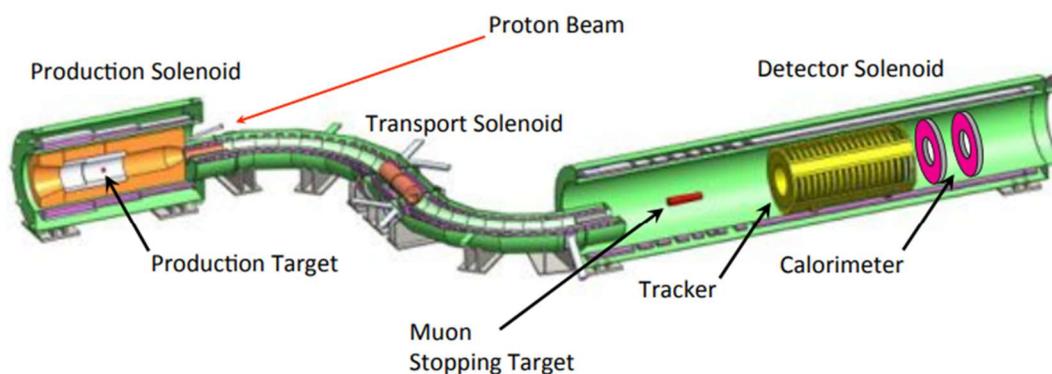


Figure 1.7: Mu2e experimental setup [6].

1.3 Production Targetry

The production target is fundamental to the success of Mu2e, as it is how we get muons in the first place, but production targetry in general is becoming increasingly important to the HEP field. Many HEP experiments, including the LHC, have made use of primary particle beams, such as proton beams, which can be produced directly via ionization. In other words, these primary beams are made up of those least massive fundamental particles which we have access to in abundance. However, HEP has largely exhausted the potential for primary beam experiments in some areas, necessitating the use of secondary particle beams, which are produced by colliding a primary beam with a target and collimating the resulting particles into a new beam, as Mu2e does.

Although the Mu2e production target will be the primary focus of this research, especially because it requires some novel target characteristics, previous HEP experiments have made use of targets to produce secondary beams, and it is worthwhile to consider historical approaches to targetry. At Fermilab, one of the first notable examples of production targetry was for the Tevatron, the high-energy collider which preceded the LHC. However, unlike the LHC, which is a proton-proton collider, the Tevatron was a proton-antiproton collider, so it required a target to produce the secondary antiproton (\bar{p}) beam. Little targetry research had been done at the time, so the only reference the \bar{p} target designers had was the CERN Antiproton Source, which had been successful with both tungsten and copper targets. The original plan was to use tungsten for the Tevatron's \bar{p} target due to its higher density and estimated secondary yield, but when tungsten was tested

under the Tevatron beam conditions, it had immediate problems due to thermal shock (i.e., it exploded), so copper was used as the target material for the Tevatron's first run. However, with increasing beam intensity, concerns about the copper target arose as a result of copper's relatively low melting point. This prompted investigation into a nickel target, which was used for the beginning of the second run. Eventually, further studies resulted in a switch from nickel to Inconel 600, a nickel-based alloy, because of its increased durability and decreased pbar yield depletion with continued use [7].

The pbar target served as a template for the production target for Muon g-2, the previously mentioned experiment at Fermilab which measured the magnetic moment of a muon to higher precision. The Muon g-2 experiment required a muon beam rather than an antiproton beam and had different proton beam parameters, but the pbar target's resilience in the Tevatron made Inconel 600 the material of choice for the Muon g-2 production target, which successfully survived the beam conditions and produced adequate muon yield for the experiment's requirements.

Upon first glance, Mu2e might seem similar enough to Muon g-2 that Inconel 600 is an obvious material choice for the Mu2e production target. However, one of the most significant differences between the two experiments is the necessary muon yield—Mu2e requires a much higher muon beam intensity than Muon g-2 did because it is searching for a rare process. The conventional wisdom in targetry has been that high atomic number (Z), high density target materials can increase production of secondary particles and result in higher intensity secondary beams. This made tungsten the initial material choice for the Mu2e production target, with its high density, high Z , and high melting point. As mentioned previously, the pbar target design team had problems with thermal shock in tungsten, but

the Mu2e team proposed an engineering fix by segmenting the cylindrical target, as shown in Fig. 1.8, to allow the material adequate room to expand without immediately having thermal shock issues. This design is referred to as the Hayman 2.

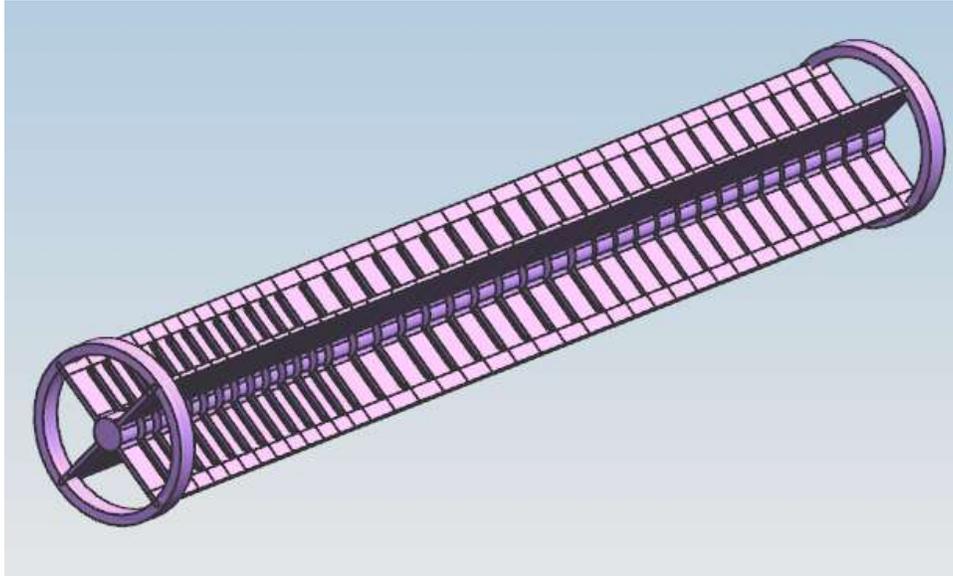


Figure 1.8: Proposed tungsten production target design for Mu2e [8]. This target is 220 mm in length with each core having a radius of 3.15 mm.

Unfortunately, recent employment of tungsten in environments subject to extreme irradiation have demonstrated its potentially detrimental susceptibility to radiation damage [9]. As a result, we have concerns about the stability of a tungsten target, especially considering the Mu2e beam parameters—subjecting tungsten to an unprecedented 8 GeV beam would likely cause an extraordinary amount of radiation damage, and this assertion has now been further substantiated by simulation [10]. In addition to instability of irradiated material due to dislocations (nuclei being displaced from the lattice) and void swelling (accumulation of dislocations in one area, causing a void), possible transmutation products of tungsten are of further concern. Rhenium, for instance, is a likely transmutation product

because of its close proximity to tungsten on the periodic table, and its melting point is about half that of tungsten. Therefore, even a small amount of rhenium in the target would be dangerous for its stability under high temperatures, and we expect the maximum temperature in the target to be significantly higher than rhenium could sustain.

Furthermore, a muon production study shown by the International Muon Collider Collaboration (IMCC) has led us to question whether an extremely high density target material like tungsten is necessary to produce adequate muon yield for Mu2e. IMCC concluded from Fig. 1.9 that tungsten was the optimal material choice for a muon production target because it has the highest yield, as expected, but we were surprised by how much more similar the muon yields for some lower density materials like nickel and molybdenum were to tungsten than we would have expected. This, in addition to our concerns about the structural integrity of a tungsten target after irradiation, has led us to consider other potential materials for the Mu2e production target.

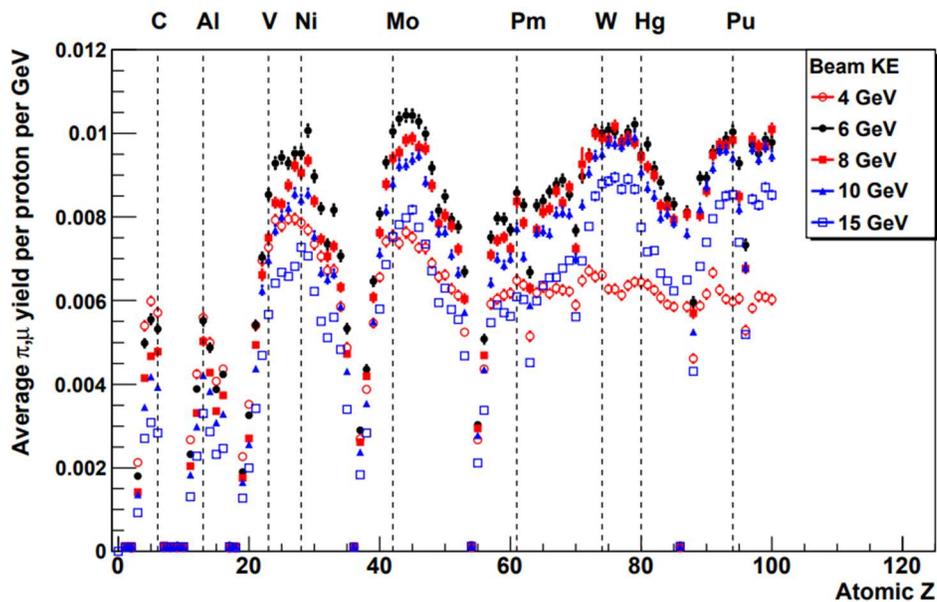


Figure 1.9: Pion and muon production study from IMCC [11].

Now that we are no longer restricted to the conventional choice of tungsten, we aim to design a framework for choosing an ideal production target material based on first principles. There has been little foundational research in targetry, as evidenced by the process of choosing materials for the pbar and Muon g-2 targets. Although we can no longer put a target directly in the beam to test its feasibility like the pbar target designers could—there are existing test beam facilities, including the Fermilab Test Beam Facility, but they operate at very different energies than the Mu2e target will experience, making them impractical for testing target materials for Mu2e—we do have a unique opportunity to capitalize on software advancements to fill the gaps in targetry knowledge as we design the Mu2e production target.

While the immediate goal for this work is designing a production target for Mu2e, it seems probable that muon production targetry will be essential to the future of HEP. The Snowmass 2021 report, which is the result of the latest planning effort of the US HEP community, slates a number of future projects which will require high intensity muon beams [12]. These include Mu2e-II and the Advanced Muon Facility, the proposed higher intensity successors to Mu2e which will utilize Fermilab's in-progress superconducting linear accelerator, PIP-II [13].

Perhaps more notably for the HEP field overall, however, is the muon collider, which is a proposed successor to the LHC. The LHC is only expected to run until 2042, and one of the biggest questions in HEP currently is what will succeed the LHC in the energy frontier, because designing and building an accelerator of even larger scale than the LHC could easily take twenty years. In the Snowmass 2021 report, the US HEP community expressed its support for a muon collider, which would allow us to probe energies far beyond what the

LHC is capable of without an unfathomably large accelerator (for reference, the LHC is on the scale of ~ 10 TeV, and a muon collider of that energy would be comparable to a proton-proton collider on the scale of ~ 100 TeV—i.e., it would have to be so big that it would be practically impossible) [12]. A muon collider, as the name suggests, would require muon beams, which means it would require a muon production target. Mu2e is requiring an unprecedented muon beam intensity, and designing the Mu2e production target is already proving difficult, so significant advances in targetry will be necessary to produce the even higher intensity muon beam required for a muon collider. The US HEP community is focusing much of its energy and accelerator frontier efforts on muon collider R&D, and targetry R&D will be crucial for a successful muon collider. Thus, the present study will not only contribute to the success of Mu2e but also to success in future accelerator projects, including a possible muon collider.

There are a few important Mu2e-specific restrictions we must consider in designing this target. For one, the target must be radiation cooled—i.e., the experiment has not been designed with any active cooling system for the target, so the target must be able to withstand extreme temperatures on its own. Additionally, the target must be corrosion resistant or stable after corrosion. The Mu2e vacuum system is limited to 10^{-5} Torr, which is not particularly high vacuum compared to many accelerator systems—the proton beamline at Fermilab, for instance, is around 10^{-10} Torr. Additionally, the tracker uses argon/CO₂, and research grade argon gas (i.e., more pure than the commercial mixture that will likely be used for Mu2e) contains around 5 ppm of O₂ [14]. Mu2e hasn't been designed with any barriers between different sections, so the entire experiment shares one vacuum system, and if any one part of the experiment leaks, the entire experiment will experience the effects.

Therefore, if any argon/CO₂ leaks from the tracker, the target will be surrounded by increased O₂, and at the extreme temperatures we anticipate in the target, any amount of O₂ is a cause for concern due to risk of corrosion. Finally, the target must fit within the confines of the remote handling system (RHS), which will be used to safely remove and replace the target as necessary. The RHS was designed for the original tungsten target design, so any new target design must be mounted on a bicycle wheel structure like the original target, shown in Fig. 1.10, and must be no longer than 22 cm.



Figure 1.10: Bicycle wheel mount for the production target [8].

To summarize, there are four main characteristics we are looking for in a target material. First, the material must be stable under extreme temperatures. While the target design can mitigate thermal shock concerns, to choose a material we must consider the melting points, expected target temperatures, and material properties (emissivity, ductility, etc.) at expected temperatures. Second, as previously explained, the material must be corrosion resistant or be stable and maintain favorable material properties after corrosion. Third, the material must be stable and maintain favorable material properties after

irradiation. Ideally, this means minimizing the amount of radiation damage the material incurs in the first place. Finally, the material should maximize muon production as much as possible while maintaining the target's structural integrity and optimizing target lifetime. Replacing the target will take, at best, around 6 weeks (and quite possibly longer), so any target replacements will result in a significant reduction in runtime.

