

A Simulation-Driven Approach to Optimize and Measure the Number of Stopped Muons in the COMET Experiment

Andreas Jansen

PhD Thesis

First referee

Prof. Dr. Arno Straessner

Second referee

Prof. Dr. Satoshi Mihara

Submitted on: 31st May 2024

Date of defence: 21st November 2024

Kurzfassung

COMET ist ein zukunftsweisendes Experiment der neuen Generation. Es wird nach dem kohärenten neutrinolosen Übergang eines Myons zu einem Elektron im Coulombfeld eines Aluminiumatomkerns suchen. Dieser Prozess eröffnet herausragende Möglichkeiten zur Erforschung und Entdeckung neuer Physik, da die kohärente Myonen-zu-Elektronen Umwandlung die geladene Leptonenzahlerhaltung verletzt und im Standardmodell der Teilchenphysik bis weit unterhalb messbarer Beiträge unterdrückt ist. Das Experiment befindet sich derzeit in der Aufbauphase. In Phase-I wird angestrebt, das bestehende Limit auf das Verzweigungsverhältnis von 7×10^{-13} um zwei Größenordnungen zu verbessern, wobei eine zusätzliche Verbesserung um mindestens zwei weitere Größenordnungen für Phase-II erwartet wird.

Zur Realisierung dieser beeindruckenden Verbesserung benötigt das COMET-Experiment einen äußerst intensiven, niederenergetischen Myonenstrahl. Über die geplante Laufzeit des Experiments von 146 Tagen erreichen etwa 1.6×10^{17} Myonen den Detektorbereich. Allerdings kann nur ein Bruchteil dieser Myonen genutzt werden, da diese zuerst gestoppt und an einen Atomkern gebunden werden müssen, bevor eine kohärente Umwandlung in Elektronen möglich ist. Diese Arbeit widmet sich der detaillierten Untersuchung der Komponente, die für diese Aufgabe verantwortlich ist: dem sogenannten Myonen-Stop Target. Mithilfe modernster Simulationstechnik werden die Parameter des Myonen-Stop Targets analysiert und optimiert, um die Anzahl der gestoppten Myonen zu maximieren, während gleichzeitig eine hohe Akzeptanz für die Signal-Elektronen gewährleistet wird. Darüber hinaus wird das Vorhaben untersucht, die Anzahl der gestoppten Myonen durch die Messung von charakteristischen myonischen Röntgenstrahlen mittels eines hochreinen Germaniumdetektors zu bestimmen.

Es wird gezeigt, dass die ‘Single Event Sensitivity’ des Experiments um einen Faktor zwei auf bis zu 1.5×10^{-15} verbessert werden kann, wenn ein Myonen-Stop Target verwendet wird, das aus 66 Aluminiumscheiben mit einer Dicke von 0.1 mm besteht. Es wird eine spezielle Aufhängestruktur zur Befestigung dieser Scheiben entworfen und mithilfe eines entwickelten Prototyps umfangreich getestet. Installations- und Stabilitätstests zeigen eine Positionsgenauigkeit von ± 1 mm in x - und ± 5 mm in y - und z -Richtung. Innerhalb dieser Genauigkeit wird keine negative Auswirkung auf die Anzahl der gestoppten Myonen erwartet. Die Aufhängestruktur wurde auf Stabilität optimiert, wobei der Verlust an Signalakzeptanz durch Interferenz mit Signalelektronen auf 3.4 % begrenzt werden konnte. Die Umsetzbarkeit des myonischen Röntgendetektors wird demonstriert. Die Platzierung des Detektors in einem Abstand von 6 m und einem Winkel von 8.5° relativ zum Myonen-Stop Target, zusammen mit einer eigens entwickelten Detektorabschirmung, die ein 2 m langes Kollimatorrohr mit einem Durchmesser von 50 mm enthält, ist ausreichend, um die Untergrundrate im Germaniumdetektor auf ein arbeitsfähiges Niveau zu senken.

Abstract

COMET is a next-generation experiment. It will search for the coherent neutrinoless transition of a muon to an electron in the Coulomb field of an atomic aluminium nucleus. This process provides a golden channel for probing and discovering new physics since coherent muon-to-electron conversion violates charged lepton flavor conservation and is suppressed below measurable contributions in the Standard Model of particle physics. During Phase-I, the experiment currently under construction seeks to improve the existing branching ratio limit of 7×10^{-13} by two orders of magnitude, with an additional improvement by at least two orders of magnitude anticipated for Phase-II.

To realize such remarkable improvements, the COMET experiment requires an incredibly intense, low-energy muon beam. Over a planned runtime of 146 days, approximately 1.6×10^{17} muons will reach the detector area. However, only a fraction of these muons are usable, as they must first be stopped and bound to a nucleus to undergo conversion. Within this thesis, the component responsible for this task is under detailed scrutiny: the muon-stopping target. A simulation-based study evaluates and optimizes the muon-stopping target's parameters to maximize the number of stopped muons while maintaining a high acceptance for signal conversion electrons. In addition, the scheme to monitor the number of stopped muons by measuring muonic X-rays with a high-purity germanium detector is investigated.

A twofold improvement in single event sensitivity to 1.5×10^{-15} is shown to be achievable by employing 66 target-disks with a thickness of 0.1 mm. A dedicated suspension structure to hold these target-disks is designed, prototyped, and tested. Installation and stability tests demonstrate a positional accuracy of ± 1 mm in x and ± 5 mm in y and z directions. Within this accuracy, no negative impact on the number of muons stopped is foreseen. The suspension structure is optimized for stability while limiting the decrease in signal acceptance from interference with signal electrons to 3.4 %. The feasibility of the muonic X-ray detector is demonstrated. Placing the detector 6 m away at an 8.5° angle relative to the muon-stopping target, along with dedicated shielding and a 2 m long, 50 mm diameter tube collimator, is sufficient to achieve workable background rates for the high-purity germanium detector.

Contents

Kurzfassung	ii
Abstract	iii
List of Tables	1
List of Figures	2
1. Introduction to Muons	5
1.1. Muons in the Standard Model	5
1.1.1. The Free Muon Decay and Flavour Conservation	6
1.1.2. Muonic Atoms	6
1.2. Muons in Beyond Standard Model Physics	9
1.2.1. Motivation behind (Charged) Lepton Flavour Violation	9
2. The COMET Experiment	14
2.1. A Staged Approach: COMET Phase-I and Phase-II	15
2.2. The μ -e Conversion Signal	16
2.2.1. Detecting Conversion Electrons: A Cylindrical Detector System .	17
2.2.2. The Single Event Sensitivity	20
2.3. Understanding and Suppressing Background	21
2.3.1. Primary Beam: Proton Target, Energy, and Timing Structure . . .	22
2.3.2. Particle Transport: Particle Selection with Bent Solenoids	24
2.3.3. Muon-Stopping Target: Shape and Material	26
3. The COMET Simulation: ICEDUST	30
3.1. Simulating the Primary Beam	30
3.2. Direct Simulation of Conversion Electrons and X-rays	33
3.3. Details of the Geometry and Magnetic Field	34
3.4. Handling of Physics in ICEDUST	36

3.5. General Optimization Strategy	38
4. The Muon-Stopping Target	40
4.1. Introduction: Status Quo and Motivation	40
4.2. Methodology	42
4.2.1. Interpreting Plots: The Centered Coordinate System	44
4.2.2. Investigating the Number of Muons Stopped	44
4.2.3. Investigating the Signal Acceptance	47
4.2.4. Investigating Detector Background Rates	52
4.2.5. Treatment of Uncertainties	56
4.3. Performance Overview: Default Configuration	58
4.3.1. Additional Information: Single Layer Occupancy	59
4.4. Starting Point: The Tube-like Disk Structure	61
4.4.1. Matching the Target Shape to the Beam Profile	61
4.4.2. Using Multiple Separate Target-Disks Forming a Tube	63
4.5. Changing the Target-Disk Radius	65
4.5.1. The Effect on Total Number of Muons Stopped	65
4.5.2. The Effect on Signal Acceptance	70
4.5.3. The Effect on Detector Background	75
4.5.4. Summary and Performance Overview	80
4.6. Material Budget: Target-Disk Thickness	81
4.6.1. Summary and Performance Overview	83
4.7. Material Budget: Total Number of Target-Disks	84
4.7.1. The Effect on Total Number of Muons Stopped	84
4.7.2. The Effect on Signal Acceptance	85
4.7.3. The Effect on Detector Background	87
4.7.4. Summary and Performance Overview	90
4.8. Disk Positioning Along the Beam Axis	91
4.9. Signal Background from Muon Decay In Orbit	94
4.9.1. Updated Results for the Default Configuration	94
4.9.2. Results for New Configurations	99
4.9.3. DIO Background from Other Isotopes	102
4.9.4. DIO Background from the Surrounding Helium Gas	104

4.10. Single Event Sensitivity for New Configurations	106
4.10.1. Calculating the SES	106
4.10.2. Summary and Interim Results	107
4.11. Production of Target-Disks	108
5. The Muon-Stopping Target Suspension	113
5.1. Technical Layout of the Suspension Structure	113
5.1.1. The Outer Ring Structure	115
5.1.2. The Three Long Bars	117
5.1.3. The Support Ring	118
5.2. Holding the Target-Disks in Place: The Spokes	118
5.3. The Installation Procedure	121
5.4. ICEDUST Study of Key Parameters	123
5.4.1. Positional Accuracy	123
5.4.2. Size of the Three Long Bars	125
5.5. Final Results and Summary	126
6. The Muon-Stopping Monitor	129
6.1. Methodology: The Two Main Parameters	131
6.1.1. Assessing the Germanium Detector Background	131
6.1.2. Calculating the Germanium Detector Signal Acceptance	134
6.2. Results: Positioning and Collimator	141
6.3. Discussion and Outlook	149
6.3.1. Moving Toward a Realistic Shielding	149
6.3.2. HPGe Data Acquisition: Read-out and Dead Time	153
6.3.3. Estimation of Accuracy	156
6.3.4. COMETs Timing Structure and Supplementary Measurements . . .	158
7. Conclusion	161
List of Acronyms	163
References	164
A. Collimator Performances	173

List of Tables

2.1. COMET's SES and contributing factors	20
3.1. Versions of the used external Geant4 datasets	37
4.1. Parameters of the default muon-stopping target design in COMET Phase-I	42
4.2. Factors contributing to the μ -e conversion signal acceptance	47
4.3. Geometric acceptance and breakdown of the tracking quality cuts	50
4.4. Updated performance overview of the 17×0.2 mm disk configuration	58
4.5. Particle separated contribution to total CDC occupancy for the default muon-stopping target . .	60
4.6. Number of follow-up target-disk hits for varying target-disk radii	74
4.7. Performance overview of configurations with varying target-disk radii	80
4.8. Number of follow-up target-disk hits for varying total number of target-disks	83
4.9. Performance overview of configurations with varying number and thickness of target-disks . .	84
4.10. Particle separated contribution to total CDC occupancy for 33 target-disks	89
4.11. Performance overview of configurations with varying total numbers of target-disks	90
4.12. Momentum cut acceptance and expected DIO background of improved configurations	100
4.13. Estimation of DIO background for a variety of dangerous elements	103
4.14. Updated geometric acceptance and breakdown of the tracking quality cuts	107
4.15. Updated results for the COMET Phase-I SES of various configurations	108
4.16. Maximum impurities of DIN standard “EN AW-1050A” and “EN AW-1200” aluminium . . .	110
5.1. Maximum permissible weight on the spokes for tungsten wires with varying diameter	120
6.1. Energy and intensity of muonic X-rays in aluminium	130
6.2. Parameters used to calculate the interaction probability of the $2p \rightarrow 1s$ X-ray	137
6.3. Contributions to the overall uncertainty on the number of muons stopped	157

List of Figures

1.1. DIO electron spectrum, normalized to the free-muon decay rate Γ_0	8
1.2. Feynman diagrams of CLFV processes allowed within the SM	10
1.3. Historical progress of muon CLFV searches	11
1.4. Complementarity and reach of past and upcoming muon CLFV experiments	13
2.1. Principle procedure of $\mu^- \rightarrow e^-$ conversion experiments	14
2.2. Schematic layout of COMET Phase-I and Phase-II	16
2.3. Schematic layout and working principle of the Cylindrical Detector system	19
2.4. Schematic of the timing structure used in the COMET experiment	23
2.5. Pion production yields after collision from 8 GeV protons with tungsten and graphite targets .	24
2.6. Cross-sectional view of the curved solenoid highlighting the dispersion effect	25
2.7. Particle yields before and after passing the bent 90° transport solenoid	26
2.8. Relationship between the lifetime of muonic atoms and the atomic number	27
2.9. Dependency of the conversion electron energy from the atomic number	28
2.10. Target dependence of the $\mu^- \rightarrow e^-$ conversion rate in single-operator dominance models .	29
3.1. COMET Phase-I geometry as used in ICEDUST	31
3.2. Demonstration of the MC5 resampling plane	32
3.3. Close up on select ICEDUST geometry details	35
4.1. Mock-up of the default muon-stopping target design	41
4.2. Momentum distribution of muons reaching the muon-stopping target	43
4.3. Side view of the CyDet area	45
4.4. Cut-open 3D view of CyDet area	45
4.5. Example of a muon x - y -profile recorded in front of the muon-stopping target	46
4.6. Demonstration of 4-fold coincidence and its background suppression	48
4.7. Momentum transfer matrix for electrons originating from the muon-stopping target	49
4.8. Comparison between momentum distributions for single- and multiple-turn tracks	50
4.9. Momentum distribution of conversion electrons at CDC entry	51
4.10. Total energy deposits per cell for signal electrons and noise hits	54
4.11. Occupancy in CDC with/without timing window cut	55
4.12. CDC layer occupancy for the default muon-stopping target design	59
4.13. Muon x - y -distribution in front of the muon-stopping target for different momentum ranges .	62
4.14. Projected muon beam x - and y -distribution along the muon-stopping target	64
4.15. Comparison of the x - y -distributions of stopped and passing muons	66
4.16. Helix radius and middle point distribution of muons traversing CyDet	67
4.17. Percentage of muon helix middle points confined within a given radius	67
4.18. Comparison of momentum distribution of stopped muons between four configurations	69
4.19. Vavilov distributions for various absorber thicknesses	70
4.20. Conversion electron momentum distribution and survival probability after X mm of aluminium	71
4.21. Paths contributing to a conversion electron's traversed material trough target-disks	72
4.22. Travel distance of conversion electrons inside target-disks	73

4.23. Helix radius and middle point distribution of beam electrons traversing CyDet	75
4.24. Beam electron x - y -distributions in front of the muon-stopping target	76
4.25. Helix radius distributions from DIO electrons inside CyDet	77
4.26. CDC layer occupancy for enlarged $r = 200$ mm target-disks	78
4.27. Muon stopping distribution for enlarged $r = 200$ mm target-disks	79
4.28. Amount of traversed target material for varying number and thickness of target-disk	82
4.29. Number of aluminium target-disks required to stop muons of varying momenta	85
4.30. Effect of an increased total number of target-disks on CDC cut survival probability	86
4.31. Summary of normalized signal for configurations with varying total number of target-disks	87
4.32. CDC and CTH background summary for varying total number of target-disks	88
4.33. CDC layer occupancy for a muon-stopping target configuration with 33 target-disks	89
4.34. Position dependent probability to fulfill the CDC and/or CTH trigger conditions	92
4.35. Signal acceptance for varying US / DS z -offsets	93
4.36. Comparison of DIO endpoint spectra for various isotopes	95
4.37. The spectrum of DIO and μ -e conversion electrons as seen by CDC	97
4.38. Impact of the lower bound of the momentum ROI (default configuration)	99
4.39. Impact of the lower bound of the momentum ROI (improved configuration)	101
4.40. Normalized DIO spectra for various elements as seen by CDC	104
4.41. Position of muon stops inside the helium gas	105
4.42. Technical drawing of the aluminium targt-disk	109
4.43. Comparison between four aluminium target-disks made from different purity and temper	112
5.1. Technical drawing of the suspension structure	114
5.2. Implementation of the muon-stopping target-disk suspension in ICEDUST	115
5.3. Photograph of the large ring component of the muon-stopping target suspension	116
5.4. Photographs of the thin tungsten spokes holding the target-disks in place	119
5.5. Photograph of the cart used to install the suspension	122
5.6. Schematic of wedges inside US CTH	122
5.7. Effects of target-disk displacements	124
5.8. Signal acceptance for varying suspension bar materials and sizes	126
5.9. Photographs of the fully assembled suspension structure	128
6.1. Technical sketch of the HPGe positioning relative to the muon-stopping target	130
6.2. Last point of interaction before particles deposit energy in the HPGe	133
6.3. Example of a calculated target-disk solid angle	135
6.4. Definition of effective detector crystal radius r_a^*	136
6.5. Probability of an X-ray reaching the HPGe detector without being absorbed or scattered	138
6.6. Intrinsic full-energy peak acceptance of the employed HPGe detector	139
6.7. Full-energy peak efficiencies of measurement, Geant4, and custom MC simulation	140
6.8. Comparison of energy depositions inside the HPGe detector for different angles	142
6.9. Schematics of the investigated collimator configurations	144
6.10. Summary for the tube collimator	145
6.11. Collimator acceptance profiles for the first, middle, and last target-disk	146
6.12. Final comparison between select collimator designs	148
6.13. Momentum distribution of particles impinging on the HPGe shielding or collimator	150
6.14. Conceptual shielding design for COMET's HPGe detector	152
6.15. Example pulse shapes of a RFP and TRP preamplifier systems	154
6.16. Decay chain of ^{27}Mg	159
A.1. Summary for the cone collimator	173

A.2. Summary for the slit-type tube collimator	174
A.3. Summary for the ellipse collimator	175
A.4. Summary for the elliptical cone collimator	176

1. Introduction to Muons

The muon made its impactful science debut in 1937 when cosmic ray experiments measured a previously unknown “heavy electron” particle [1–3]. Initially mistaken as the predicted Yukawa particle (later identified as the pion [4]), its distinct properties soon identified it as a new kind of elementary particle. The history of this particle, eventually named the muon, spans over vastly different domains of physics. Slowing down and capturing muons within an atomic orbit covers solid-state and atomic physics. The subsequent electromagnetic cascade extends into the field of quantum electrodynamics, while nuclear effects, such as the nuclear charge distribution, become prominent as the muon gets closer to the nucleus. Finally, the realm of weak interactions and neutrino physics is entered when the muon decays or is captured by the atomic nucleus. Additionally, a steady growth in the number of practical applications of the muon can be witnessed over recent decades. For instance, techniques such as muon tomography or muography, utilized for generating three-dimensional images of volumes, could be recently applied to discover hidden rooms within the 4500-year-old Great Pyramid of Giza [5]. Overall, it is no surprise that countless experiments have dedicated their efforts to the muon, leading to an ever-growing understanding of the nature of subatomic particles and their interactions.

Despite these years of research, muons are still a prominent research topic in modern particle physics. They will also take the central stage within the COMET experiment introduced in Chapter 2. Hence, the beginning of this thesis will briefly establish the results of past research and cover the muon’s “known aspects” encompassed within the so-called Standard Model (SM) of particle physics. The second half of this chapter will then delve into still-open questions surrounding the muon to highlight its enduring relevance in discovering and exploring new physics.

1.1. Muons in the Standard Model

Within the SM, the muon is categorized as second-generation charged lepton. That means the muon is a spin- $\frac{1}{2}$ elementary particle with an electric charge of $-1e$. The only fundamental difference to an electron is its approximately 207 times larger mass and its “muonness”. The large mass causes the muon to be unstable. In a vacuum, it decays through the weak force with a mean lifetime of $2.2\,\mu\text{s}$ into the lighter electron. The “muonness” is categorized more physically as lepton flavour \mathcal{L}_μ .

1.1.1. The Free Muon Decay and Flavour Conservation

As theorized only four years after the muon's discovery, the (missing) energy of the decay electron indicates that a free muon decay must be accompanied by the “simultaneous creation of a number of neutrinos” [6]. In today’s physics, the free decay is expressed as:

$$\begin{array}{ccccccc}
\mu^- & \rightarrow & e^- & + & \nu_\mu & + & \bar{\nu}_e \\
\mathcal{L}_\mu: & +1 & \rightarrow & 0 & +1 & & 0 \\
\mathcal{L}_e: & 0 & \rightarrow & +1 & 0 & & -1
\end{array} \tag{1.1}$$

As predicted, two neutrinos are emitted¹. The necessity for two neutrinos – a muon neutrino ν_μ and an electron antineutrino $\bar{\nu}_e$ – were the first hints for the “muonness” (or the muon flavour)². The SM incorporates this characteristic as *lepton flavour*. Each set of charged lepton, namely electron e ; muon μ ; and tau τ , and their matching uncharged partner neutrinos (ν_e , ν_μ , and ν_τ) is assigned an individual lepton flavour quantum number $\mathcal{L}_{e/\mu/\tau}$. As indicated at the bottom of decay Equation 1.1, flavour quantum numbers equal +1 for particles and neutrinos, and -1 for antiparticles and antineutrinos. Within a decay or reaction, the sum of each type of lepton number is conserved, a principle known as *lepton flavour conservation*.

1.1.2. Muonic Atoms

Apart from the shown free decay, *negative* muons can also be captured within the atomic orbit of an atom. Excellent step-by-step introductions to this process are provided in [9] and [10]. Initially, the muon must be slowed down to a few keV. At energies above several hundred GeV, energy loss occurs primarily through radiative processes, while at lower energies, it happens due to ionization or excitation of surrounding matter [11]. Once the muon loses all its kinetic energy and comes to a stop in the material, it is promptly (around 10^{-13} s) captured into a (high) atomic orbit, thereby forming a *muonic atom*. The probability of atomic capture roughly depends on the number of loosely bound electrons [12]. The principal quantum number n_μ in which the muon enters the atomic orbit is usually around $n_\mu \gtrsim 20$ [13].

Muonic X-rays Once captured, the muon will quickly cascade down from these high orbits to lower energy because all lower muonic states are unoccupied. This near-instantaneous cascade happens at a time scale of around 10^{-13} s. The most common transitions occur in steps of $\Delta n_\mu = 1$, but transitions with $\Delta n_\mu \geq 2$ may also happen at a reduced intensity from 20 % to 5 % [14]. Muonic orbits around $n_\mu \simeq 14$ are of comparable size to K -shell electron

¹In addition to the Michel decay shown in decay Equation 1.1, alternative decay modes including a photon or electron-positron emission are possible but exhibit extremely low branching ratios of less than 1 % [7]. For the sake of simplicity, these decay modes are disregarded in the discussion.

²It took almost 20 more years until 1962 to validate the existence of muon neutrinos experimentally in the Brookhaven muon-neutrino discovery experiment [8].

orbits ($n_e = 1$). Hence, the muon cascade at high n_μ releases energy comparable to the electron binding energies and mainly occurs by Auger electron emission. As the transition energy rises rapidly toward lower orbits, shell electron interactions become suppressed, and around $n_\mu \simeq 5$ radiative transitions characterized by X-ray emission start to dominate³.