The present study will focus on three main factors: muon production, radiation damage, and energy deposition. Radiation damage and energy deposition studies will be grouped together under material damage studies. In general, we aim to choose a target material and design to maximize muon production while minimizing material damage.

Chapter 2

Methods

2.1 Software

Five software programs were used throughout this research. The first is the Mu2e Offline, which is a simulation of the entire Mu2e experiment, including the nominal beam, targets, and solenoids. The Offline is based in Geant4, a C++ toolkit for Monte Carlo simulations of particles moving through matter [15]. The second software is G4beamline, a simplified, user-friendly version of Geant4 which doesn't require programming in C++ [16]. Both the Offline and G4beamline were used for simulating production of secondary particles, namely negative pions (particles made up of a down quark and an up antiquark) and muons. Data from the Offline was analyzed in ROOT, a program designed to interface well with C++ specifically for the purpose of particle physics data analysis [17]. Data from G4beamline was analyzed in Python. The final software is FLUKA, which, similarly to Geant4, is a package used for Monte Carlo simulations of particles' interactions with matter [18]. FLUKA was used to simulate radiation damage and energy depositions in the target.

2.2 Muon Production

Many iterations of secondary production studies will be discussed here, so I want to begin by defining some terms to distinguish them. “Stopped muons” are muons which reach the stopping target (i.e., muons which are actually used in the experiment). The number of stopped muons can only be simulated in the Offline since it is the only software which includes the entire experiment, including the stopping target. For some simulations, “total muon production” will be calculated, meaning all muons produced by the beam hitting the production target will be counted (limited by some arbitrary boundary for simulation purposes), regardless of whether they reach the stopping target. This can be done in any software which can simulate secondary production and does not depend on simulating all of Mu2e. “Protons on target,” or “POT,” refers to the number of protons which collide with the production target. Secondary production will often be given in secondaries per POT as a means of normalization (e.g., I might give the number of muons per POT instead of the total number of muons produced in a simulation). For the purpose of this research, “low momentum muons” will refer to muons which have a momentum less than 100 MeV/c. This is a rough estimate of the muons which will be within the allowance of the transport solenoid—i.e., only certain muons will actually be directed into the transport solenoid to form the beam because of how the magnetic fields work out, and low momentum muons are a reasonable proxy for the muons which will actually enter the transport solenoid and continue to the rest of the experiment.

Additionally, it should be noted that muons are produced via pion decay. When the beam protons interact with the target, they first produce pions. These pions then decay into

muons (and neutrinos, which are of interest to the extensive group of people studying neutrino physics). Negative pions are the ones which produce muons, so when I refer to pions here, I am referring specifically to them, but there are other types of pions with different combinations of up and down quarks and antiquarks. At some points, pion production and/or characteristics of produced pions will be discussed—for the purpose of the present studies, this is not significantly different from muon production, as they are directly correlated. Whether I look at pions or muons is entirely dependent on the time and position in the experiment that I want to consider, but this choice has no bearing on the conclusions I draw from the results (i.e., greater pion production and greater muon production are more or less the same—the pions are just produced slightly earlier and closer to the target than the muons).

All simulations used the Mu2e beam parameters: an 8 GeV gaussian proton beam with a standard deviation of 1 mm. In the case of the Offline simulations, the full experimental parameters are in place (i.e., everything shown in Fig. 1.7 is included in the simulation), and in all other cases, a proton beam with the above parameters was directed down the central axis of a target in a vacuum. The Offline simulations also used the original target geometry, shown in Fig. 1.8, because the simulation was built around that target. All other simulations could use varying target geometries, which will be specified.

The first set of muon production simulations were simulations of muon production with respect to material. First, G4beamline simulations of total muon production per POT were run for cylinders (3 mm radius, 160 mm length) of different materials. These G4beamline simulations were then repeated with a momentum cut of less than 100 MeV/c to estimate low momentum muon production per POT. Additionally, muon stops per POT

were simulated in the Offline for targets of different materials. Error for all of these simulations was calculated as \sqrt{N}/POT , where N is the number of muons, low momentum muons, or stopped muons, accordingly [19].

The next set of muon production simulations were simulations of muon production for various target volumes. These were done exclusively in G4beamline, with total muon production and low momentum muon production per POT being simulated for cylinders of varying length and radius for a few materials of interest (namely graphite, molybdenum, and tungsten). Again, the error was calculated as \sqrt{N}/POT .

The final set of muon production simulations examined specific properties of interest for stopped muons to more specifically characterize changes in muon stops by material. Muon stops from the Offline simulations were analyzed in ROOT to understand factors like initial position of produced pions and muons, initial momentum of produced muons, and stopping position of muons in the stopping target.

2.3 Material Damage

As with the muon production studies, there will be a number of material damage simulations discussed here, so I will begin with a few definitions. “Displacements per atom,” or “DPA,” describes the average number of times each atom in a material is displaced from the lattice during irradiation. DPA serves as a quantitative measure of radiation damage. Here, DPA will refer specifically refer to Norgett-Robinson-Torrens (NRT) DPA, which does not account for displaced nuclei recombining into the lattice (as opposed to ARC-DPA, where ARC stands for

athermal recombination corrected) [20]. DPA will generally be given in simulations as DPA per POT, and I will sometimes convert this number to DPA per year by multiplying by the approximate number of POT Mu2e expects per year ($\sim 1 \times 10^{20}$). “Energy deposition” is a measure of how much energy is transferred from the irradiating particles to the material during irradiation. Energy deposition can also serve as a measure of radiation damage, but here it will mainly be used to estimate temperatures in the target.

The first set of material damage simulations were FLUKA simulations of DPA for cylinders (3 mm radius, 160 mm length) of different materials. This included total DPA and a breakdown of DPA by particle causing the displacement, including protons, neutrons, deuterons, alpha particles, and heavy ions. The error for these simulations is given by FLUKA [18].

The other set of material damage simulations were FLUKA simulations of energy depositions for cylinders (3 mm radius, 160 mm length) of a few materials of interest. These energy depositions were then used to approximate expected temperatures using the Stefan-Boltzmann law:

$$\frac{P}{A} = \epsilon \sigma T^4.$$

The average energy deposition per POT was converted to average power using the number of POT Mu2e expects per year, and the power was then divided by the lateral surface area of the target cylinder to get P/A. This value could then be used to approximate a temperature for each material.

Chapter 3

Muon Production

3.1 Results

The muon production results from the G4beamline simulations were first plotted against Z to verify the results that IMCC presented regarding muon production by target material. This yielded the same trends that IMCC saw, with peaks around $Z=30$, $Z=47$, and $Z=77$, as shown in Fig. 3.1. They were also plotted against density, and there were diminishing returns on total muon production, low momentum muon production, and muon stops with increasing density, as shown in Figs. 3.2-3.4. The specific trend here is less important than the fact that muon production clearly did not vary linearly with density.

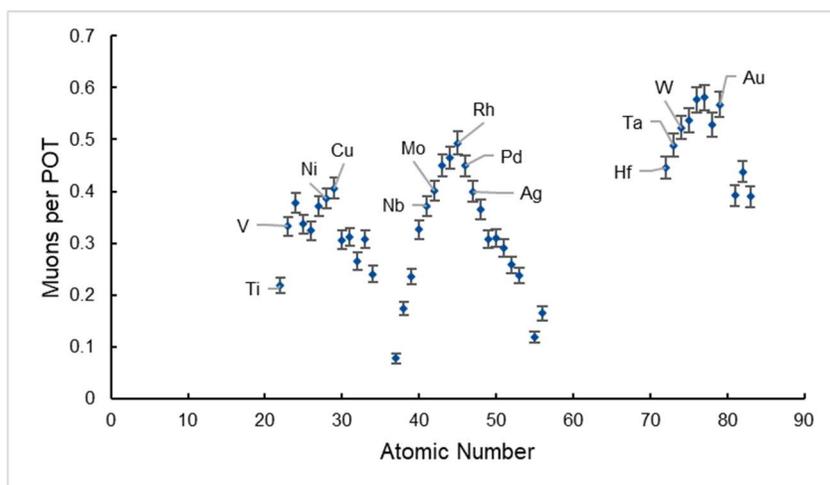


Figure 3.1: Total muon production per POT versus Z.

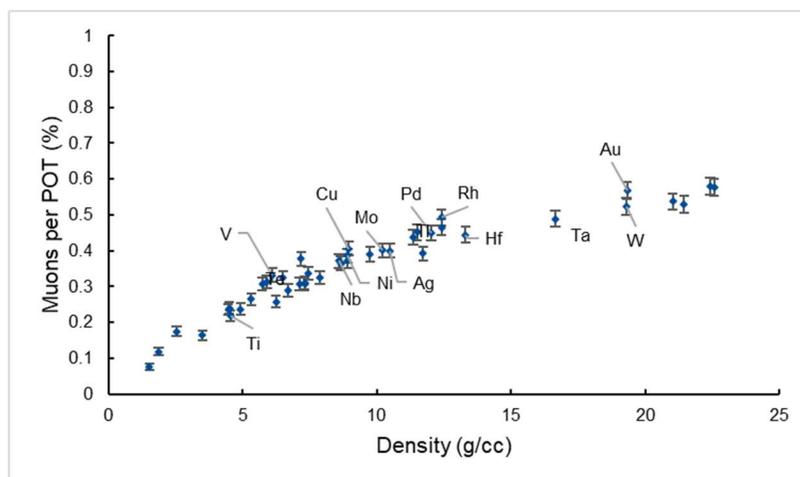


Figure 3.2: Total muon production per POT versus density.

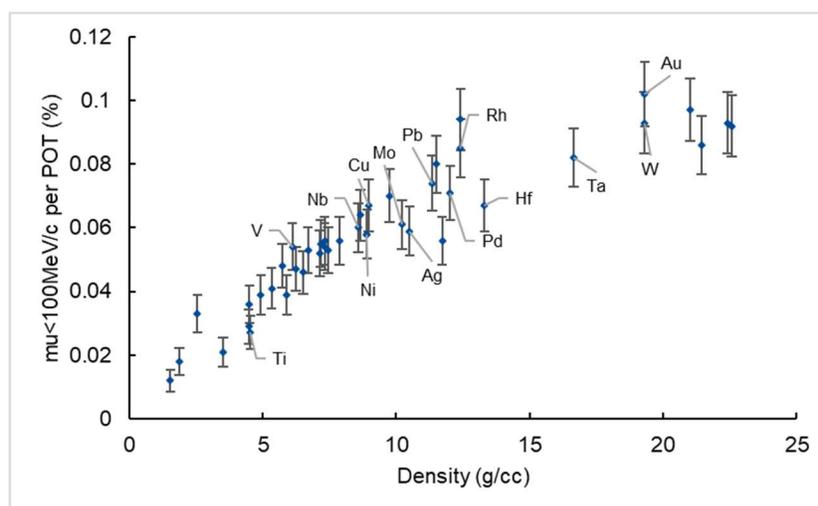


Figure 3.3: Low momentum muon production per POT versus density.

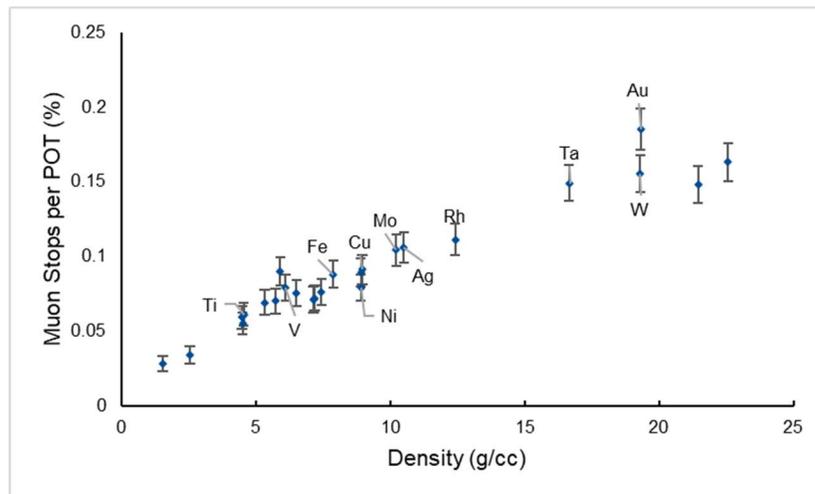


Figure 3.4: Muon stops per POT versus density.

In an attempt to find a variable for which muon production did vary linearly, the muon production results were plotted against the number of nuclear interaction lengths in each target. Nuclear interaction length is the average distance a hadron (a particle made of quarks, such as a proton) will travel through a material before interacting with the material via the strong force, which is the force which holds quarks together in hadrons [19]. The number of nuclear interaction lengths was determined by dividing the target length by the material's characteristic nuclear interaction length, so a material with a shorter nuclear interaction length has a greater number of nuclear interaction lengths for a certain sized target. Muon production varied far more linearly with number of nuclear interaction lengths, as shown in Figs. 3.5-3.7, and no significant diminishing returns on muon production were observed with increasing number of nuclear interaction lengths.