Muon Decay in Orbit In the context of SM physics, the muon's journey after reaching the lowest $1s$ atomic orbital can continue (or rather end) in one of two ways: the first, previously expressed with decay Equation 1.1, is the muon's decay into an electron and two neutrinos. However, as the Coulomb field of the nucleus strongly influences the bound muon, the process is now named Decay In Orbit (DIO) and shows some distinct features compared to the free decay.

One minor aspect of DIO is a slight reduction in decay rate, approximately 0.5 %, due to lower available muon energy from the binding ($E = m_\mu - E_{\text{binding}} < m_\mu$) and relativistic time dilation (average velocity $\langle \beta \rangle \simeq Z\alpha$). More significant is the change in the decay electron spectrum, as shown in Figure 1.1. The sharp edge of the free-decay energy spectrum around $E_e \sim m_\mu/2 \sim 53$ MeV is softened when accounting for nuclear recoil effects and the finite size and Coulomb field of the nucleus. The most crucial aspect for the later discussion is the formation of a high-energy tail extending to energies of almost $m_\mu = 105.66 \text{ MeV}/c^2$. Calculations of this effect were first introduced in [15] and recently updated in [16] to account for more isotopes.

Muon Nuclear Capture Instead of DIO, a bound muon can also be captured within the atom's nucleus. The reaction occurs through the charged current of the weak interaction and is written as:



With the exception of very light nuclei, muon nuclear capture is far more likely than decay. The capture process can occur in addition to DIO, thus both rates Γ simply add, and the total lifetime τ_{total} of bound muons is reduced:

$$\underbrace{\left(\frac{1}{\tau_{\text{total}}} \right)}_{\Gamma_{\text{total}}} = \underbrace{\left(\frac{1}{\tau_{\text{capture}}} \right)}_{\Gamma_{\text{capture}}} + Q \cdot \underbrace{\left(\frac{1}{\tau_{\text{free-decay}}} \right)}_{\Gamma_{\text{free-decay}}}. \quad (1.3)$$

³Note that the situation is considerably more complicated than this simple description. For one, in cases where Auger electrons cannot be sufficiently replaced, such as in gases, radiative transitions might become more dominant. For another, fine splitting of levels due to the muon's spin can noticeably broaden measured X-ray energies, especially for heavier elements (around $Z = 22$ the level separation is 2.2 keV already). Fortunately, both these effects are negligible for the choice of aluminium (solid metal, $Z = 13$) targets in the COMET experiment, thus they are disregarded throughout later discussions.

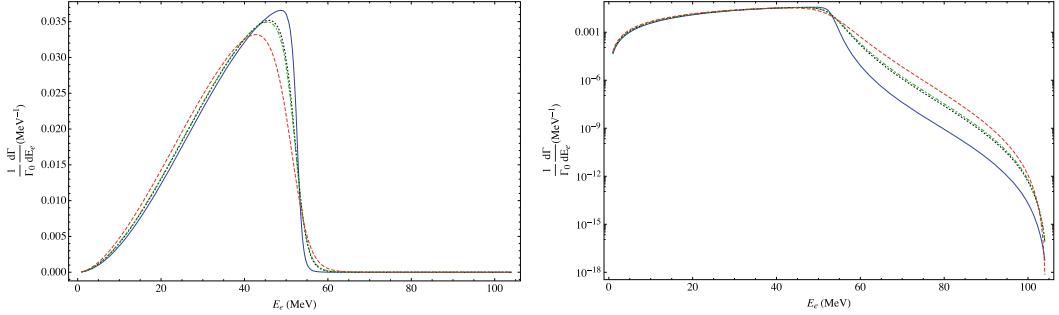


Figure 1.1.: DIO electron spectrum, normalized to the free-muon decay rate Γ_0 .

The colored lines correspond to carbon (solid blue), aluminium (black dotted), silicon (green dot-dashed), and titanium (red dashed). As discussed in the text, the bound decay spectrum exhibits a softer edge around the otherwise sharp edge around 53 MeV. Leading order radiative corrections cause a long high-energy tail up to almost $m_\mu = 105.66 \text{ MeV}/c^2$. Figures were taken from [17].

The formula incorporates a small correction factor Q , known as the Huff factor, which adjusts for the slight reduction of the bound muons' decay rate due to the effects discussed in the previous paragraph [18].

Following muon nuclear capture, the nucleus is left in a highly excited state. The neutrino carries away most of the released 105 MeV, but the nucleus can and does absorb a significant amount of energy, typically ranging from 10 MeV to 20 MeV. Deexcitation proceeds through various reactions, including the emission of neutrons, protons, or even alpha particles.

Investigating these deexcitation particles is valuable for learning more about the nucleus. The muon's approximately 207 times greater mass than the electron not only scales up the muonic atom's binding energies but also reduces the dimension of bound muon orbits. Historically, the close proximity of K -shell muons to the nucleus served as an ideal probe for nuclear charge and momentum distributions. Nowadays, the appreciably long period that muons spend "within" the nucleus serves as an excellent test to evaluate calculations of other weak nuclear reactions, such as electron capture in stars relevant for core-collapse supernovae [19]. Moreover, capture reactions without particle emission that leave the nucleus in a bound but highly excited state are of considerable interest. The high momentum exchange of several MeV makes these reactions the ideal tool to benchmark nuclear matrix element calculations, which are essential in the modern search for the neutrinoless double-beta decay [20].

1.2. Muons in Beyond Standard Model Physics

The SM, formulated throughout the 1960s and early 1970s, stands as a cornerstone of modern particle physics. Undoubtedly, the SM represents a remarkable achievement that has withstood numerous experimental tests. However, it has long been recognized that such comprehensive theories are either flawed or incomplete. Many lingering questions, such as the nature of dark matter or the puzzling matter-antimatter imbalance in the cosmos [21], strongly suggest the existence of more, the existence of *physics beyond the Standard Model* (BSM).

The muon stands at the frontier of such BSM searches. For instance, a long-persisting discrepancy between SM predictions and experimental findings concerns the muon's anomalous magnetic moment, referred to as $g-2$. The gyromagnetic ratio g of the spin- $\frac{1}{2}$ muon was originally predicted as $g = 2$ by Dirac in 1928 [22]. The value has since been refined and increased to account for radiative corrections (additional interactions with virtual particles) described by relativistic quantum field theory [23]. Nonetheless, the difference between the observed muon magnetic anomaly $a_\mu = (g - 2)_\mu/2$ and SM predictions persists. While initially less than 3 standard deviations [24], this deviation has increased over the following years, first to the level of 3.7 standard deviations [25], and most recently up to 5.1 standard deviations [26]. Thus, despite a remarkably high precision (experiment and SM predictions are within parts per million), the muon's $g-2$ anomaly stands as one of the most significant discrepancies between data and SM and embodies a promising candidate to observe BSM physics.

1.2.1. Motivation behind (Charged) Lepton Flavour Violation

Of comparable relevance for BSM searches as the anomalous muon $g-2$ is the violation of lepton flavor conservation introduced in Section 1.1.1. Strong evidence for the existence of *lepton flavour violation* (LFV) dates back to the late 1960s when the Homestake experiment measured a significantly lower solar neutrino flux than predicted by solar model calculations [27]. This solar neutrino problem was resolved only a few decades later, which culminated in the award of the Physics Nobel Prize in 2015 for the groundbreaking discovery of neutrino oscillation by the Super Kamiokande and SNO (Sudbury Neutrino Observatory) experiments [28, 29].

Within the “original” SM, neutrino oscillation (or the mixing of neutrino flavours) is forbidden because neutrinos are assumed to be massless [30]. However, neutrino oscillation was now experimentally confirmed. In a simplified scenario, assuming only two neutrino mixing (i.e. only one nonvanishing mixing angle θ_{ij}), the probability of observing neutrino oscillation from flavour ν_α to ν_β over a traveled distance L through vacuum is given by [31]:

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sin^2(2\theta_{ij}) \cdot \sin^2\left(\frac{L \cdot \Delta m_{ij}^2}{4E}\right) \quad (1.4)$$

It becomes clear from this formula that neutrino oscillation is only possible for non-zero neutrino masses ($\Delta m_{ij}^2 \neq 0$). Consequently, the SM had to be extended. The established approach to accommodate neutrino mixing within the SM is introducing a unitary 3×3 mixing matrix. Similar to the mixing of quarks [32], this Pontecorvo–Maki–Nakagawa–Sakata matrix connects the neutrino flavour eigenstates (ν_e, ν_μ, ν_τ) with their mass eigenstates (ν_1, ν_2, ν_3). In essence, it embeds (neutral) LFV within the SM.

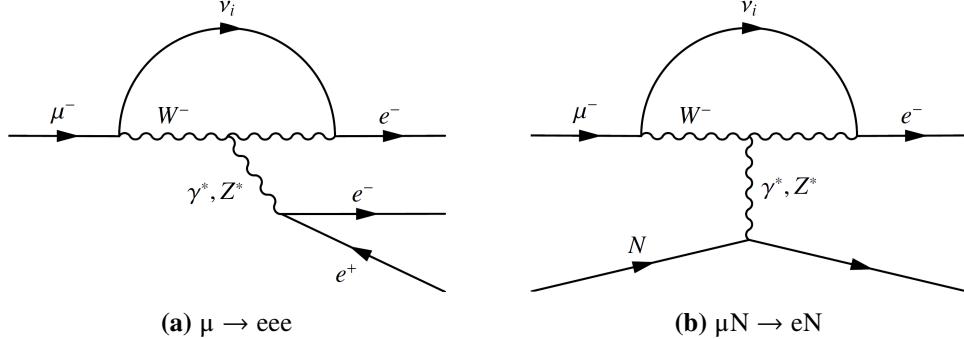


Figure 1.2.: Feynman diagrams of charged lepton flavour violating processes allowed within the SM after its extension for neutrino mixing. As discussed in regards to Formula 1.5, branching ratios of these processes are suppressed below measurable levels.

Charged Lepton Flavour Violation With the existence of (neutral) LFV proven and the mixing of quarks and neutrinos already permitted, questions naturally turn toward the possibility of lepton mixing within the charged sector. If the SM is minimally extended to include neutrino oscillation and Dirac neutrino masses, such a *charged* lepton flavour violation (CLFV) can already occur through one-loop diagrams involving a couple of $W\nu\nu$ vertices, as shown in the Feynman diagrams in Figure 1.2. However, these CLFV rates are heavily suppressed to unobservably small levels due to the low (limit on) neutrino masses ($m_\nu < 0.8$ eV [33]). For instance, the predicted branching ratio for the decay $\mu \rightarrow e\gamma$ is [34, 35]:

$$\mathcal{B}(\mu \rightarrow e\gamma) = \frac{3\alpha}{32\pi} \left| \sum_{i=2,3} U_{\mu i}^* U_{ei} \frac{\Delta m_{i1}^2}{M_W^2} \right|^2 \approx 10^{-54}, \quad (1.5)$$

and other muon CLFV processes similarly remain well below sensitivities achievable by modern experiments [36].

Perhaps surprisingly, however, these incredibly small branching ratios for CLFV caused by the tiny non-zero neutrino mass turn out to be a boon in disguise rather than a hindrance. In the absence of competing SM processes, a measurement of CLFV can be directly connected

to BSM physics. Moreover, in various theoretical BSM models such as supersymmetry-based grand unification theories [37] or models involving a new gauge boson Z' [38], leptoquarks [39], or heavy right-handed sterile neutrinos [40], CLFV rates are inherently enhanced. These models introduce new couplings and lepton interactions, resulting in new sources of CLFV, and raise CLFV rates to potentially accessible levels for next-generation experiments [41].

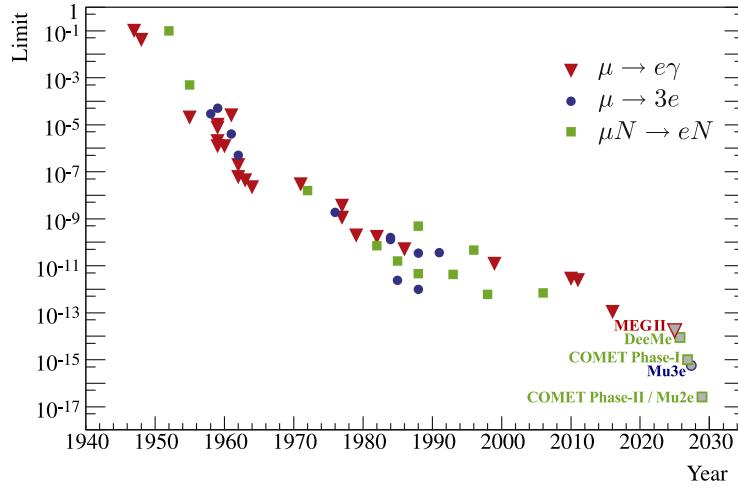


Figure 1.3.: Historical progress of experiments searching for muon CLFV. Following the discovery of the muon in cosmic rays, searches for muon CLFV advanced rapidly around the 1950s and 1960s with the development of synchrocyclotrons. After a brief plateau, a second phase of rapid progress occurred in the 1980s and 1990s due to advancements in equipment and the advent of modern computing, which lowered the branching ratio limits on these processes to approximately $O(10^{-12})$. Subsequently, the rate of improvement began to level off. However, the confirmation of neutrino mixing and neutral LFV rekindled the interest in flavour physics. A series of modern experiments, such as DeeMe [42]; COMET [43]; or Mu2e [44], are currently underway to search for muon CLFV with branching ratios up to $O(10^{-17})$. The Figure was adapted from [35] and supplemented with limit predictions from these next-generation experiments.

However, with an abundance of available BSM theories trying to explain the neutrino masses, the question for an experimental physicist becomes: *where should one look for CLFV?* The answer to this question far exceeds the scope of this introduction, but it is incredibly well discussed in [45]. For the muon, three channels have been extensively studied, namely $\mu \rightarrow e\gamma$, $\mu \rightarrow eee$, and $\mu N \rightarrow eN$ (called μ -e conversion). The historical progress and prospects of experiments searching for these three processes is shown in Figure 1.3.

Following [46], the three processes can be described (at an experimental scale $\sim m_\mu$) by the effective Lagrangian:

$$\begin{aligned} \delta\mathcal{L} = \frac{1}{\Lambda_{\text{LFV}}^2} \Big[& C_D \left(m_\mu \bar{e} \sigma^{\alpha\beta} P_R \mu \right) F_{\alpha\beta} \\ & + C_S (\bar{e} P_R \mu) (\bar{e} P_R e) \\ & + C_{VR} (\bar{e} \gamma^\alpha P_L \mu) (\bar{e} \gamma_\alpha P_R e) \\ & + C_{VL} (\bar{e} \gamma^\alpha P_L \mu) (\bar{e} \gamma_\alpha P_R e) \\ & + C_{\text{N-light}} \mathcal{O}_{\text{N-light}} + C_{\text{N-heavy}} \mathcal{O}_{\text{N-heavy}} \Big]. \end{aligned} \quad (1.6)$$

Parameters $\vec{C} = \{C_D, C_S, C_{VR}, C_{VL}, C_{\text{N-light}}, C_{\text{N-heavy}}\}$ correspond to the most relevant (dimensionless) Wilson coefficients for muon CLFV, while Λ_{LFV} defines the mass scale of the New Physics. Note that the dipole operator in the first term of this Lagrangian mediates the $\mu \rightarrow e\gamma$ decay but also contributes to both $\mu \rightarrow eee$ and $\mu N \rightarrow eN$ in cases where the capturing nucleus absorbs the photon. On the contrary, the following three contact operators for $\mu \rightarrow eee$, as well as both $\mathcal{O}_{\text{N-light}}$ and $\mathcal{O}_{\text{N-heavy}}$ operators probing muon conversion on light ($Z \simeq 13$) and heavy ($Z \simeq 79$) conversion targets have no photonic dipole contribution. Hence, a combination of experiments measuring different channels will prove invaluable for differentiating various BSM models.

This complementarity⁴ of muon CLFV experiments is illustrated in Figure 1.4. In order to produce this plot, the Wilson coefficients \vec{C} are transformed to the LFV scale Λ_{LFV} via renormalization group equations, parameterized in spherical coordinates, and expressed in terms of $\kappa_D = \cot(\theta_D - \pi/2)$ (see [46] for details). $|\kappa_D| \ll 1$ corresponds to a dominant dipole operator in Formula 1.6, while for $|\kappa_D| \gg 1$ the four-fermion operators dominate.

Several conclusions can be drawn from Figure 1.4. Firstly, past experiments have already covered New Physics mass scales up to around 1000 TeV, and upcoming experiments will probe regions of 10^4 TeV and beyond. Secondly, experimental bounds from the $\mu \rightarrow e\gamma$ process provide the most stringent constraints on the dipole operator. However, probing branching ratios below $\mathcal{O}(10^{-14})$ in this channel proves experimentally challenging due to a significant increase in accidental $e\gamma$ -coincidences at higher muon beam intensities⁵. In contrast, $\mu \rightarrow eee$ conversion best constrains the four-fermion operators. Moreover, this channel outperforms $\mu \rightarrow e\gamma$ for most values of κ_D if one can achieve around two orders of magnitude higher sensitivity. As discussed in Section 2.2, an inherent advantage when

⁴The authors of [46] define complementarity as “the ability of different experiments to probe independent operator coefficients”.

⁵Using detector parameters similar to those of the MEG experiment and a muon rate of $R_\mu \simeq 10^8 \mu/\text{s}$, one expects one background and one signal event for a branching ratio of $\mathcal{B}(\mu \rightarrow e\gamma) = 10^{-14}$. Probing $\mathcal{B}(\mu \rightarrow e\gamma) = 10^{-16}$, an expectation value of one signal corresponds to a muon rate of $R_\mu \simeq 10^{10} \mu/\text{s}$. However, at such a muon rate, accidental backgrounds would be increased to 10^4 events [47].

measuring μ -e conversion is the process' clean experimental signature of a single electron. Consequently, achieving branching ratios of $\mathcal{B}(\mu N \rightarrow e N) \approx 10^{-16}$ is feasible, and μ -e conversion represents a golden channel to discover BSM physics in the coming decade.

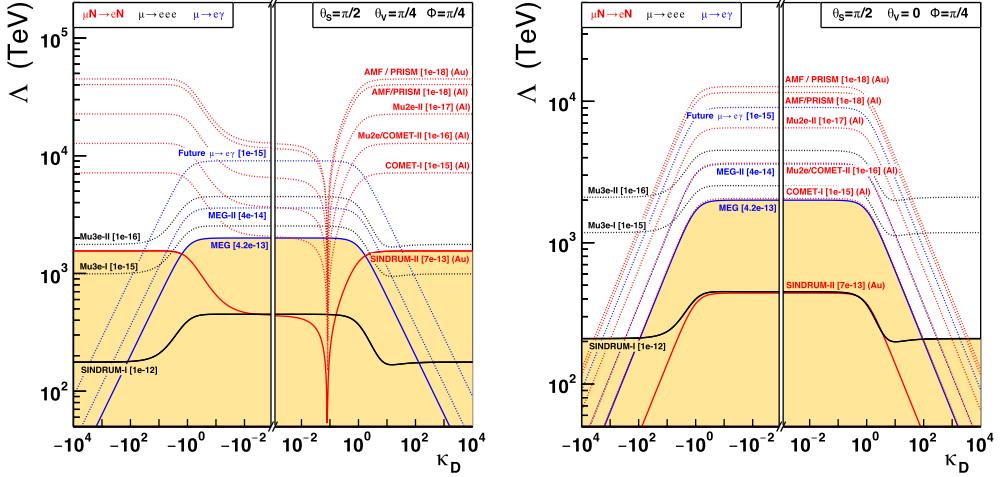


Figure 1.4.: Complementarity and reach of past (solid lines) and upcoming (dotted lines) muon CLFV experiments. The solid yellow region is currently excluded. The New Physics scale Λ_{LFV} and the variable κ_D , which parametrizes the relative magnitude of dipole (dominant for $|\kappa_D| \ll 1$) and four-fermion (dominant for $|\kappa_D| \gg 1$) coefficients, are defined over Formula 1.6. The remaining parameters are set to typical values as described in [46].

2. The COMET Experiment

The COherent Muon to Electron Transition (COMET) experiment aims to investigate the elusive phenomenon of muon-to-electron ($\mu^- \rightarrow e^-$) conversion. The experiment is currently under construction at the Japan Proton Accelerator Research Complex (J-PARC) in Tokai, Ibaraki, Japan. COMET seeks to improve the current limit on the conversion rate, set by the SINDRUM-II experiment as 7×10^{-13} (90 % C.L.) in muonic gold [48], by at least four orders of magnitude. To achieve the desired sensitivity in the order of $O(10^{-17})$, the design of COMET must meet several stringent requirements.

First, an intense source of negatively charged muons is essential. As further elaborated in the introduction Section 1.1, muons, due to their higher mass compared to electrons, are unstable particles with a mean free lifetime of only 2.2 μs and thus are not abundant on earth. Although secondary muons from cosmic rays constantly bombard the earth's surface, their properties, such as fluence and energy distribution, cannot be controlled, rendering them unsuitable for the experiment. Therefore, the adopted procedure for producing muons in COMET, as well as in other muon experiments like Mu2e [44] or DeeMe [42], is based on the sequence shown in Figure 2.1:

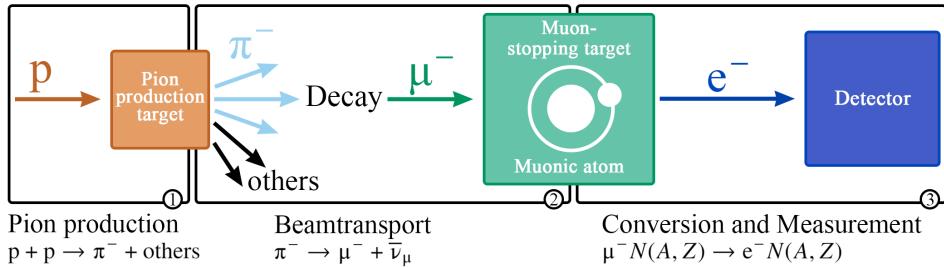


Figure 2.1.: Principle procedure of $\mu^- \rightarrow e^-$ conversion experiments.

In step one of Figure 2.1, an intense 8 GeV proton beam obtained from the J-PARC main ring synchrotron impinges on a graphite pion production target. Negatively charged pions produced in this process are then collected, formed into a beam by quadrupole magnets, and transported along a beamline toward the detector systems.

During step two, the desired muon beam is produced from in-flight decays of pions, taking advantage of the pions' short lifetime of 26 ns. However, the proton-proton collision in the graphite target also generates various undesired particles, including neutrons, antiprotons, or differently charged pions (referred to as *others*). These particles could easily overwhelm

or even damage the detectors following later in the chain. Thus, shielding and mitigation measures to address these backgrounds, as well as strict muon filtering via charge and momentum selection during beam transport are crucial for the experiment’s success.