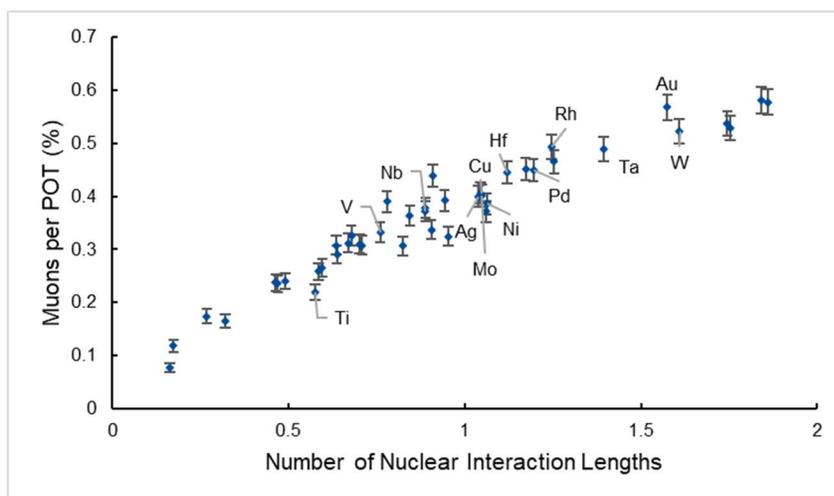


Figure 3.5: Muons per POT versus number of nuclear interaction lengths.

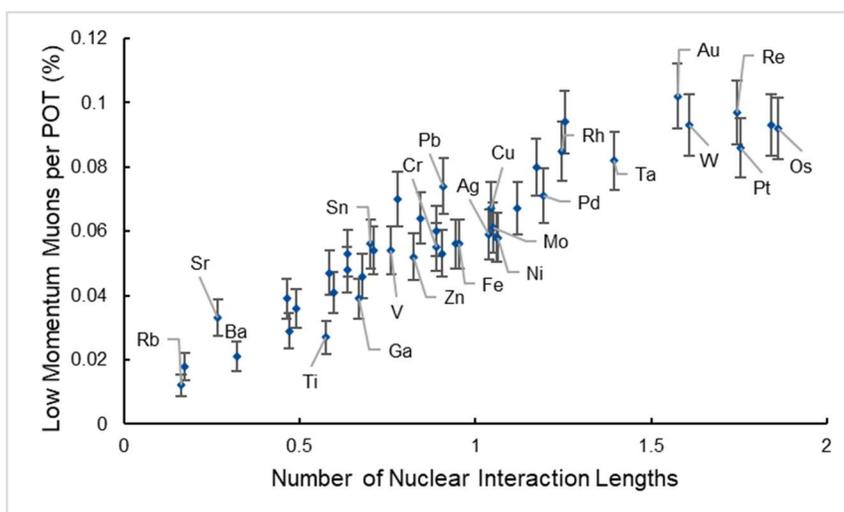


Figure 3.6: Low momentum muons per POT versus number of nuclear interaction lengths.

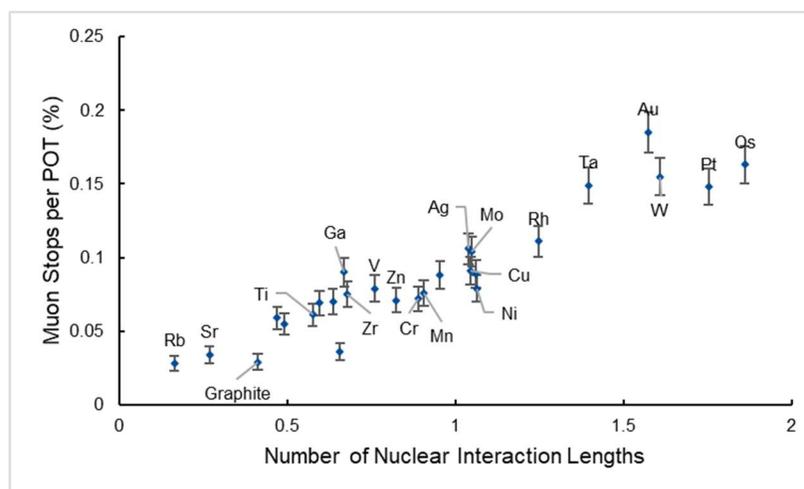


Figure 3.7: Muon stops per POT versus number of nuclear interaction lengths.

A linearly increasing trend in muon production with increasing number of nuclear interaction lengths indicated that it might be possible to make up some of the lost muon production from material by increasing the length (and possibly volume) of the target. The low momentum muon production per POT versus number of nuclear interaction lengths for graphite, molybdenum, and tungsten cylinders of 3 mm radius and lengths from 60 mm to 300 mm in 10 mm increments is shown in Fig. 3.8. The muon production per POT for graphite, molybdenum, and tungsten cylinders of different volumes (radii from 2 mm to 4 mm in 0.5 mm increments and lengths from 120 to 200 mm in 10 mm increments) is shown in Fig. 3.9. The approximate volume of the Hayman 2 target is also labeled in Fig. 3.9.

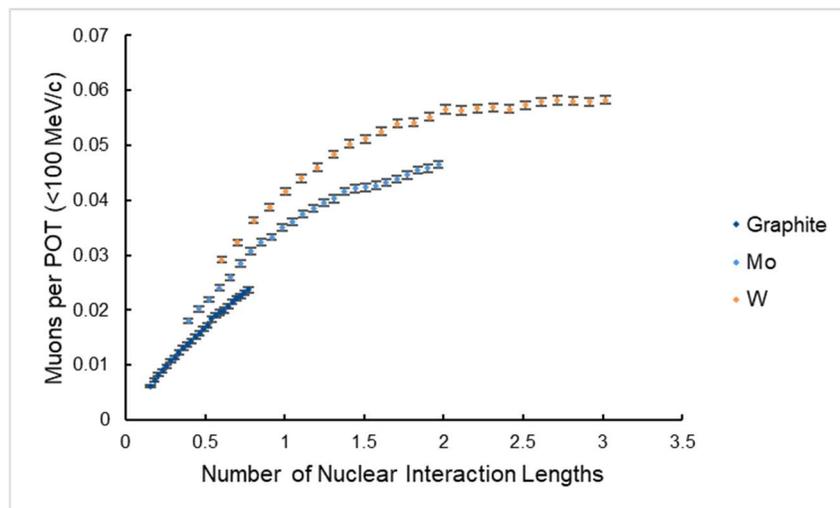


Figure 3.8: Low momentum muons per POT versus number of nuclear interaction lengths for graphite, molybdenum, and tungsten cylinders of varying length.

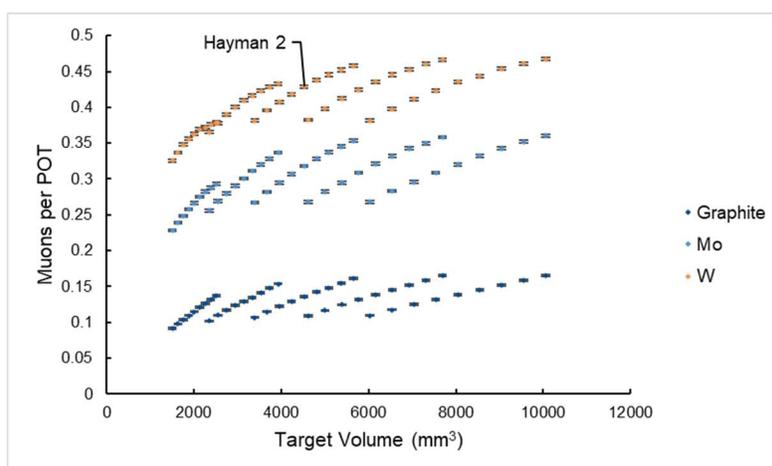


Figure 3.9: Total muon production per POT versus target volume for graphite, molybdenum, and tungsten cylinders of varying volume.

A number of plots were produced from ROOT analysis of the Offline simulations (the full list is available in Appendix C), but the most notable results are the histograms of pion production for tungsten, Inconel, molybdenum, and graphite targets, shown in Fig. 3.10. The axes here give the position in the universal Mu2e coordinates (for this purpose, we can consider them arbitrary position axes), and the color bars to the right of each plot give the scales for number of produced pions. The targets are oriented diagonal to the axes here (you can make out the segments of the Hayman 2 geometry where the higher density of pion production is), and the proton beam comes in from the top right. All of the targets had their peak pion production in the first third or so of their length, but the tungsten target's pion production at any given point in the target is the highest, followed by molybdenum, Inconel, and finally graphite.

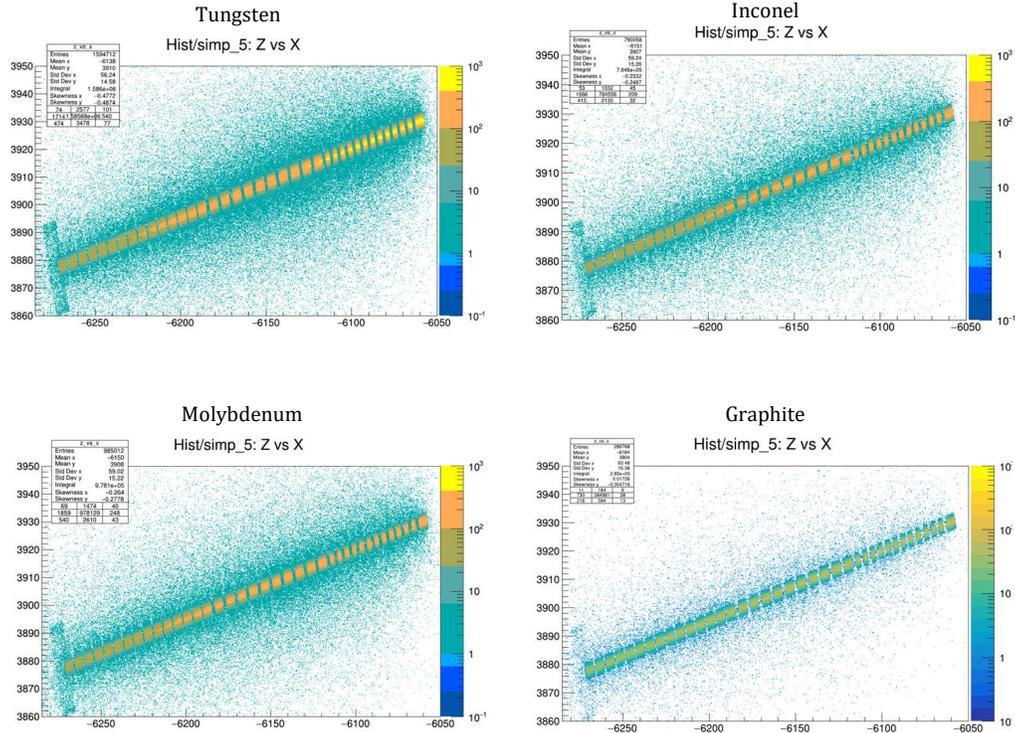


Figure 3.10: 2D histogram of position of pion production for tungsten, Inconel, molybdenum, and graphite targets with the Hayman 2 target geometry. The axes give position in the universal Mu2e coordinates, and the color bars give the scales for number of produced pions. The inset boxes give statistical information and information specific to this viewing window which are not necessary for the purpose of the present work.

3.2 Discussion

The successful reproduction of the IMCC plot, shown in Fig. 3.1, confirmed that there was not as significant of a difference in muon production with increasing Z as we expected. The plots of muon production versus density (Figs. 3.2-3.4) provide further context for this result: there were diminishing returns on muon production with increasing density. Instead, we found that muon production increased linearly with number of nuclear interaction lengths (Figs. 3.5-3.7). This can be understood if we return to the interactions which produce muons

in the first place. As mentioned previously, muons are produced from pions, which are produced from the beam protons' interactions with the target material. More specifically, pions are produced from strong force interactions between the protons and the nuclei of the target material. Nuclear interaction length is a measure of exactly that: how far a proton will go in a material before experiencing a strong force interaction. Therefore, in hindsight, it makes perfect sense that the number of nuclear interaction lengths in a target would correlate directly with the number of pions produced and, in turn, the number of muons produced.

This linear correlation with number of nuclear interaction lengths then raises the question of whether increasing the length of a target would increase muon production for a given material. In other words, is it possible for, say, a molybdenum target with the same number of nuclear interaction lengths as a tungsten target (i.e., a longer molybdenum target) to produce the same number of muons as the tungsten target? Fig. 3.8 suggests that there is an increase in muon production with increasing target length for the same material, but having the same number of nuclear interaction lengths doesn't necessarily mean the same amount of muon production between materials. Molybdenum reached around 84% of the muon production that tungsten did for the same number of nuclear interaction lengths, and graphite reached around 65%. Although it might not be possible to exactly replicate tungsten's muon production in a different material by increasing the length of the target, it is possible to increase muon production significantly by increasing the length—for comparison, a molybdenum target of the same length as a tungsten target only reached 65% of the tungsten's muon production, and a graphite target of the same length only reached 33%. This further corroborates our initial hypothesis from the IMCC plot that a lower density

target material might not result in significant muon production loss: not only are there diminishing returns on muon production with increasing density, but it is also possible to increase muon production for a lower density material by increasing the target length. Therefore, there could be an even less significant difference in muon production between the original tungsten target design and a new target made of a different material than we initially expected.

In addition to the length of the target alone, we wanted to explore whether changing the overall volume, including the radius, would result in differences in muon production. Although increasing the radius of the target provides more material for the beam protons to interact with and produce pions (i.e., there are more nuclear interaction lengths in other directions), increasing the amount of material the produced pions and/or muons have to move through could decrease their chances of getting out of the target at all. We expected this might be of particular concern for higher density materials because there would be greater hindrance to particles moving out of the target. However, we observed a relatively similar trend for graphite, molybdenum, and tungsten targets, as shown in Fig. 3.9. Each group of points in Fig. 3.9 corresponds to a different radius, and there is a clear increase in muon production with increasing radius across all three materials up to 3 mm. After that, muon production plateaus or even decreases slightly with increasing radius, which does align with our expectation. There is no clear difference between materials of different densities, though, so a target radius of around 3 mm seems optimal across a wide range of materials.