In the third step shown in Figure 2.1, muons that reach the detector area are stopped inside a muon-stopping target, where they might undergo conversion to electrons. COMET must optimize both the available amount of stopped muons as well as the detector acceptance for a potential conversion signal. This last stage is the primary subject of this thesis. However, all building blocks leading up to this part of the experiment are essential. In order to guide the understanding for later discussions, this chapter will briefly introduce COMET’s most important aspects and design decisions.

2.1. A Staged Approach: COMET Phase-I and Phase-II

The COMET experiment will be carried out in a two-stage approach. The initial stage, Phase-I (shown on the left in Figure 2.2), will be operated with a reduced proton beam intensity and a short 90° muon-beamline. Phase-I has two primary objectives: First, the measurement of $\mu^- \rightarrow e^-$ conversion at a Single Event Sensitivity (SES) of $O(10^{-15})$ (see Section 2.2.2), utilizing a Cylindrical Detector system (CyDet). Second, a direct measurement of the muon beam using a combination of straw-tube tracker stations followed by an electromagnetic calorimeter, called StrECAL. The Phase-I beam measurement is intended to enhance the understanding of the pion production system, including particle yields and background rates, as well as particle transportation using bent solenoids. Furthermore, the measurement will be utilized to verify the underlying Monte Carlo (MC) simulations and to establish the foundation for future upgrades in COMET’s second stage, Phase-II.

In Phase-II, the muon transport solenoid will be extended by an additional 90° . Moreover, a 180° bent solenoid is introduced after the muon-stopping target to improve the measurement of conversion electrons and further suppress background, as shown on the right in Figure 2.2. The experience gained from Phase-I will be used to fine-tune the StrECAL detector, and an improved version¹ will be employed for the conversion electron measurement. Consequently, Phase-II can utilize a higher beam intensity to achieve COMET’s overall design SES of $O(10^{-17})$.

Exact values of SES and proton beam parameters are summarized in Table 2.1 of Section 2.2.2. Details of experimental components are provided throughout the following sections, whereby the focus is placed on the currently constructed Phase-I, in particular the CyDet system used for the $\mu^- \rightarrow e^-$ conversion measurement. For further details regarding Phase-II, the reader is directed toward the official COMET Technical Design Report (TDR) [43] or a recent Phase-II study provided in [49].

¹Possible improvements include more straw-tube tracker stations and thinner diameter tubes to reduce electron multiple scattering for a better position and momentum resolution.

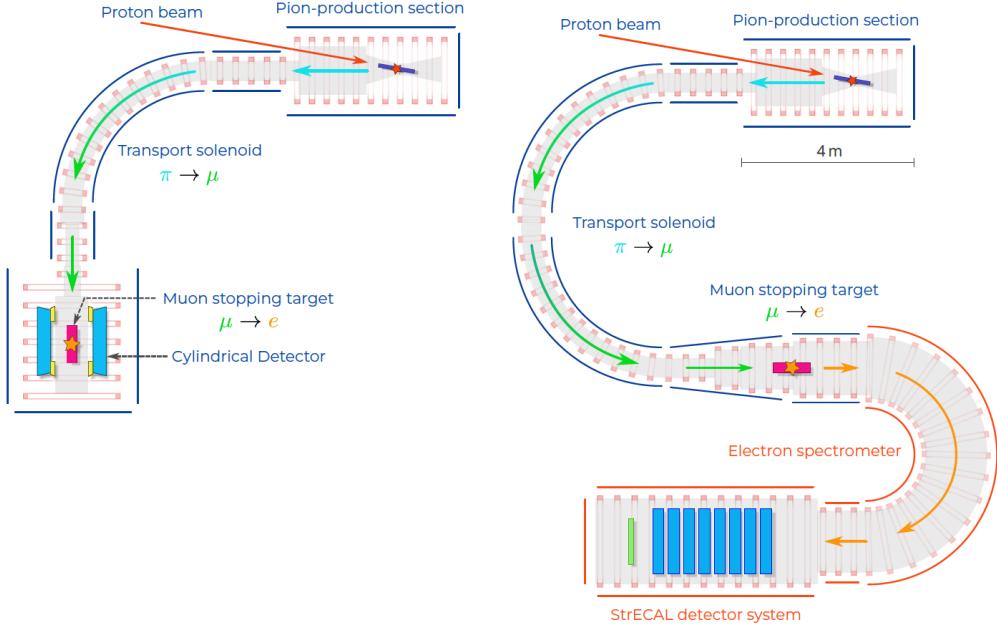


Figure 2.2.: Schematic layout of COMET Phase-I (left) and Phase-II (right). A semi-transparent grey path indicates the beamline. Solenoids for the magnetic field are shown as empty red rectangles. For details of the named components, see Sections 2.2 and 2.3. Figure taken from [50].

2.2. The μ -e Conversion Signal

To understand the fundamental design decisions behind the COMET layout, one must first understand the properties of the desired signal. Coherent muon-to-electron conversion, the BSM process in which a bound muon within a muonic atom decays without emitting a neutrino, can be expressed by:

$$\mu^- N(A, Z) \rightarrow e^- N(A, Z). \quad (2.1)$$

The nucleus $N(A, Z)$, characterized by its mass number A and atomic number Z , is essential in this process. As discussed in Section 1.2.1, μ -e conversion can occur through either photonic dipole or non-photonic four-lepton contact interactions. In the former scenario, the nucleus is necessary to absorb the photon, whereas in the latter, CLFV arises from the interaction of heavy virtual particles with the nucleus' quark system. Regardless of the interaction type, the nucleus either remains in its ground state or becomes excited after the decay.

COMET focuses specifically on μ -e conversion where the nucleus remains in its ground state, called a *coherent* transitions. Not only does the coherent channel dominate the total

decay width², but it also provides a significant experimental advantage: the emitted signal (or conversion) electron is mono-energetic. Its exact energy depends solely on the nucleus and can be calculated using:

$$E_e = m_\mu - E_{\text{binding}} - E_{\text{recoil}}. \quad (2.2)$$

The Formula includes corrections due to the muon's binding, $E_{\text{binding}} \approx \alpha^2 Z^2 m_\mu/2$, as well as the nuclear recoil, $E_{\text{recoil}} = (m_\mu - E_{\text{binding}})^2 / (2m_N)$, which depend on the fine structure constant α [52] and the recoiling nucleus, more precisely its atomic number Z or mass m_N . However, both E_{binding} and E_{recoil} remain in the order of only a few MeV and are small compared to the muon mass $m_\mu = 105.66 \text{ MeV}/c^2$. Consequently, the expected energy of the conversion electrons is in the order of 100 MeV. This places the signal firmly above typical Standard Model (SM) physics backgrounds, such as the free muon decay ($\mu^- \rightarrow e^- \nu_\mu \bar{\nu}_e$), where the electron energy spectrum has a clear cut off around $m_\mu/2 = 53 \text{ MeV}/c$ (see Section 1.1.1).

It should be highlighted again that the $\mu^- \rightarrow e^-$ signature consists of a *single* electron. This eliminates the need for a complex coincidence setup as it would be necessary for other CLFV channels, such as $\mu \rightarrow eee$ [53]. Furthermore, light muonic atoms have half-lives of hundreds of nanoseconds, which can be utilized to further eliminate the prompt background associated with muon production or beam transport. The details of such a background suppression are discussed in Section 2.3.

2.2.1. Detecting Conversion Electrons: A Cylindrical Detector System

In COMET Phase-I, the detector system follows after the 90° transport solenoid. This is particularly beneficial for the planned beam measurement with the StrECAL system because a position close to the pion production target simplifies studies of beam and solenoid dynamics. However, particle and momentum filtering with a short 90° beamline is also limited (further details see Section 2.3.2). Consequently, the expected flux of particles during the conversion measurement is too large to use the StrECAL system for a precise enough measurement and without risking damage to the detector. Instead, a dedicated detector system called CyDet will be used for the μ -e conversion search.

The layout and general working principle of the CyDet system are shown in Figure 2.3. The system consists of a Cylindrical Drift Chamber (CDC) used for particle trajectory reconstruction, as well as upstream and downstream scintillator rings for time-stamping and readout-triggering, referred to as CTH. The muon-stopping target necessary for the conversion search is positioned at the center of CyDet. Details of the muon-stopping target will be thoroughly discussed in Section 4.

²The coherent channel is dominant over the whole periodic table. For instance, in non-photonic interactions, the ratio of coherent to incoherent transitions in ^{12}C is approximately a factor of 5. The factor increases quickly up to about 9 around ^{13}Al and reaches its maximum of about 15 at the ^{82}Pb region [51].

The whole CyDet system is installed inside the warm bore³ of a large 1 T superconducting solenoid, causing charged particles to follow helical trajectories (details can be found in Section 4.5. However, both CDC and CTH are placed at a sufficiently large radius to ensure that prompt particles, following the proton collision in the pion production target, pass through their middle without causing any hits, thus avoiding the high-flux issues associated with the StrECAL system. The radius of the detector system is selected to suppress electrons with momenta lower than 60 MeV/c (see details in Section 4.5.3). Consequently, the majority of DIO background electrons are suppressed, while conversion electrons with around 100 MeV and possibly large transverse (outward-directed) momentum components can reach both detectors.

The CDC design is guided by requirements on momentum resolution. In particular, a resolution better than 200 keV for 105 MeV electrons must be archived to suppress the high-energy tail of the DIO spectrum (compare Figure 1.1 in Section 1.1.2). Low-mass materials must be utilized to reduce electron multiple scattering, which dominates the momentum resolution at this energy. Therefore, the inner (outer) wall of CDC is made from lightweight 0.5 mm (5 mm) thick carbon fibre-reinforced plastic. A total of 4986 sense and 14 562 field wires are arranged to form 20 concentric layers, with the inner- and outermost layer acting as so-called guard layers to remove space charge that would otherwise accumulate on the CDC walls. The 25 μm diameter gold-plated tungsten sense wires are placed at a high voltage of up to 1900 V. Each sense wire is surrounded by eight 126 μm diameter aluminium field wires at ground potential, forming a grid of almost-square CDC cells with nearly constant dimensions of 16.8 mm width and 16 mm height. Isobutane (He:C₄H₁₀ (90:10)) is selected as the counting gas. Cosmic ray tests conducted with the completed CDC since summer 2016 demonstrate a maximum drift time of 700 ns (around cell corners), a gain of approximately $1 \sim 4 \times 10^4$, and confirm the required momentum resolution of 200 keV [43].

The CTH consists of two modules located at each end of the CDC tube. A module consists of two rings, and each ring is segmented into 64 rectangular plastic scintillation counters (BC-408, Saint-Gobain crystal). Each counter is tilted at a specific angle and shifted roughly half its width against the previous counter. This alignment enables the suppression of significant parts of the background originating from the muon-stopping target by enforcing a 10 ns coincidence between hits in multiple counters, further discussed in Section 4.2.3. Utilizing the scintillators' fast timing resolution of 1 ns, the CTH is used to provide the time-stamp for conversion events and the first-level trigger for the CyDet system. This first-level trigger helps to identify and limit proton background occurring after nuclear muon capture at the muon-stopping target. The CTH scintillators are read out using long 5 to 10 m plastic fiber bundles connected to multi-pixel photon counters to mitigate radiation levels on readout electronics. However, optimization of this design is still ongoing, based on light yield, time resolution, and radiation performance [54].

³The warm bore of a solenoid is the area within the coil where the magnetic field is produced when the solenoid is energized. It's called "warm" because it tends to heat up due to the electrical current passing through the coil.