The pion production position plots from the ROOT analysis (Fig. 3.10) illustrate the changes in pion production across the length of the target—in general, pion production

decreases as the beam moves further down the target. This is expected, since the beam protons will spread out and lose energy as they continue to interact with the target. I will discuss this further when I get to the radiation damage results, but this raises the possibility of using different target materials for different parts of the target. Some newer target designs have separate core segments held together by an external support structure, so the core segments could be made of different materials. It might be necessary to use a lower Z material in the first part of the target's length (again, more on this later), but it also might be possible to use a higher Z material later in the target's length and increase overall pion production compared to a target made up of only the lower Z material.

Chapter 4

Material Damage

4.1 Results

Fig. 4.1 shows the peak DPA per year for graphite, titanium, nickel, molybdenum, and tungsten cylinders plotted against density. Although more data points and an increased number of simulated POT to decrease the error are likely needed to fully make sense of the trend here, it does seem as though DPA per year increases faster than linearly with increasing density. Fig. 4.2 shows the same data plotted against Z^2 , and Fig. 4.3 shows total DPA except for DPA caused by heavy ions plotted against Z . Again, more data points and POT. are likely needed to be more confident in any trend here, but both of these plots show a more linear trend than the plot of DPA versus density.

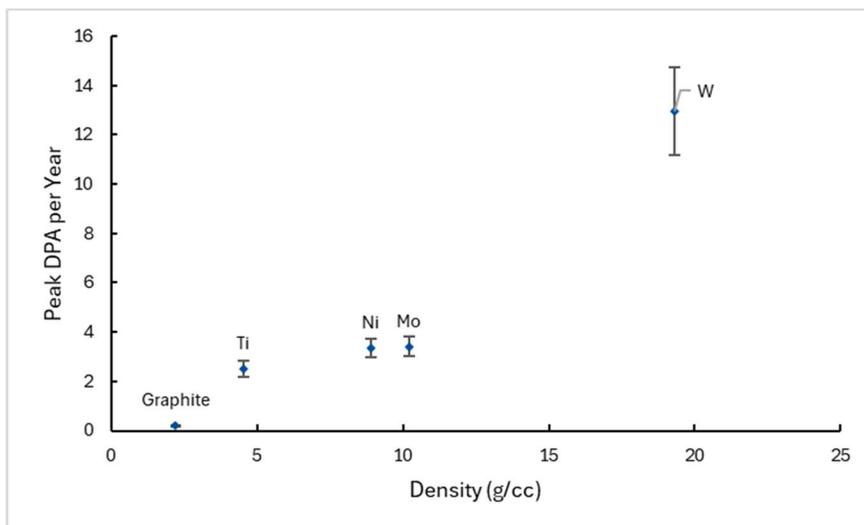


Figure 4.1: Peak DPA per year versus density.

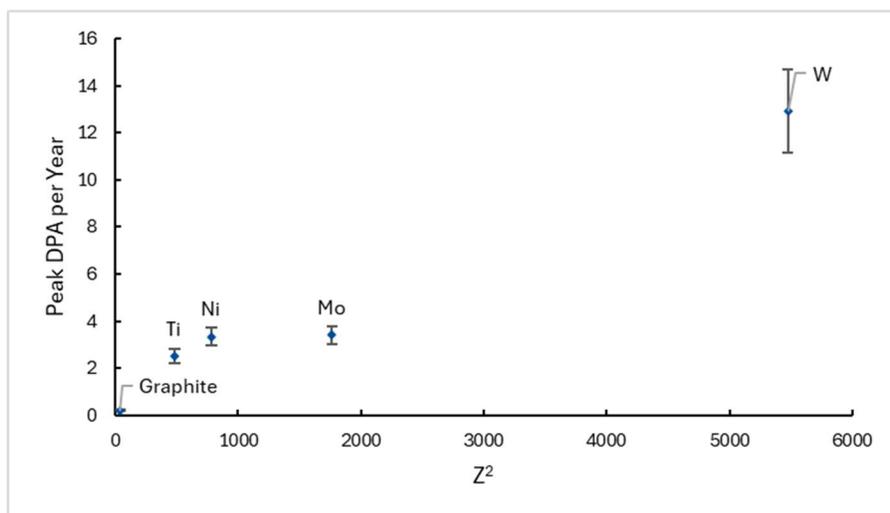


Figure 4.2: Peak DPA per year versus Z^2 .

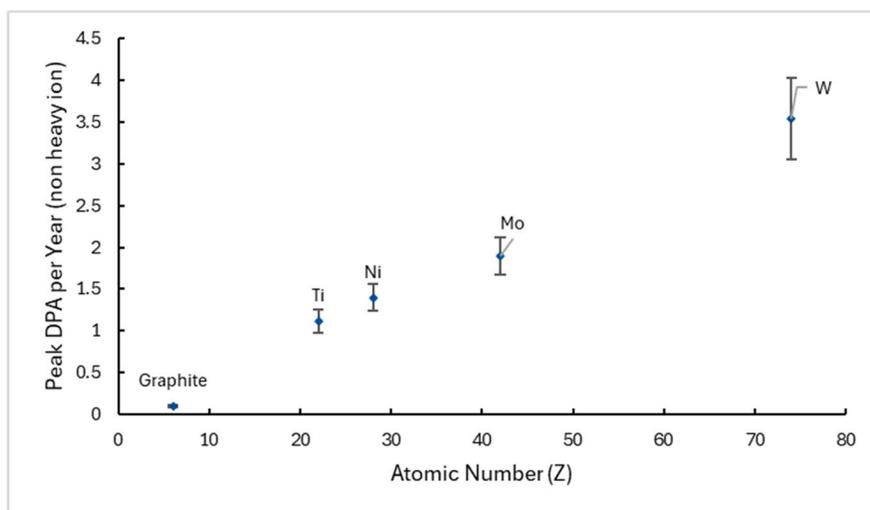


Figure 4.3: Peak DPA per year except DPA caused by heavy ions versus Z.

Fig. 4.4 shows the DPA per year across a tungsten cylinder divided into five segments. In addition to the total DPA per year, the DPA per year by particle causing the DPA is shown. Fig. 4.5 shows the same information for a graphite cylinder. In both targets, DPA is dominated by heavy ions, though the tungsten target has a greater proportion of DPA caused by heavy ions. Additionally, DPA decreases more drastically across the target's length in tungsten than in graphite.

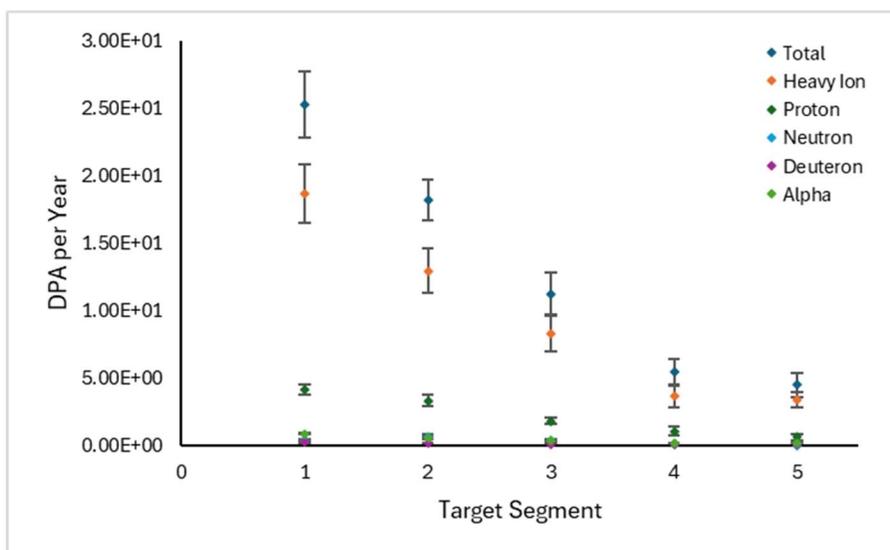


Figure 4.4: DPA per year by particle for a tungsten cylinder.

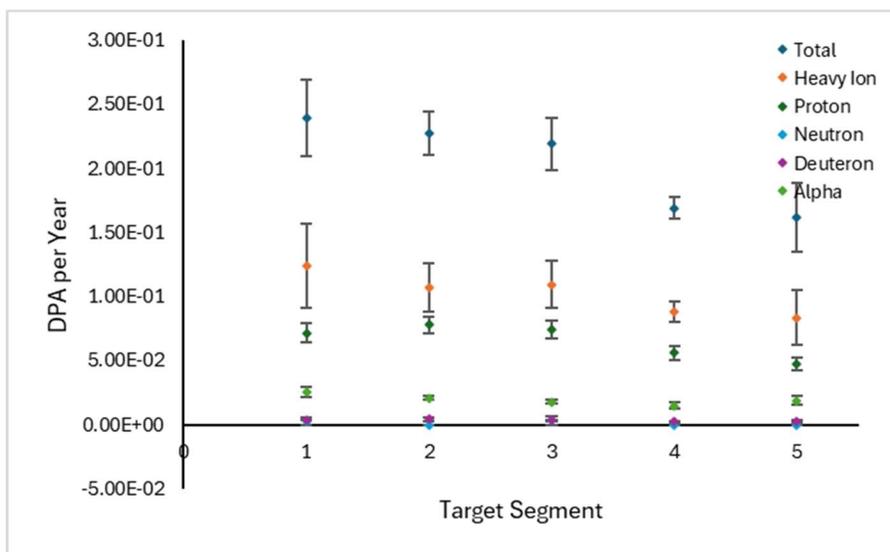


Figure 4.5: DPA per year by particle for a graphite cylinder.

The estimated temperatures for graphite, nickel, molybdenum, and tungsten cylinders are shown in Fig. 4.6 plotted against number of radiation lengths. Radiation length is a measure of how far a charged particle can travel through a material before interacting electromagnetically with the material (more specifically, it's the distance an electron can travel before its energy is decreased by a factor of $1/e$, but the specifics are not important here) [19]. As with number of nuclear interaction lengths, the number of radiation lengths

was calculated by dividing the target length by the characteristic radiation length of the material. In general, there seems to be an increase in target temperature with increasing number of radiation lengths.

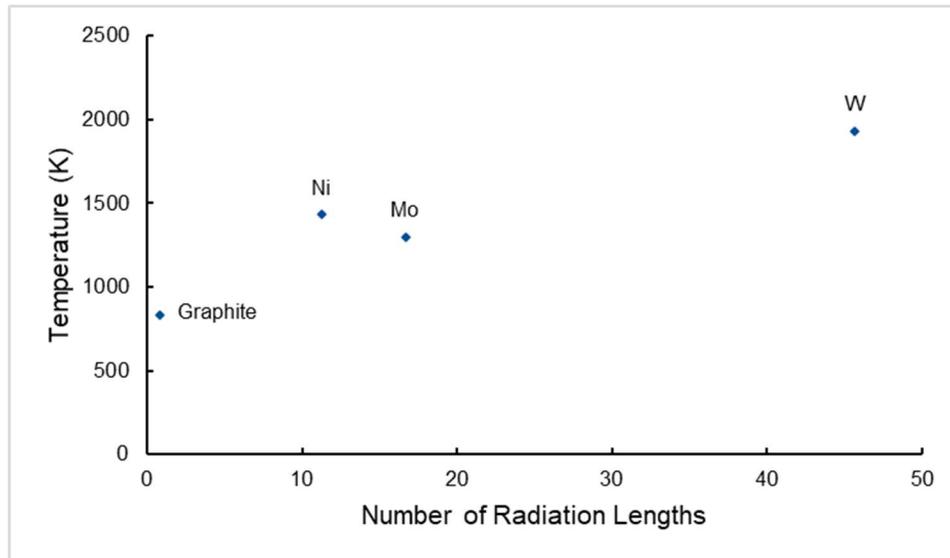


Figure 4.6. Approximate target temperature versus number of radiation lengths.

4.2 Discussion

As I mentioned previously, it would likely be clearer to draw conclusions from Figs. 4.1-4.3 with more data points and increased POT, but the current data seems to suggest that peak DPA per year varies linearly with Z^2 if we include heavy ions and linearly with Z if we discount them. The full derivation of this is beyond the scope of this work, but it is possible to derive an expression for the number of displacements over the ion fluence for a beam of a single species of ion, which comes out to

$$\frac{R_d}{NI} = \frac{\pi Z_1^2 Z_2^2 \varepsilon^4}{4E_i E_d} \frac{M_i}{M} \ln \left(\frac{\gamma E_i}{E_d} \right) \frac{\text{dpa}}{\text{ion/cm}^2} \quad [21].$$

Describing all of these variables is also beyond the scope of this work (you can see the cited textbook for a full explanation), but the necessary things to know are that Z_1 is the atomic number of the ion, Z_2 is the atomic number of the target material, and M/M_i is $1/A_2$, where A_2 is the mass number of the target material. The important takeaway here is

$$N_d \propto \frac{Z_1^2 Z_2^2}{A_2},$$

where N_d denotes the number of displacements per ion. This relationship gives us some insight as to why the DPA scales the way it does.

I'll begin with the DPA excluding DPA caused by heavy ions, which, as evidenced by Figs. 4.4 and 4.5, is predominantly caused by protons. If we assume a pure proton beam, then Z_1^2 is simply 1. Then N_d scales with Z_2^2/A_2 , and since Z roughly scales with A , N_d will roughly scale with Z_2 , which is exactly what we saw in Fig. 4.3.