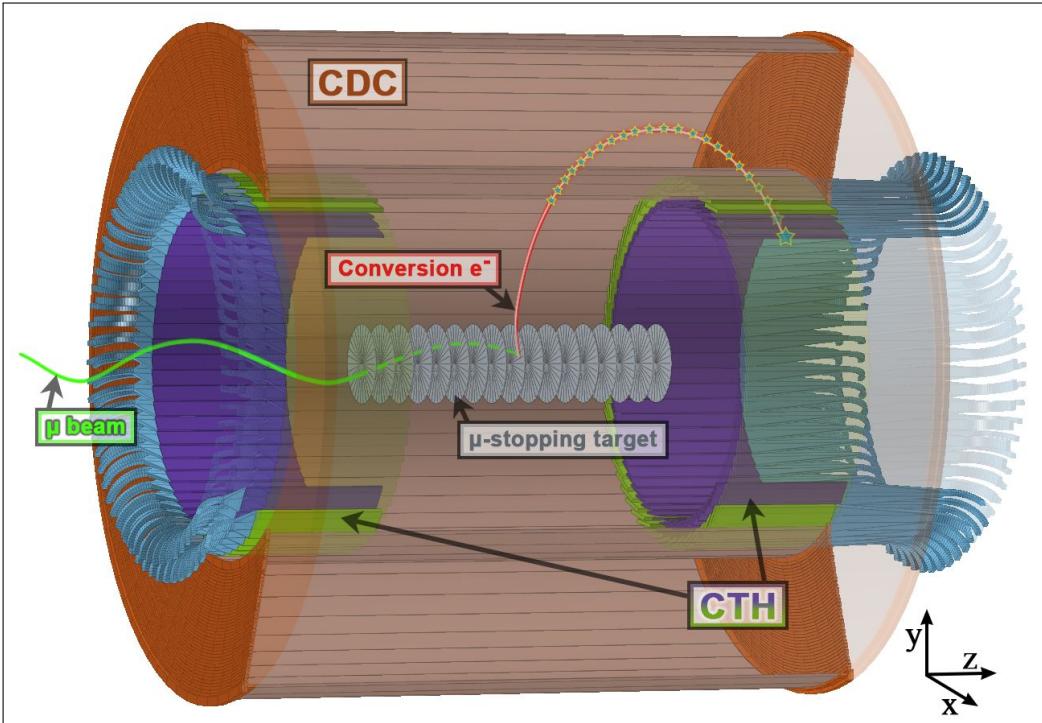


Figure 2.3.: Schematic layout and working principle of the Cylindrical Detector system (CyDet) system. A muon trajectory (green line), representing the beam, arrives from the transport solenoid (located on the left side) at the detector area. After a muon is stopped inside the muon-stopping target (grey disks) positioned at the center of CyDet, a conversion electron (red line) is emitted from the muon-stopping position. A 1 T magnetic field forces the electron on a helical trajectory. For clarity, the superconducting solenoid is omitted in the Figure. To be *accepted* by the detector, the electron must first spiral through the Cylindrical Drift Chamber (CDC) (orange tube). The electron momentum can be reconstructed from CDC hits (blue stars). Afterward, the electron must hit either the DownStream (DS) (right side) or UpStream (US) (left side) CyDet Trigger Hodoscope (CTH) counters to trigger the readout. The CTH is shown as green (outer) and dark blue (inner) rings, with light guides for readout in cyan. Note that an electron must first traverse the CDC and subsequently hit the CTH from the CDC side (larger radius) as illustrated. Heavy shielding on the inside (small radius) and front-side of CTH, designed to suppress background originating from the muon-stopping target, would otherwise alter the electron's momentum and prevent a subsequent momentum measurement.

2.2.2. The Single Event Sensitivity

The ability of COMET to measure μ -e conversion is expressed in terms of the SES. The SES is defined as the value for the partial branching ratio:

$$\mathcal{B.R.} = \frac{\Gamma(\mu\text{-e conversion})}{\Gamma(\text{nuclear capture})} \quad (2.3)$$

corresponding to an expectation value of one event, and it can be calculated using the formula:

$$\mathcal{B.R.} \stackrel{N_{\mu e}^{\text{obs.}}=1}{\equiv} \text{SES} = \frac{1}{N_{\mu} \cdot A_{\mu \rightarrow e} \cdot \mathcal{B}_{\text{capture}} \cdot f_{\text{coherent}}}. \quad (2.4)$$

Here, N_{μ} is the total number of muons stopped in the muon-stopping target, $A_{\mu \rightarrow e}$ is the detector's total signal acceptance, $\mathcal{B}_{\text{capture}}$ is the branching ratio for muon nuclear capture, and f_{coherent} is the fraction of conversions estimated to occur coherently, leaving the nucleus in its ground state.

Table 2.1.: Estimated COMET SES and contributing factors, as taken from the COMET Phase-I TDR [43] and a recent Phase-II study [49]. For details regarding both phases, see Section 2.1.

Parameter	Phase-I	Phase-II	Comment
(1) Number of stopped muons	1.5×10^{16}	3.3×10^{18}	
Total POT* (N_{POT})	3.2×10^{19}	8.8×10^{20}	
Beam current (I_p)	0.4 μA	7.0 μA	8 GeV proton beam
Run time (t_{run})	$\sim 146 \text{ d}$	$\sim 231 \text{ d}$	
Muon stop rate per POT ($R_{\mu/p}$)	4.7×10^{-4}	3.8×10^{-3}	
(2) Total detector signal acceptance	0.041	0.034	Details in Table 4.2
(3) Muon nuclear capture BR	0.61	0.61	For $^{27}_{13}\text{Al}$, taken from [55]
Single event sensitivity (SES)			
a) $f_{\text{coherent}} = 0.9$	3.0×10^{-15}	–	
b) $f_{\text{coherent}} = 1.0$	2.7×10^{-15}	1.4×10^{-17}	

* POT stands for Proton On Target.

The exact values of the SES and its contributing factors are presented in Table 2.1. Three aspects are essential to highlight regarding the SES. Firstly, despite being named *sensitivity*, its value should be as small as possible. Lower values equal a better performance because the SES corresponds to the lowest partial branching ratio $\mathcal{B.R.}$ for which one event is still expected to be observed.

Secondly, Formula 2.4 includes the model and nucleus dependent coherence factor f_{coherent} . Theoretically estimated incoherence rates are often limited to a small fraction ($\lesssim 10\%$) [51, 56]. For aluminium, the COMET TDR assumes $f_{\text{coherent}} = 0.9$ [43]. However, predictions

for this factor vary between different BSM models [57, 58]. To eliminate the dependency, this thesis adopts the approach of prior studies [59] and [49] and excludes this factor in the SES estimates. Therefore, the presented values for the SES should be understood as *only for coherent conversion*, where the state of the nucleus remains unchanged.

Lastly, one may observe that the SES only provides information about the experiment's ability to measure the signal. It does not include any information about experimental backgrounds. As the SES represents the probability to correctly identify *single* signals, using it as a figure of merit assumes a comprehensive understanding of the expected background. In fact, the number of background events must be demonstrated to remain comparably low (if possible far below one) in order to confidently label an observation as signal.

2.3. Understanding and Suppressing Background

Given COMET's reliance on the SES as a figure of merit, a concise understanding of the experimental background is essential. The unique properties of the μ -e conversion signal offer some advantages in distinguishing it from common SM backgrounds. Even though the remaining background contributions that can mimic a single electron around 100 MeV are small, they must not be disregarded. These backgrounds can be broadly categorized into three distinct groups:

- **Beam related:** This first group of background arises from impurities present in the muon beam. It is further divided into *prompt* and *delayed* background, based on their timing with respect to the arrival of the proton beam. Prompt backgrounds result from pion or muon decays-in-flight or radiative pion capture at the muon-stopping target. To produce signal-like high-energy electrons, the momentum of the in-flight decaying pions (muons) must exceed 50 MeV (70 MeV). Prompt backgrounds can be suppressed by limiting the flux of those higher-energy components and utilizing a bunched proton-beam structure in combination with stringent detector timing cuts, discussed in Section 2.3.1.

Delayed backgrounds exhibit a timing structure that precludes the application of timing cuts based on the bunched proton beam because their arrival time at the CyDet system overlays with follow-up bunches. This category includes low-energy neutrons, which can slowly travel and bounce along the beamline, leading to delayed signal-like events after neutron capture. Furthermore, the annihilation of low-energy antiprotons can directly or indirectly produce 100 MeV electrons. Due to the antiproton's negative charge and considerably larger mass compared to muons, they cannot efficiently be filtered by the beamline and arrive at much later times at the detector region. To mitigate antiproton components, a thin absorber material in the beamline or a lower proton beam energy could be considered.

- **Cosmics:** As recently shown in [50], cosmic rays are one of the most critical background source for COMET. Cosmic muons can mimic the conversion signal after decaying in flight or when interacting with the muon-stopping target or nearby surrounding materials. The

extent of this background mainly depends on the level of shielding above the detector system. For further mitigation, a dedicated cosmic ray veto system surrounding the entire detector region will be employed.

- **Intrinsic physics:** Intrinsic backgrounds arise from the interactions of muons with the muon-stopping target, which is an indispensable part of the experiment. Through processes like asymmetric e^+e^- conversion or Compton scattering, high-momentum electrons resembling the signal can be produced from gamma-rays emitted during radiative muon capture. Additionally, reaction chains initiated by high-energy charged particles or neutrons emitted from the nucleus after muon capture can produce signal-like backgrounds. Fortunately, these contributions are reasonably well suppressed and amount to less than 0.005 estimated events throughout the whole Phase-I runtime [43].

The dominant intrinsic background arises from DIO electrons. While the electron momentum in a free muon decay is limited to roughly half the muon mass, the nuclear recoil in DIO allows the electron to surpass this limit significantly. In fact, the endpoint of the DIO energy spectrum reaches up to the conversion electron energy (compare Section 1.1.2). Careful selection of materials (and purity) of the muon-stopping target and its surrounding, and the excellent 200 keV CDC momentum resolution in combination with stringent cuts on the electron momentum are used to suppress this contribution. Details and an updated estimate of DIO background are discussed in Section 4.9.

2.3.1. Primary Beam: Proton Target, Energy, and Timing Structure

Beam-related backgrounds are mitigated through careful adjustments to the proton beam and the pion production target properties. Pion production increases almost linearly with proton energy. However, the maximum permissible proton energy is constrained by the antiproton background. Antiproton production from proton-proton collisions within the pion production target becomes possible starting from 6.56 GeV and rapidly grows above this threshold. To limit antiproton production but also overall background radiation while still producing enough pions to achieve COMET’s design sensitivity within a reasonable timeframe, the optimal proton beam energy for COMET was determined to be 8 GeV. COMET Phase-I will utilize a beam intensity of 0.4 μA , translating to a power of 3.2 kW. The planned updates for COMET Phase-II, briefly discussed in Section 2.1, will enable an increase in beam intensity for this second phase of up to 7 μA or 56 kW.

As briefly mentioned in the last section, prompt beam backgrounds are suppressed using a bunched beam structure. Proton bunches from the J-PARC accelerator are separated by 550 ns. COMET intends to utilize every second bunch, resulting in proton bunches with a duration of 100 ns and a separation of at least 1.17 μs . This timing structure is transferred to the muon beam. Utilizing the lifetime of bound muons (further discussed in Section 2.3.3), a delayed measurement time window between proton bunches is implemented, in which beam background intensity is significantly reduced. A schematic of this principle is shown in Figure 2.4.