The total DPA including heavy ions, which is predominantly caused by heavy ions, is a little more complicated, as it appears to scale roughly with Z_2^2 . If we assume that all of the heavy ions which cause displacements are displaced nuclei with the same Z as the original target material, then $Z_1=Z_2$, and N_d should scale with Z_2^3 . However, not all of the heavy ions will have the same Z as the original target material, so it is likely that Z_1 will actually be less than Z_2 on average, so N_d will scale with something less than Z_2^3 . Additionally, while the total DPA is dominated by heavy ions, there are DPA contributions by other particles, mainly protons, which will further decrease the average Z_1 . Therefore, it makes sense that N_d will scale with something around Z_2^2 , as shown in Fig. 4.2. Furthermore, it is worth noting that radiation length is proportional to Z^2 , so the total DPA varying approximately

linearly with Z^2 means it will also vary roughly with number of radiation lengths, so we may be able to use number of radiation lengths as a proxy for DPA [19].

The majority of the radiation damage being caused by heavy ions here (Figs. 4.4-4.5) also provides some clarity about discrepancies in radiation damage calculations between softwares. Our FLUKA DPA estimates for tungsten have consistently been much higher than other estimates, most of which have been done in other softwares, such as MCNP (another Monte Carlo particle transport code which is specifically designed for radiation) [22]. It has been suggested that MCNP does not by default include displacement cross-sections for all of the particles which FLUKA includes displacement cross-sections for, and heavy ions are notably missing from the MCNP cross-sections [23], which would mean that MCNP is not able to calculate DPA caused by heavy ions. If this is the case, the majority of our FLUKA DPA calculations being attributed to heavy ions would explain the drastic differences between our FLUKA estimates of total DPA and other calculations from other softwares.

Additionally, Figs. 4.4 and 4.5 illustrate the decrease in radiation damage across the length of the target (it's more drastic in tungsten, but still present in graphite), just as Fig. 3.10 showed a decrease in pion production across the length of the target. As I mentioned previously, new target designs give us the possibility of using different core materials throughout the target. Because radiation damage is more severe in the beginning of the length of the target, a material which incurs less radiation damage overall, such as graphite, might be more ideal for that part of the target to improve target lifetime. However, a material like graphite also produces fewer pions, so to increase overall pion production without incurring a detrimental amount of radiation damage, it might be favorable to use a material which produces more pions, like tungsten, in the later part of the target where radiation

damage is less significant. Graphite and tungsten are two extremes here and are just examples for the sake of demonstration, but there could be room here to use more than two target materials, and it's likely that neither graphite nor tungsten would be the ideal.

Like Figs. 4.1-4.3, Fig. 4.6 would likely benefit from more data points, but the general trend does make sense. Increasing the number of radiation lengths in the target should increase the number of electromagnetic interactions between the beam and the target materials, which should increase the amount of energy deposited in the target material from the beam and thus increase the target's temperature. In general, increased temperature will result in less favorable material properties like ductility and emissivity in addition to increased oxidation, both of which can result in a less stable target. Additionally, I will note once again that these are rough estimates, and a more thorough temperature analysis would be necessary to validate a target's viability. This would require a finite element analysis software as well as far more time and computing power than these estimates took, so it's not feasible to do for broad studies like this but will certainly be done for a more concrete target design before any target is manufactured.

Chapter 5

A General Framework

Perhaps the most significant takeaway from this work is that number of nuclear interaction lengths is a proxy for pion and muon production and number of radiation lengths is a proxy for material damage, including DPA and temperature. Therefore, nuclear interaction length and radiation length provide a starting point for a general framework for choosing a target material. Fig. 5.1 shows nuclear interaction length plotted against radiation length for a variety of materials. Fig. 5.2 shows the same plot on broader axes to include a few materials far from the bulk of the data points.

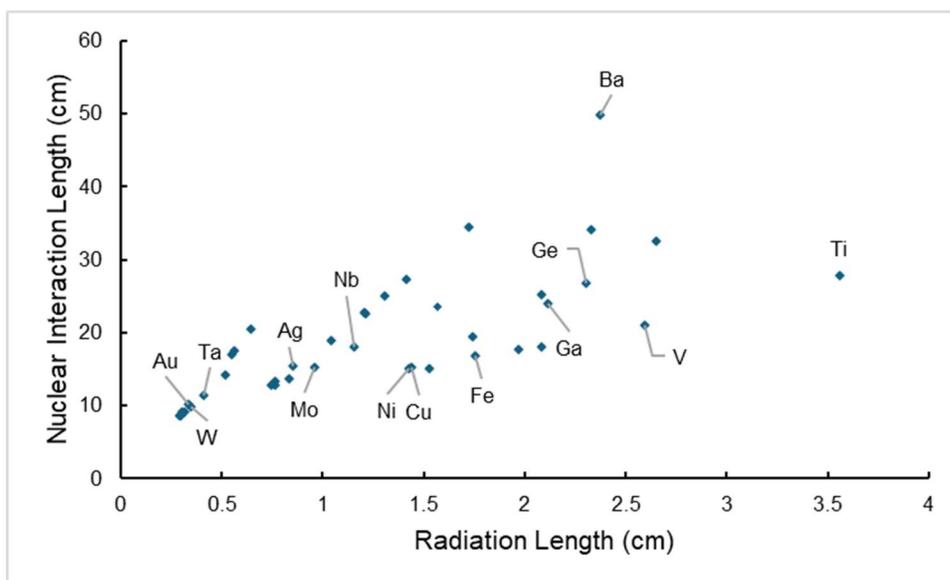


Figure 5.1: Nuclear interaction length versus radiation length for a variety of materials.

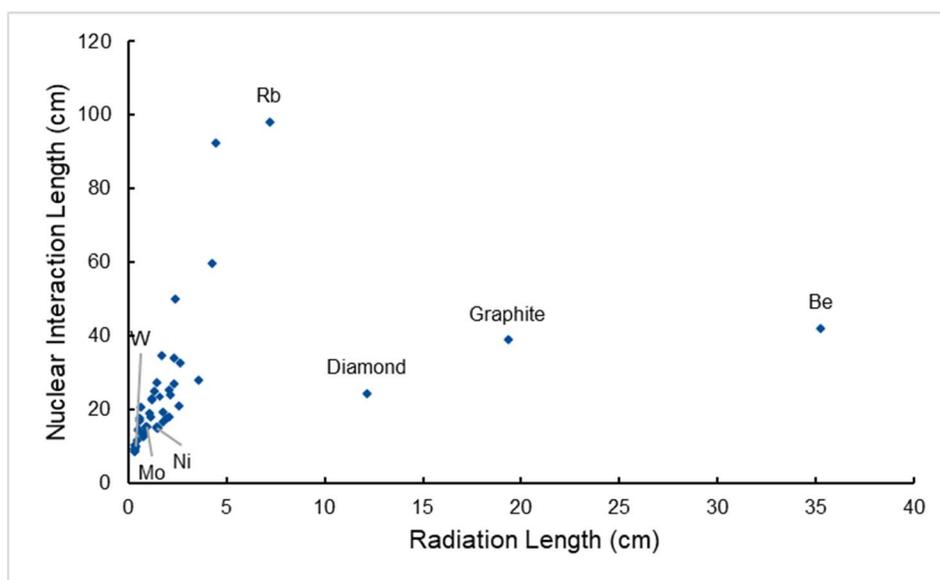


Figure 5.2: Nuclear interaction length versus radiation length for a variety of materials (on broader axes).

For Mu2e, we are trying to maximize muon production as much as possible, which means we want a shorter nuclear interaction length (to increase the number of nuclear interaction lengths in the target). Additionally, because we have an unprecedented proton beam energy and want to maximize target lifetime as much as possible, we want to minimize

any damage to the material, which means we want a longer radiation length (to decrease the number of radiation lengths in the target). In general, materials with shorter nuclear interaction lengths tend to have shorter radiation lengths as well, so Mu2e will require us to balance those two factors and find a compromise which provides adequate muon production and target lifetime.

Figs. 5.1 and 5.2 can also provide a more general framework for choosing target materials for applications other than Mu2e, and, in hindsight, they provide valuable insight about previously used materials. It makes sense that Inconel, which is primarily composed of nickel, was a successful material for the pbar and Muon g-2 targets, as nickel has a reasonably low nuclear interaction length but a longer radiation length than most other materials with similar nuclear interaction lengths. Graphite is another interesting material here (only visible in Fig. 4.2) because it has such a long radiation length compared to other materials with similar nuclear interaction lengths. Graphite is used for neutrino production targets, which are subject to extremely high energy beams, and those targets have survived very well for extended periods of time, which makes sense because graphite's long radiation length would make it less susceptible to material damage. Beryllium is also notable (also only visible in Fig. 4.2), as it has the longest radiation length of any of the materials considered here. Beryllium is often used for beam windows, which serve as barriers between different vacuum areas along the beamline while minimizing interaction with the beam. In other words, the beam is supposed to pass right through them, so it makes sense that beryllium would be a good material for that purpose because it has such a long radiation length that the beam will have little interaction with it.

Chapter 6

Conclusions

Although the conventional wisdom in targetry has been that high intensity secondary beams require high Z , high density production target materials, we have found that there are diminishing returns on pion and muon production with increasing density. Instead, number of nuclear interaction lengths is a better proxy for muon production. As a result, it is possible to make up some of the muon production lost with a material with a longer nuclear interaction length by increasing the length of the target, though increasing the radius only increases muon production up to about 3 mm. Furthermore, we have found that radiation length could be a proxy for factors which result in damage to the target material, namely DPA and temperature. Additionally, displacements are predominantly caused by heavy ions here, which could explain discrepancies between our FLUKA DPA calculations and DPA calculations from other softwares.

Overall, we have begun to develop a general framework for choosing a target material, both for Mu2e and for other applications, such as future iterations of Mu2e or a possible muon collider. Considering nuclear interaction lengths and radiation lengths is a

starting point for determining an appropriate target material, and this analysis provides valuable insight about materials which have been previously used in the path of a beamline.

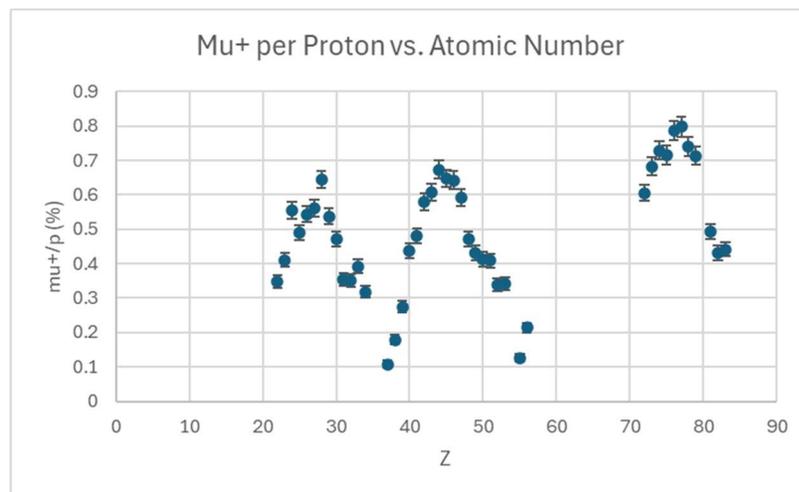
For Mu2e, our next steps will be working toward a more specific target design and running more rigorous simulations now that the material choices have been narrowed down. Inconel, for instance, is a contender for the first Mu2e target material, so future directions will include higher statistics muon production and radiation damage studies of Inconel with a more realistic target design than a simple cylinder. Additionally, energy depositions for the Inconel target design will be analyzed in a finite element analysis software to give a more accurate picture of temperature and structural parameters (stress, strain, etc.). In short, we will be performing a far more in depth analysis of one or very few specific designs and materials to validate the integrity of the target before Mu2e runs for the first time.

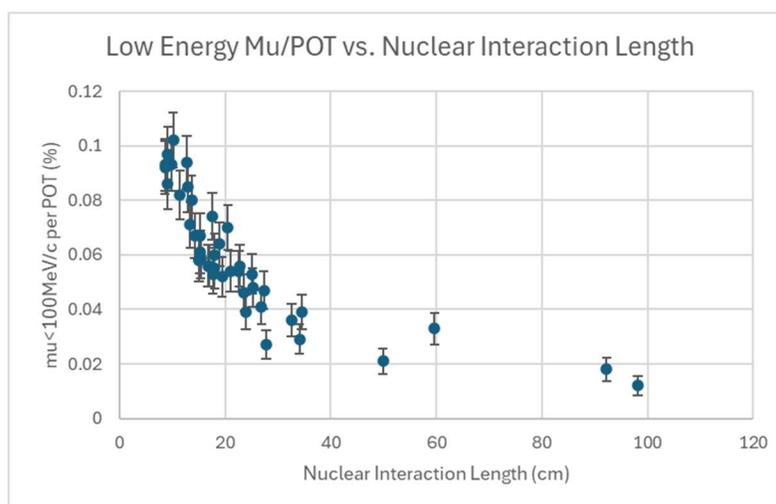
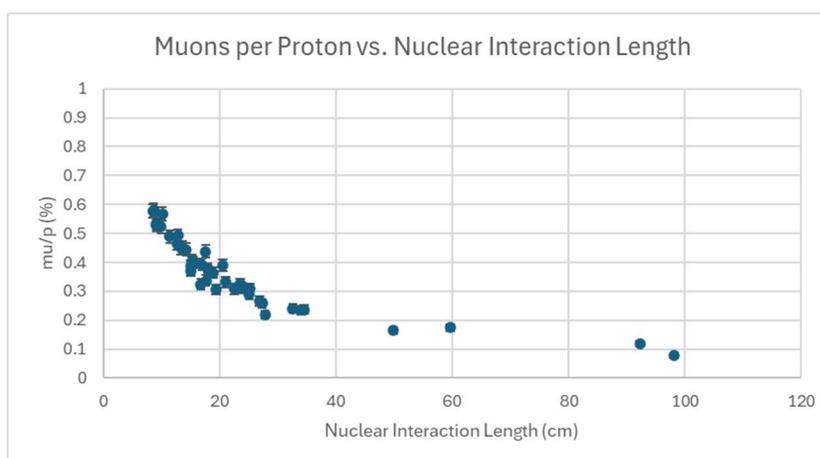
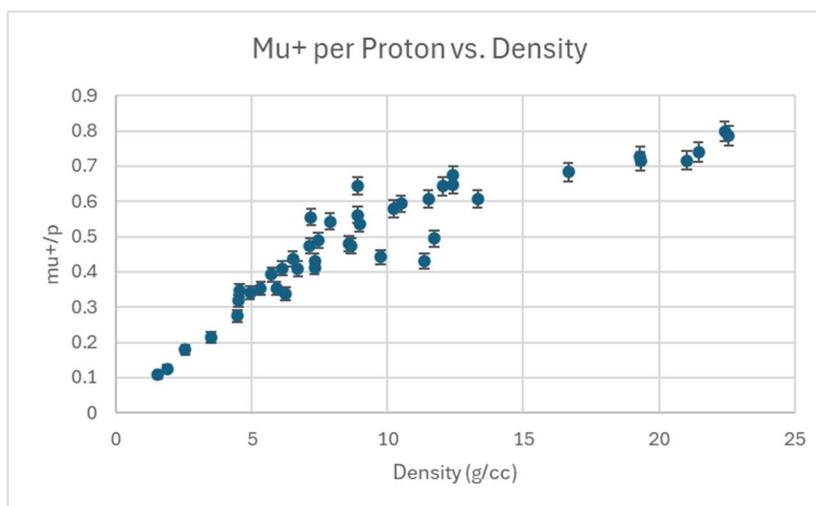
Furthermore, I want to reiterate that high power production targetry is becoming increasingly important to the HEP field as a whole, with many of the slated future projects requiring high intensity secondary beams. In addition to Mu2e, the US HEP community has stated intentions of working toward Mu2e-II, the Advanced Muon Facility, and a muon collider. None of these endeavors will be possible without significant advancements in production targetry, and we hope that the more general analysis of target materials presented in this work will provide a foundation for successful target designs in future applications.