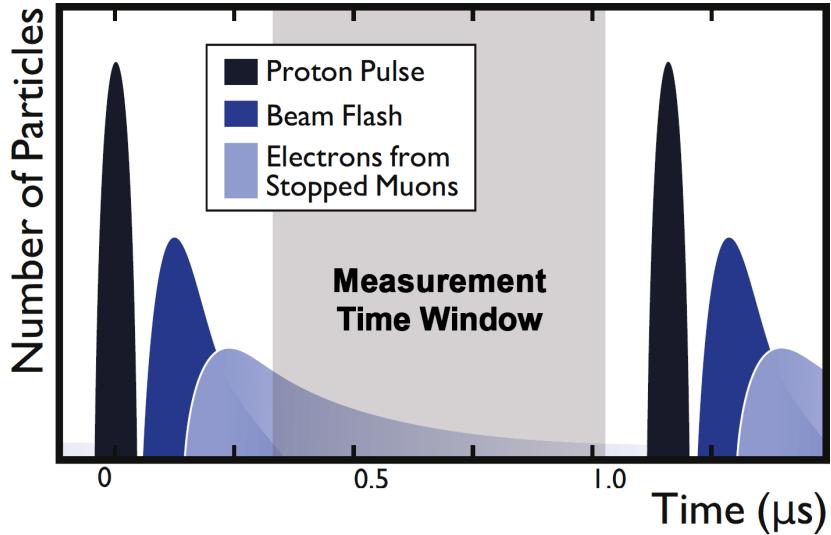


Figure 2.4.: Schematic of the timing structure used in the COMET experiment, as adopted from [49]. 100 ns proton bunches (black) impinge on the pion production target in 1.17 μ s intervals, each causing a flash of particles in the detectors (dark blue) which lasts for approximately 300 μ s. Utilizing the lifetime of bound muons, a delayed measurement time window (grey) between proton bunches is adopted, in which beam background intensity is significantly reduced.

The choice of the pion production target material is constrained by its mechanical strength and thermal resistance. Although targets with a high number of nucleons are preferable to increase pion yields, COMET Phase-I will use a lightweight graphite target. Graphite withstands temperatures up to a few thousand degrees. Considering the reduced Phase-I beam power, radiation cooling alone is sufficient for such a target. For Phase-II, tungsten has been selected, but a dedicated cooling scheme will be necessary.

For both materials, pion production yields after collision with 8 GeV protons are presented in Figure 2.5. The production rate is separated into forward and backward direction with respect to the proton beam. It can be seen that more pions are produced in forward than backward direction, but most of them have very high momentum. Besides high-momentum pions posing the risk of producing a signal-like background, they decay into comparably high-momentum muons with a low probability of being stopped in the muon-stopping target. In contrast, yields of low-momentum pions are similar in both directions. Consequently, the pion production and capture system is designed around backward-directed pions, utilizing a solenoid magnet of 5 T gradually decreasing to 3 T toward the matching beam solenoid. This focus on backward-produced particles additionally reduces a large number of other background particles boosted in the forward proton beam direction.

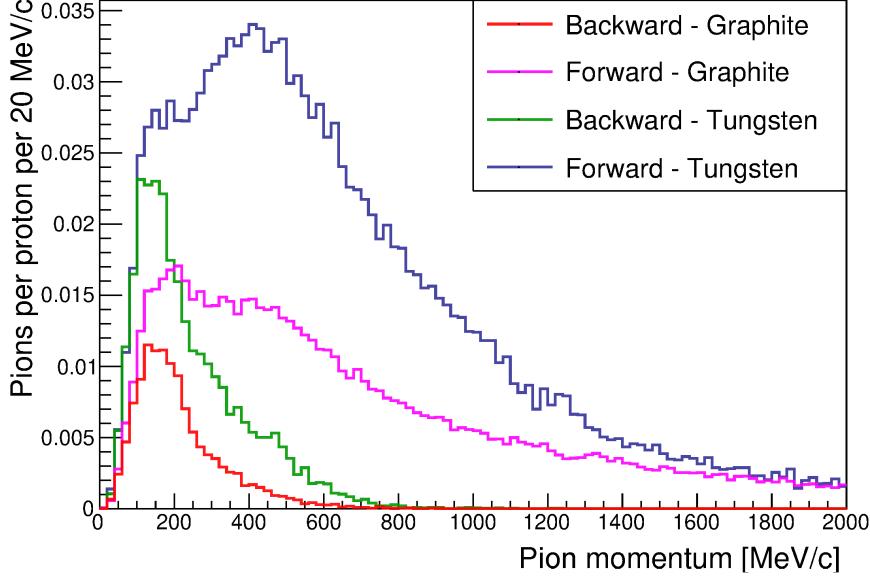


Figure 2.5.: Pion production yields after collision of 8 GeV protons with tungsten and graphite targets. The production rate is separated into forward and backward direction with respect to the proton beam. Data was generated using a Geant4 simulation with the QGSP_BERT hadron model [60]. Compared to graphite, tungsten targets yield overall more pions due to their higher atomic number. The image is adapted from [43].

2.3.2. Particle Transport: Particle Selection with Bent Solenoids

In Phase-I (Phase-II), pions originating from the pion production target are transported to the detector area by a 90° (180°) transport solenoid. The purpose of this curved solenoid design is twofold, namely allowing for pion decay to muons and particle filtering.

The adopted muon beam transport for Phase-I consists of superconducting solenoid magnets of 3 T over a length of 7.6 m. As charged particles traverse the solenoid's magnetic field, they are forced onto helical trajectories. Because of the solenoid's bend, the magnetic field strength is slightly higher at smaller radii compared to larger radii. Both effects combined cause a vertical drift of the helical trajectory's center. The drift is perpendicular to the plane of curvature and can be calculated with:

$$\begin{aligned}
D &= \frac{1}{qB} \left(\frac{s}{R} \right) \frac{p_L^2 + 0.5p_T^2}{p_L} \\
&= \frac{1}{qB} \left(\frac{s}{R} \right) \left(\frac{p}{2} \right) \left(\cos(\theta) + \frac{1}{\cos(\theta)} \right).
\end{aligned} \tag{2.5}$$

Notably, the magnitude of this drift D depends on the particle's charge q , the magnetic field at the axis B , the path length through the curved solenoid s , the radius of curvature R , and the particle's momentum p , where p_L and p_T denote the longitudinal and transverse components with respect to the solenoid's axis [49]. The trajectory's pitch angle θ is defined as $\theta = \tan^{-1} \left(\frac{p_T}{p_L} \right)$.

The dependence on particle charge causes negatively and positively charged particles to drift in opposite directions. Thus, in combination with suitable collimators, the curved solenoid can be utilized for particle selection. Furthermore, the drift's dependence on particle momentum results in a strong momentum dispersion for same-charged particles, as visualized in the left panel of Figure 2.6. An additional compensating dipole field parallel to the drift direction is introduced to counteract the downward drift of negatively charged particles. Its field strength is tuned to keep trajectories of particles with select momenta central in the bending plane, as illustrated in the right panel of Figure 2.6. Hence, it focuses low-momentum muons onto the muon-stopping target, while undesirable high-momentum particles can still be collimated. Figure 2.7 illustrates the effects of particle transport and collimation by comparing particle yields before and after the transport solenoid.

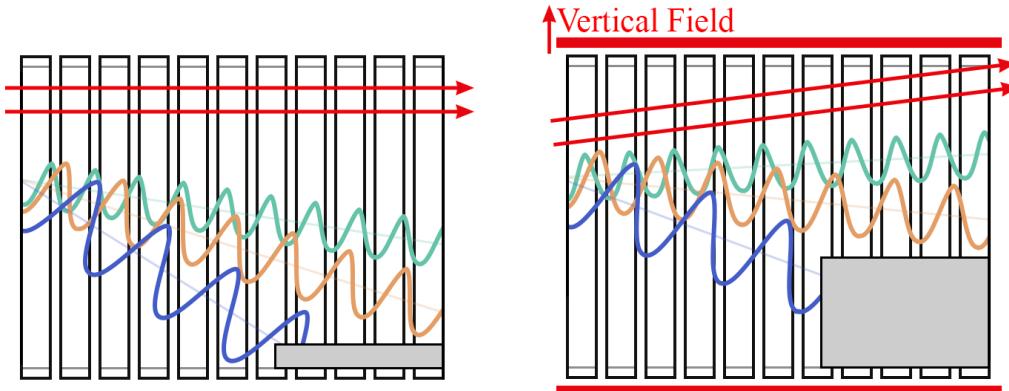


Figure 2.6: Cross-sectional view of the curved solenoid highlighting its dispersion effect on charged particles. Black squares represent the solenoid coils. Red arrows indicate magnetic field lines. Helical trajectories of low (green), medium (orange), and high (blue) momentum particles are overlaid. Their helix center drifts according to Formula 2.5. By utilizing an additional vertical field, particles of select momenta can be centered, while collimators (grey boxes) can be used to suppress the high-momentum particle background.

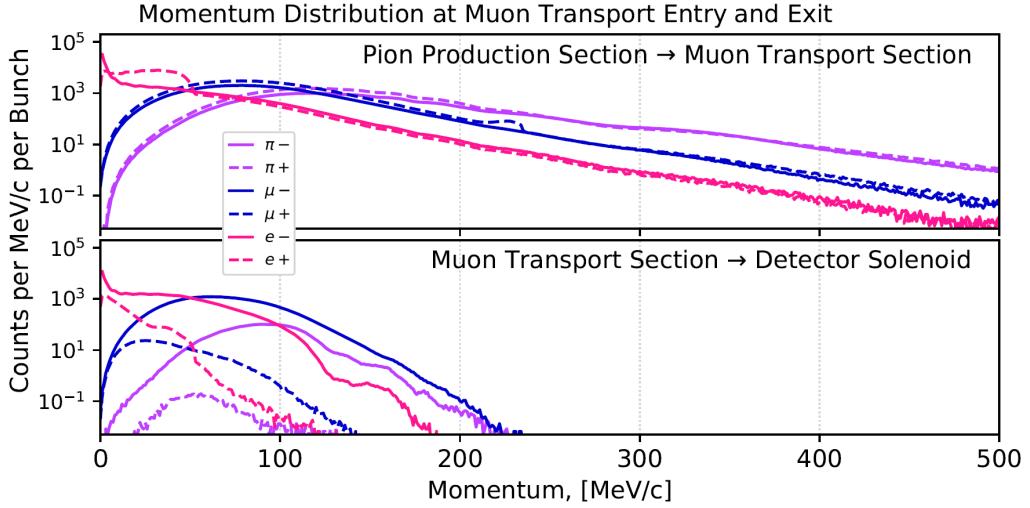


Figure 2.7.: Particle yields before and after passing the bent 90° transport solenoid.

Plots are normalized to one bunch and taken from [61]. While the flux of low-momentum muons is preserved, high-momentum particles and positively charged backgrounds are efficiently filtered away.

2.3.3. Muon-Stopping Target: Shape and Material

As its name suggests, the muon-stopping target's primary objective is to stop muons in the center of the detector system. It consists of 17 thin target-disks, with a radius of 100 mm, thickness of 0.2 mm, and a disk-to-disk distance of 50 mm. The muon-stopping target's configuration and placement, including its impact on detector background, is one of the main topics of this thesis and extensively discussed in subsequent Sections 4 and beyond. For now, this section closely examines a key parameter unchanged throughout the later discussion: the material of the target-disks.

The muon-stopping target material is governed by various constraints. Naturally, mechanical stability, abundance, purity, commercial availability, and correlated material costs must be considered. From a particle physics standpoint, three essential aspects have to be balanced: conversion electron energy, muonic atom lifetime, and muon conversion branching ratio.