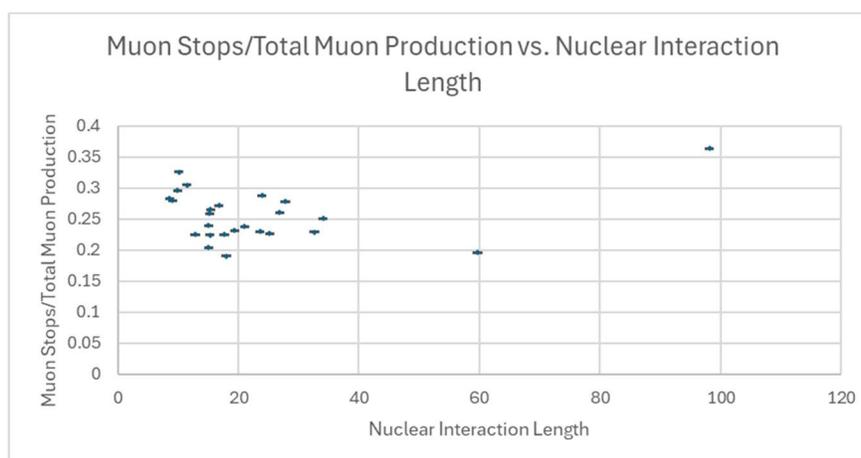
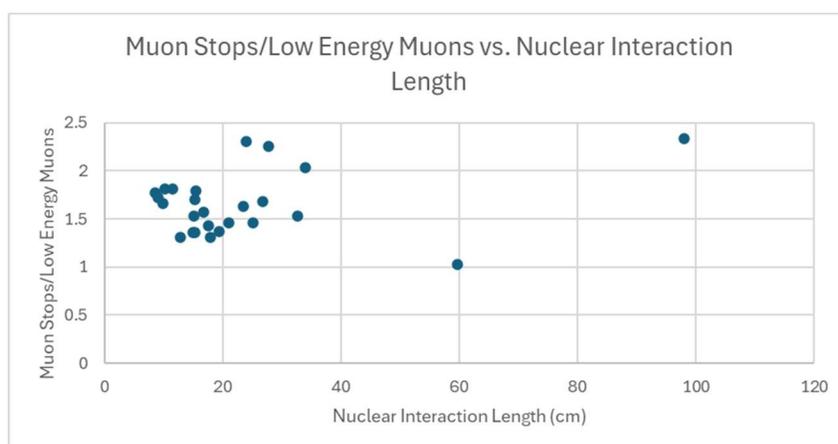
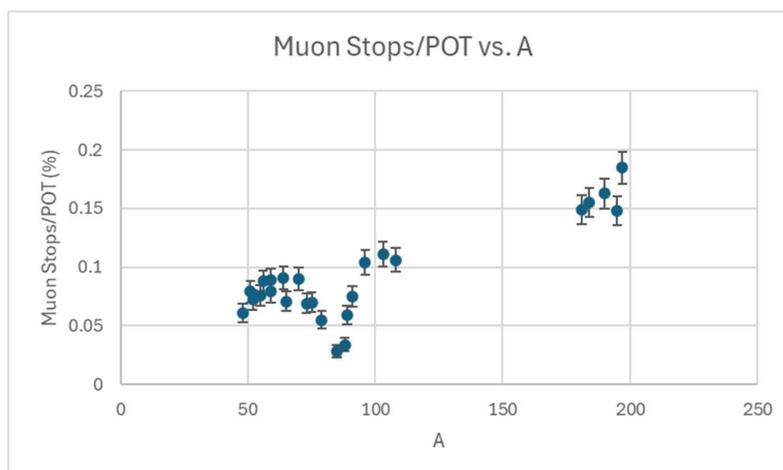
Appendix A

Muon Production

Various plots of muon production (where low energy muons and muon stops are labeled, and all other plots are of total muon production) versus atomic number, mass number (A), density, and nuclear interaction length from the G4beamline muon production simulations. Additionally, some plots include antimuon production, denoted “mu+”.



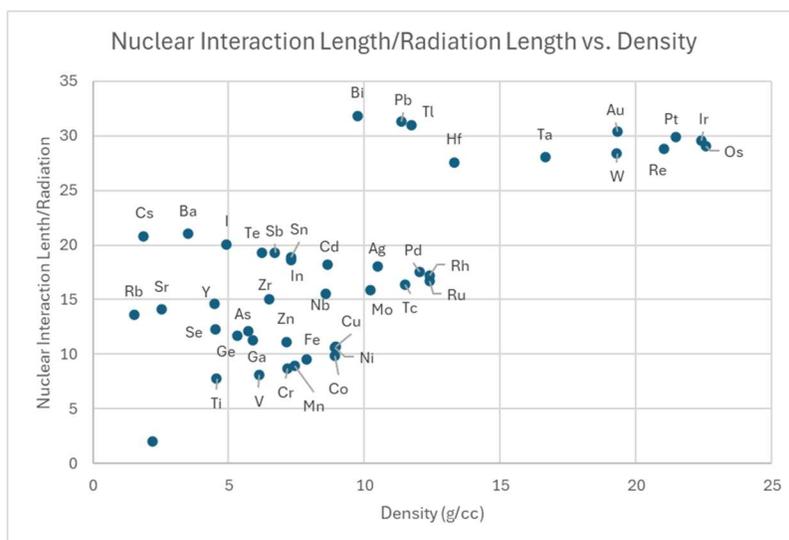


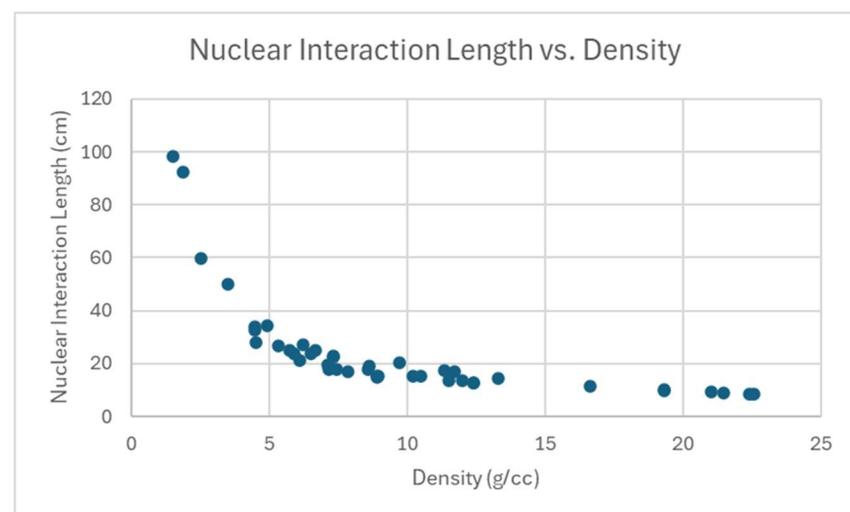
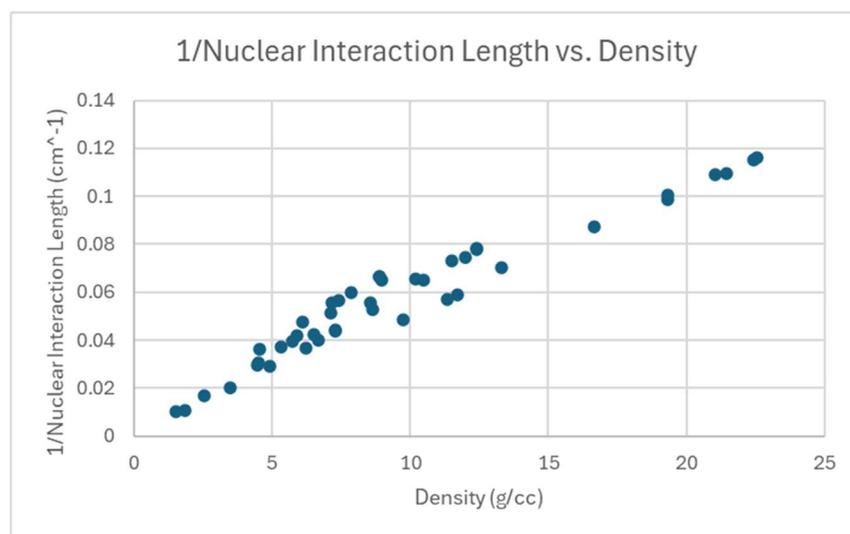
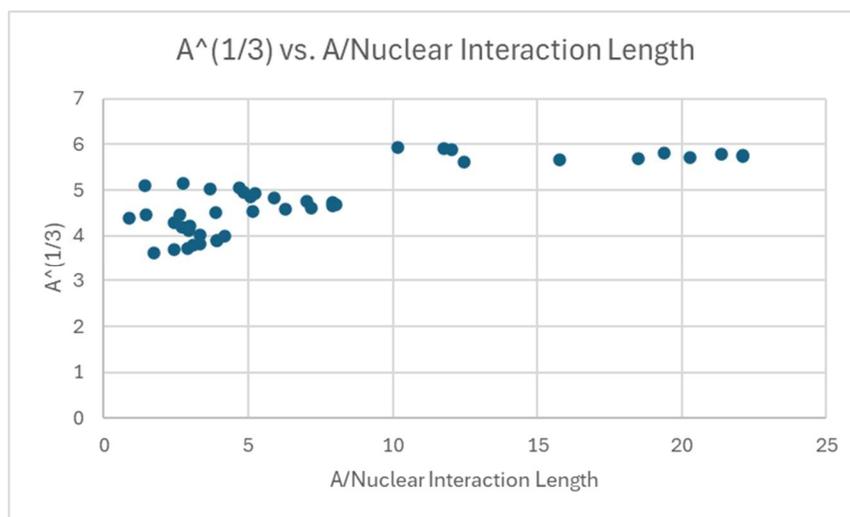


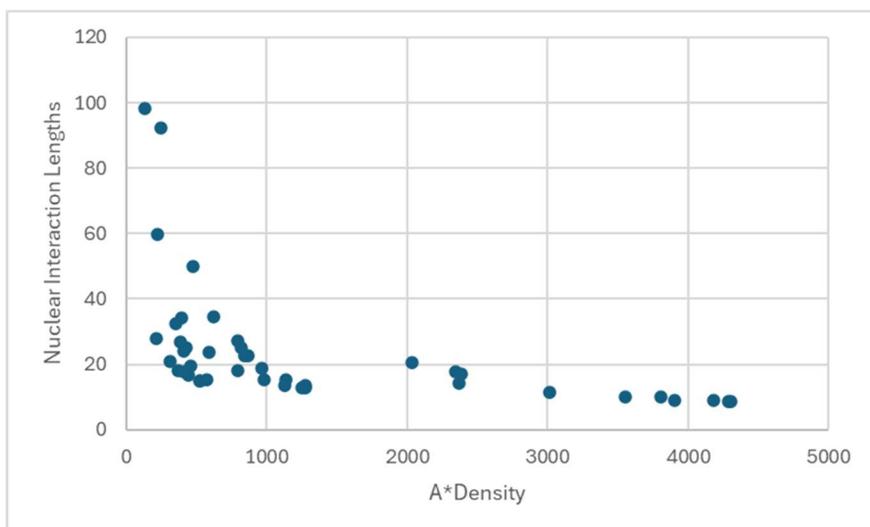
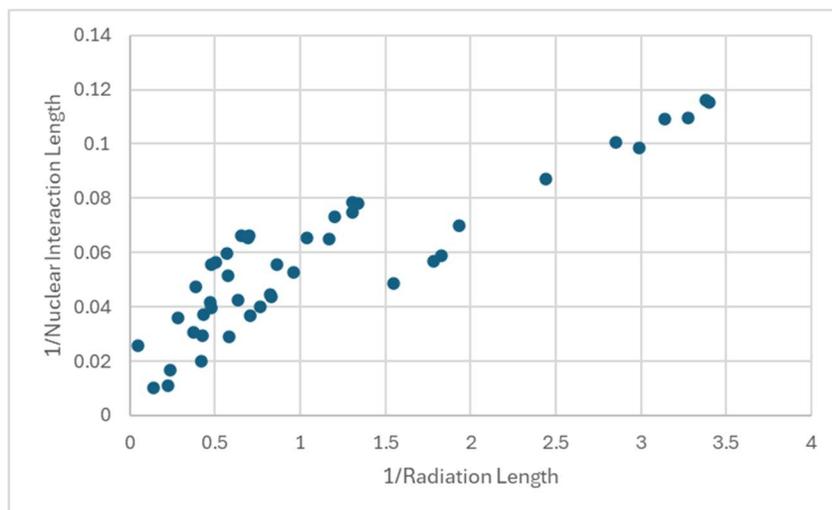
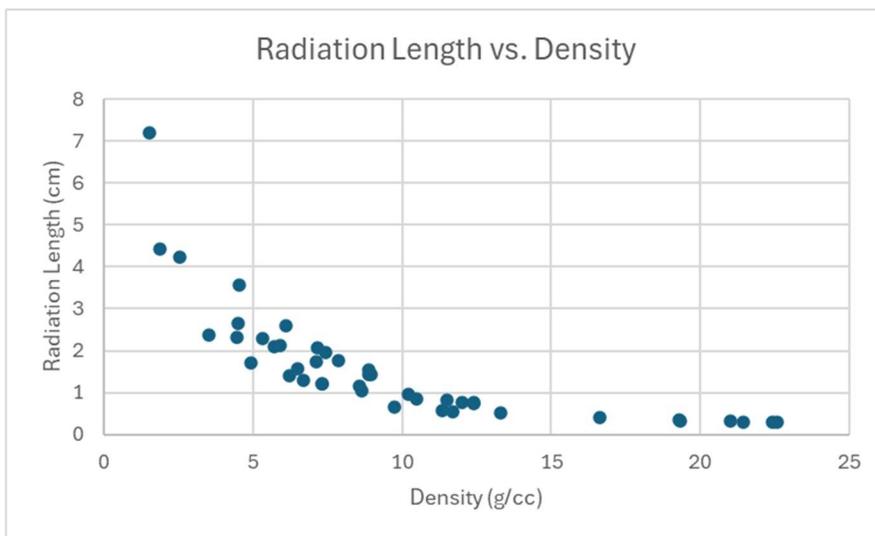
Appendix B

Material Properties

Various plots of the relationships between material properties, including nuclear interaction length, radiation length, density, and A .



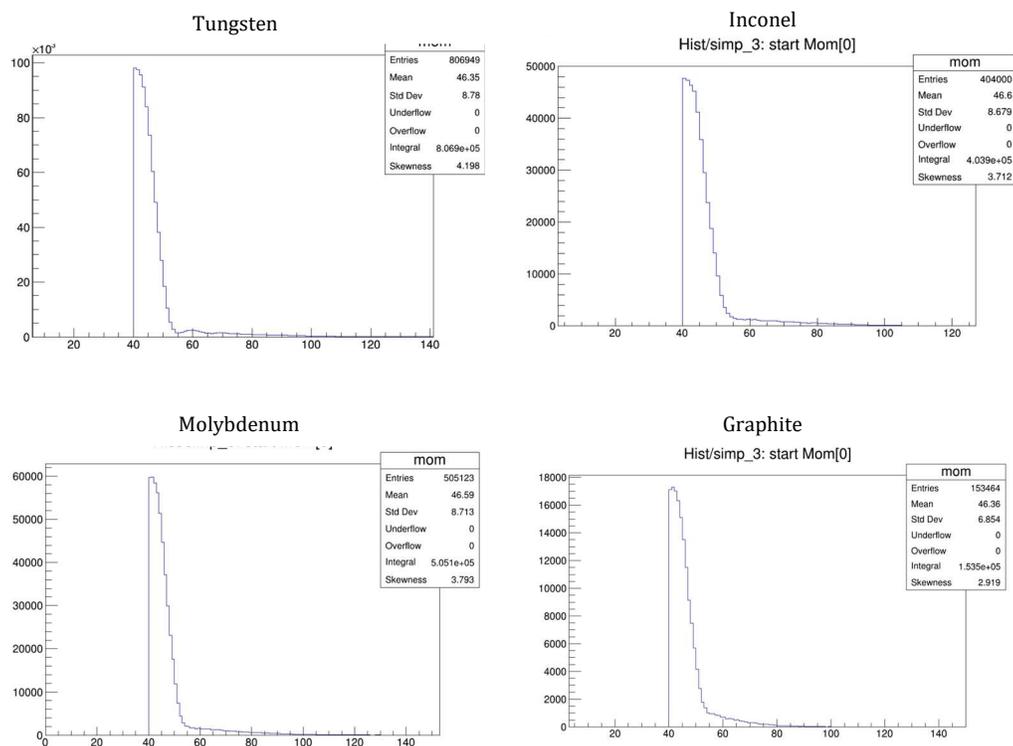




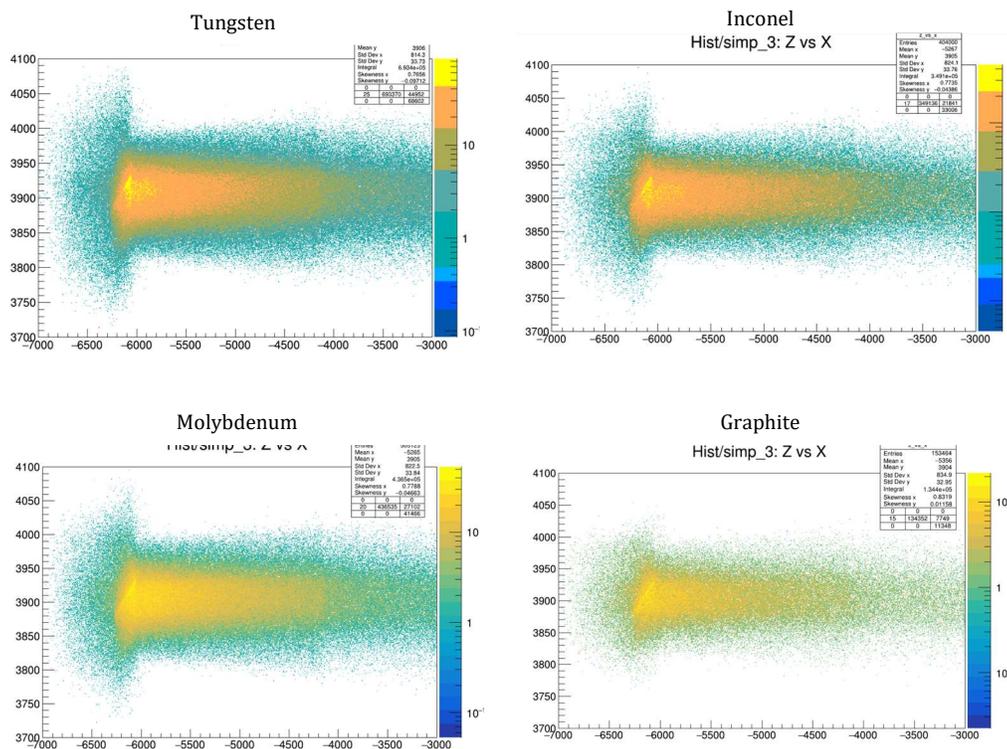
Appendix C

ROOT Analysis

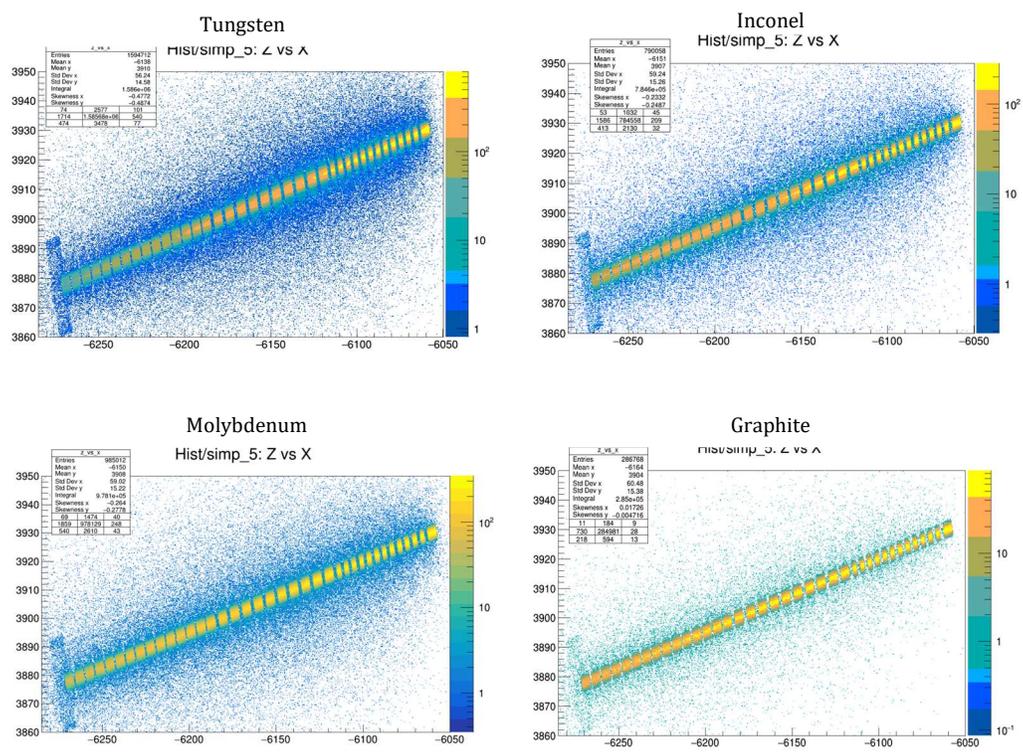
The full list of plots generated from the ROOT analysis of the Offline simulations of muon stops.



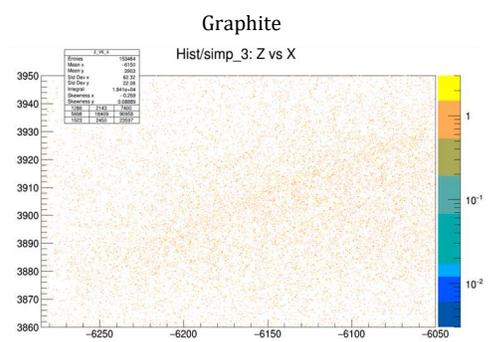
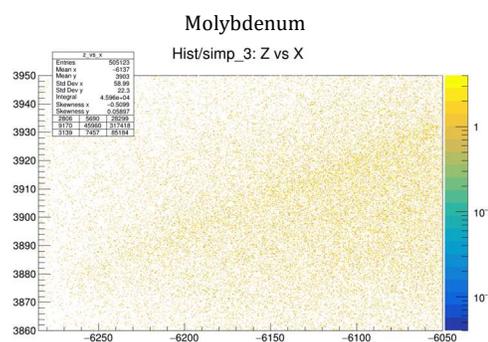
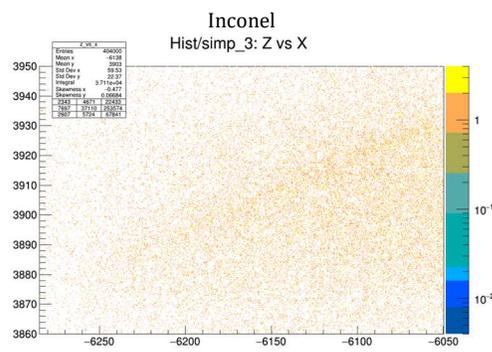
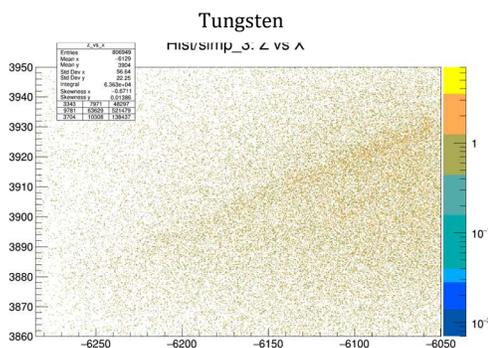
1D histogram of initial momentum of muons which stop. The x-axis shows momentum in MeV/c, and the y-axis is the number of muons with said initial momentum.



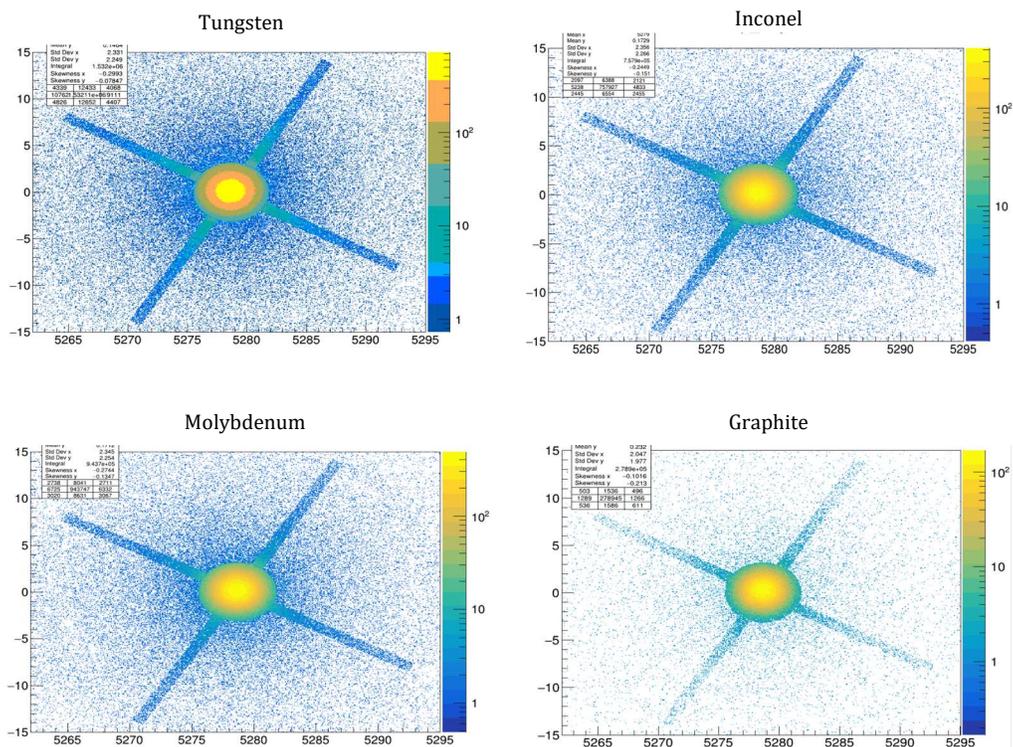
2D histogram of position of muon production in the production solenoid. The axes give position coordinates in the MuZee coordinate system, and the area shown is roughly the area of the production solenoid viewed from the top down. The color bar to the right of each plot describes the number of produced muons.



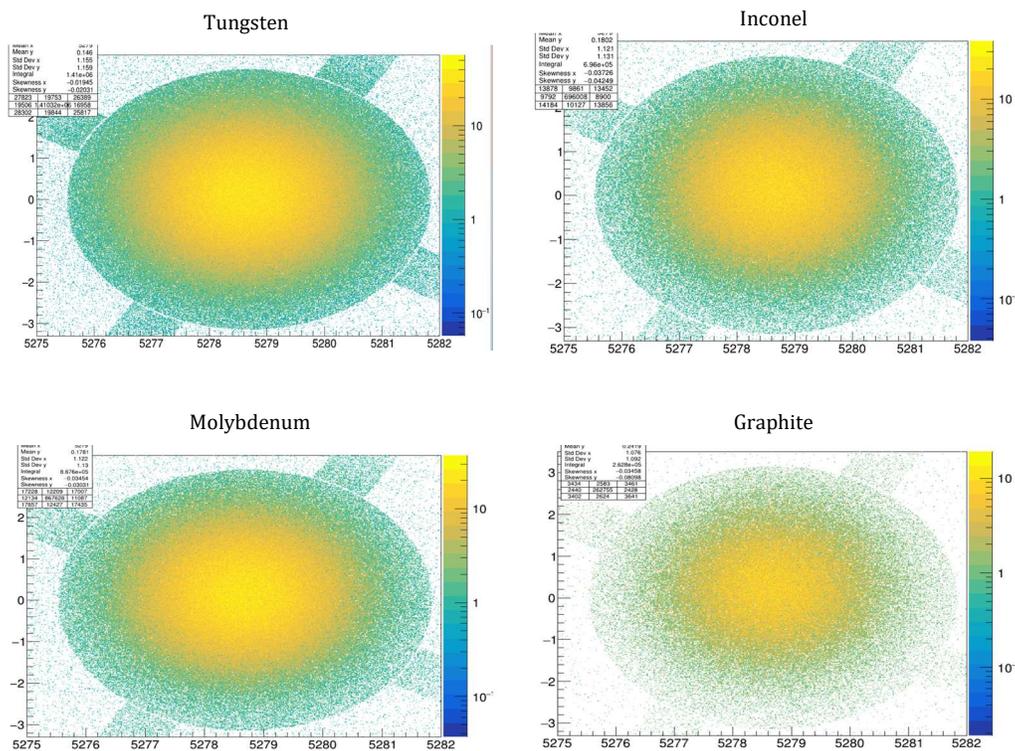
2D histogram of pion production in the production target. The axes give position coordinates in the MuZe coordinate system, and the color bar to the right of each plot describes the number of produced pions. The target is oriented diagonally across the plot (roughly where the highest concentration of produced pions is visible). The proton beam comes in from the top right corner of this plot and runs through the target toward the bottom left.



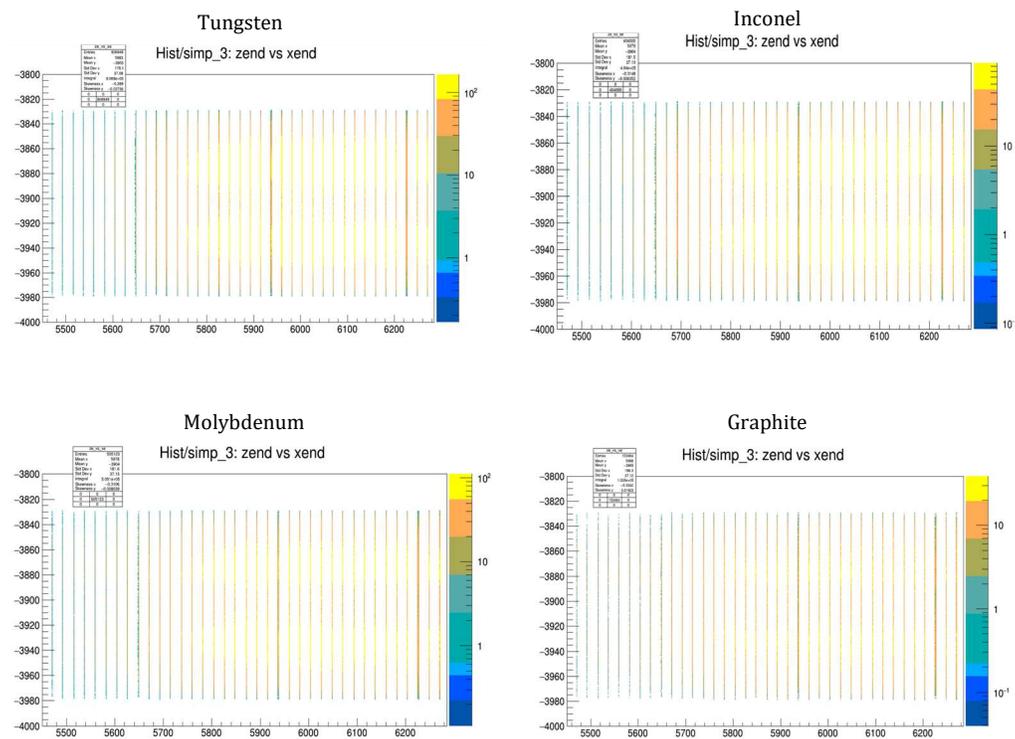
2D histogram of muon production in the production target. The axes here are identical to the axes on the histograms of pion production in the target shown above, though the target is less visible here. The color bar to the right of each plot describes the number of produced muons.



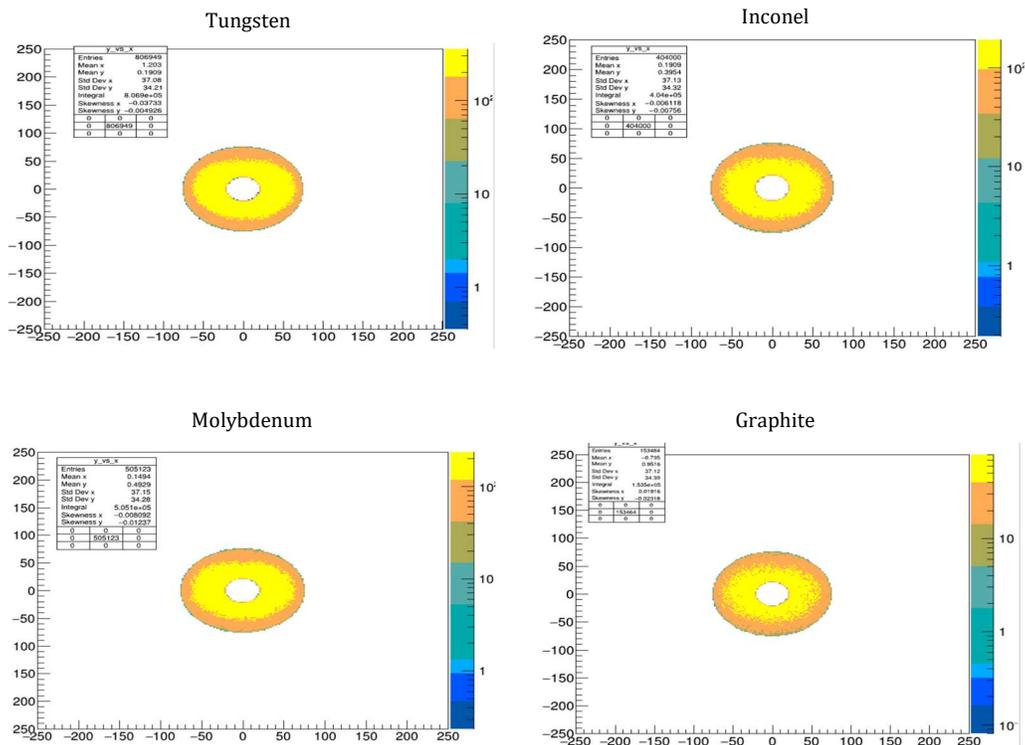
2D histogram of the radial distribution of pion production in the production target. The axes here are transformed from the Mu2e position coordinates such that we can view the target down its central axis (i.e., we have the perspective of the incoming proton beam). The color bar to the right of each plot describes the number of produced pions.



2D histogram of the radial distribution of pion production in the production. The axes are the same as they were for the above plot but zoomed in on only the target core (the circle visible in the center of these plots). The color bar to the right of each plot describes the number of produced pions.



2D histogram of the position of stopped muons in the stopping target, which consists of 34 aluminum foils. The axes here are in the Mu2e coordinates, and this is a view of the stopping target from above (each foil is visible as a line on the plot). The color bar to the right of each plot describes the number of stopped muons.



2D histogram of the radial distribution of stopped muons in the stopping target. The axes here are in the Mu2e coordinate system, this time looking down the central axis of the stopping target, (i.e., we have the perspective of a muon entering the stopping target) which consists of a series of disks with holes in the center. The color bars to the right of each plot describe the number of stopped muons.

References

- [1] *File: Standard Model of Elementary Particles.Svg*, https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg.
- [2] K. Lynch, *The Search for CLFV with the Mu2e Experiment*.
- [3] *File: Muon Decay.Svg*, https://commons.wikimedia.org/wiki/File:Muon_Decay.svg.
- [4] K. Ieki, *BSM Search with High Intensity Muon Beam in MEG II Experiment*.
- [5] M. Hedges, *The Mu2e Experiment — Searching for Charged Lepton Flavor Violation*.
- [6] L. Bartoszek et al., *Mu2e Technical Design Report*.
- [7] J. Morgan, Pbar Production Target Performance in Run II, FERMILAB-PBAR-NOTE-683, 2003.
- [8] C. N. Steffel, *Target Practice: Perfecting the Mu2e Production Target*, <https://news.fnal.gov/2019/11/target-practice-perfecting-the-mu2e-production-target>.
- [9] J. Habainy, Y. Dai, Y. Lee, and S. Iyengar, Mechanical properties of tungsten irradiated with high-energy protons and spallation neutrons, *Journal of Nuclear Materials* **514**, 189 (2019).
- [10] M. T. Hedges and M. Bloomer, *High-Power Targetry for Muon Production*, in *Proceedings of the 25th International Workshop on Neutrinos from Accelerators (NuFact 2024)* (Argonne National Laboratory, 2025).
- [11] C. Densham, *Prospects and Considerations for a Fluidized Tungsten Target for the Muon Collider*.
- [12] J. N. Butler et al., Report of the 2021 U.S. Community Study on the Future of Particle Physics (Snowmass 2021) Summary Chapter, 2023.
- [13] M. Artuso, R. H. Bernstein, and A. A. Petrov, Report of the Frontier for Rare Processes and Precision Measurements, 2022.

- [14] J. Habainy, Characterisation of Mechanical and Thermal Properties of Tungsten for High Power Spallation Target Applications, Lund University, 2018.
- [15] *Geant4*, <https://geant4.web.cern.ch/>.
- [16] *G4beamline*, <https://www.muonsinc.com/Website1/G4beamline>.
- [17] *ROOT: Analyzing Petabytes of Data, Scientifically.*, <https://root.cern/>.
- [18] *CERN FLUKA*, <https://fluka.cern/>.
- [19] Particle Data Group, Review of Particle Physics, Progress of Theoretical and Experimental Physics **2020**, 083C01 (2020).
- [20] S. Chen, Effective primary radiation damage model parameters for neutron irradiation experiments, Journal of Nuclear Materials **568**, 153883 (2022).
- [21] G. S. Was, *Fundamentals of Radiation Materials Science: Metals and Alloys* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2007).
- [22] *The MCNP® Code*, <https://mcnp.lanl.gov/>.
- [23] Y. Iwamoto, Y. Celik, F. Cerruti, R. Froeschl, T. Lorenzon, N. Mokhov, F. S. Pujol, V. Vlachoudis, and L. Yao, *Intercomparison of Radiation Damage Calculations in Target Materials at Proton Accelerator Facilities Using Various Monte Carlo Particle Transport Codes*, in *Proceedings of 15th Workshop on Shielding Aspects of Accelerators, Targets, and Irradiation Facilities (SATIF-15)* (2022).