As discussed in Section 2.3.1 and shown in Figure 2.4, a bunched proton beam is adopted to suppress the background from prompt beam particles. Detector hits from prompt beam particles are expected to occur for up to approximately 200 ns after the proton beam's arrival. Thus, the lifetime of muonic atoms must exceed this duration for effective background and signal separation. The relationship between the lifetime of muonic atoms and their atomic number Z is given in Figure 2.8. For light nuclei ($Z \leq 7$), the radius of the muon orbit is large, and the overlap with the nucleus is small. Thus, the total lifetime is dominated by

the free muon lifetime of $2.2\ \mu\text{s}$. As Z increases, the probability of finding the muon at the nucleus rises with Z^3 , and the muon capture probability rises with Z^4 . The increase levels off once the muon radius becomes comparable to the size of the nucleus. As the total lifetime simply depends on the sum of decay and capture rate (compare Formula 1.3), it falls off steeply for larger nuclei. Consequently, a sufficient background and signal separation in COMET requires the use of light elements ($Z < 26$).

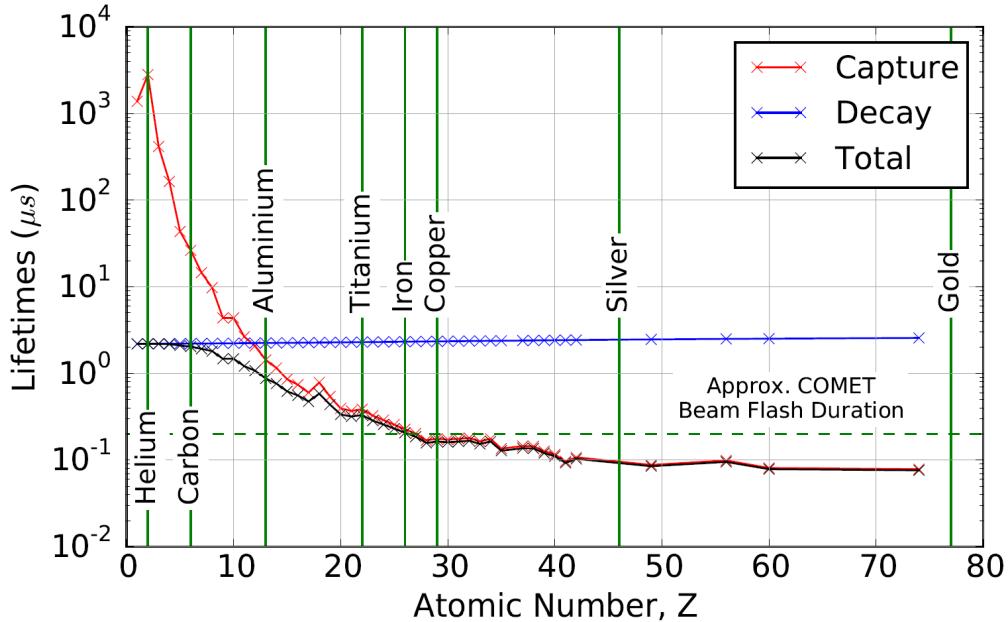


Figure 2.8.: Relationship between the lifetime of muonic atoms and the atomic number Z . The reciprocal of the total lifetime equals the reciprocal sum of the decay and capture lifetimes (compare Formula 1.3 in Section 1.1.2). The plot was adapted from [59] and shows capture and decay rates, as taken from Geant4 [60]. The approximate duration of detector hits from prompt beam particles (denoted beam flash) at 200 ns is included. Notably, lifetimes of elements heavier than iron ($Z > 26$) fall below this value and thus are unsuitable for signal-background suppression utilizing the bunched beam structure.

The energy of the conversion electron E_e can be derived from Formula 2.2. It is in the order of the muon mass but slightly reduced due to the binding of the muon and nuclear recoil from the conversion. Both corrections depend on the atomic number Z . The relationship between E_e and Z is shown in Figure 2.9. Background from DIO can reach up to this conversion electron energy, making higher values of E_e preferable to avoid signal contamination from DIO in materials with higher endpoint energies. Consequently, again, lighter elements are favored as target material, particularly in the range between $Z = 5$ and $Z = 12$.

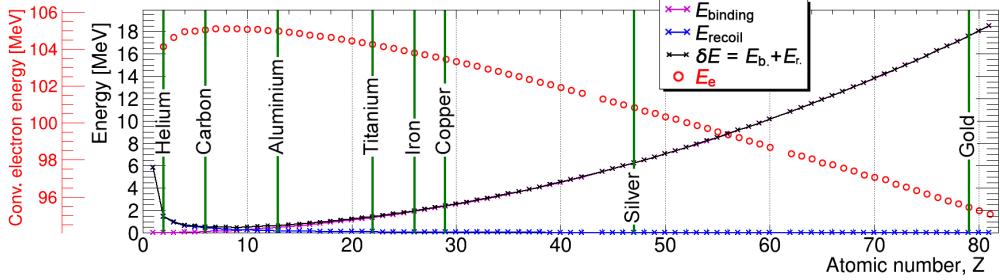


Figure 2.9.: Dependency of the conversion electron energy from the atomic number Z . Values of E_{binding} and E_{recoil} are calculated according to Formula 2.2. Values for the conversion electron energy E_e correspond to the red y-axis and are taken from [16] to additionally include small corrections due to vacuum polarization effects.

Lastly, theoretical predictions of the muon conversion branching ratio for different materials and models are shown in Figure 2.10. Clearly, the plot indicates an increase in conversion rates toward heavier nuclei. A particularly noteworthy feature is the increased discriminating power between shown models for larger Z . However, constraints from DIO contamination and especially lifetime restrictions outweigh such potential improvements from heavier elements.

In conclusion, lighter targets appear better suited for COMET. The adopted material for the muon-stopping target is aluminium ($Z = 13$). Apart from the material's high abundance and the possibility of achieving purities of up to 99.999 %, the lifetime of muonic aluminium is 864 ns, aligning with the planned bunch structure of the proton beam. Conversion electrons from aluminium have an energy of $E_{\mu e, \text{Al}} = 104.97 \text{ MeV}$. Hence, DIO electrons from heavier elements like iron ($E_{\mu e, \text{Fe}} = 103.82 \text{ MeV}$) or lead ($E_{\mu e, \text{Pb}} = 94.89 \text{ MeV}$), used for the construction and shielding of the CyDet detector parts, are well separated from the signal. However, the elements oxygen ($E_{\mu e, \text{O}} = 105.11 \text{ MeV}$) and nitrogen ($E_{\mu e, \text{N}} = 105.09 \text{ MeV}$) must be avoided, thus the muon-stopping target area will utilize a helium atmosphere instead of air.

To bring this discussion to a close, it should be noted that the conversion rate shown in Figure 2.10 heavily depends on the target nucleus, specifically its spin, mass number, and proton-to-neutron ratio [62–64]. If μ -e conversion in aluminium were observed, comparisons with different target materials will be necessary to distinguish underlying BSM models. Complementary to data from the DeeMe experiment, lighter elements such as carbon or silicon carbide could be considered as muon targets. However, these elements have a high probability of DIO, a long lifetime of the muonic atom, and a reduced stopping power for muons, which would require significant changes in COMET's design for only a small difference in the predicted conversion rate. A measurement using titanium targets ($Z = 22$, light nucleus without spin, $E_{\mu e, \text{Ti}} = 104.26 \text{ MeV}$, $\tau_{\mu^-} = 329 \text{ ns}$) is often discussed and

might be a feasible follow-up step for COMET. However, the suppression of prompt beam background and the higher DIO contamination must be carefully investigated for such a heavy element. Even heavier targets like gold ($Z = 79$) or lead ($Z = 82$) are completely disfavored by the current COMET design, as any conversion signal would be drowned out by prompt beam background.

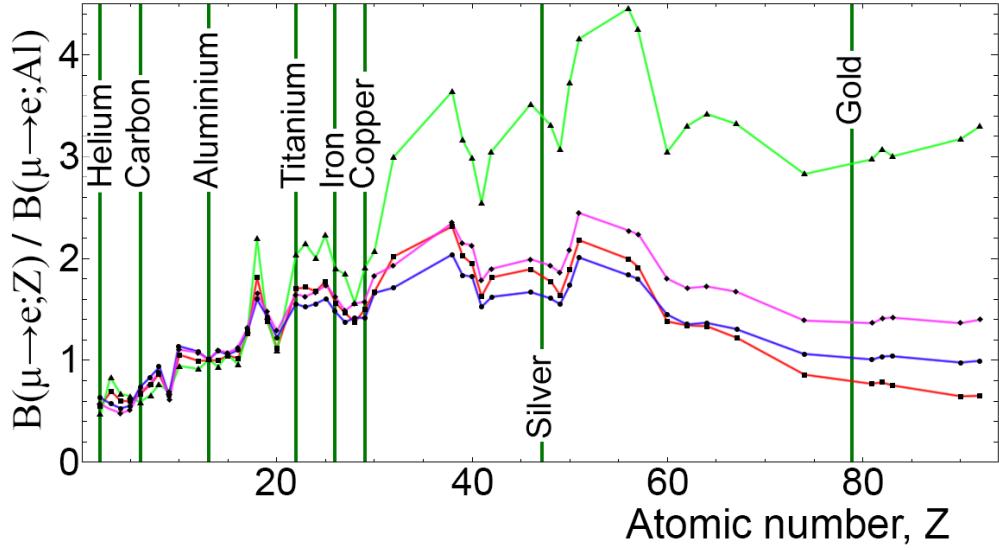


Figure 2.10.: Target dependence of the $\mu^- \rightarrow e^-$ conversion rate in different single-operator dominance models. The plot is adapted from [65]. The conversion rate is normalized to the rate in aluminium ($Z = 13$) and plotted versus the atomic number Z for four theoretical models: Dipole interaction (blue), Scalar interaction (red), Vector interaction with photons (magenta), vector interaction with Z bosons (green).

3. The COMET Simulation: ICEDUST

As new high-energy particle physics experiments continue to grow in scale, conventional ‘build and error’ approaches for their design become increasingly difficult due to budget constraints or interdependences between unfinished components. Thus, a modern solution for the development, analysis, and optimization of such experiments often involves the use of so-called Monte Carlo (MC) simulations.

The MC framework developed by the COMET collaboration is called the “Integrated COMET Experiment Data User Software Toolkit”, in short ICEDUST. The software relies in large parts on established MC simulation code bases, in particular Geant4 [60] libraries, to manage the experiment’s geometry, magnetic field, particle interactions, and particle tracking. A snapshot of the simulated geometry of COMET Phase-I in the conversion-searching design using CyDet is illustrated in Figure 3.1. ICEDUST serves as a fundamental cornerstone in guiding optimizations throughout this thesis. Therefore, this chapter provides a brief overview of its most relevant components. Particular focus is given to details required to interpret and compare this thesis with future results. For a more general description of the simulation see [59] or [50].

3.1. Simulating the Primary Beam

The first particle created in the simulation is called the “primary”. In a standard ICEDUST run, this refers to the initial proton that impinges on the pion production target. Particles created from interactions of the primary are called “secondaries”, and the timing of these secondary particles is determined with respect to the POT time. The primary proton’s energy and position distributions are taken from the COMET Phase-I TDR [43]. The primary and all its secondaries are contained within so-called “events”. Each event, and thus each proton, is simulated independently. The bunched structure containing 16×10^6 POT per bunch is created post-simulation, where the timing of each event is shifted slightly to mimic a more realistic timing distribution. However, no correlations between events, such as primary protons interacting with each other, can be included.

While starting with the initial proton is the most accurate, it has one major constraint: computing time. On average, each simulated proton requires 2.7 s. This thesis primarily focuses on the muon-stopping target and the CyDet region. However, only a fraction of

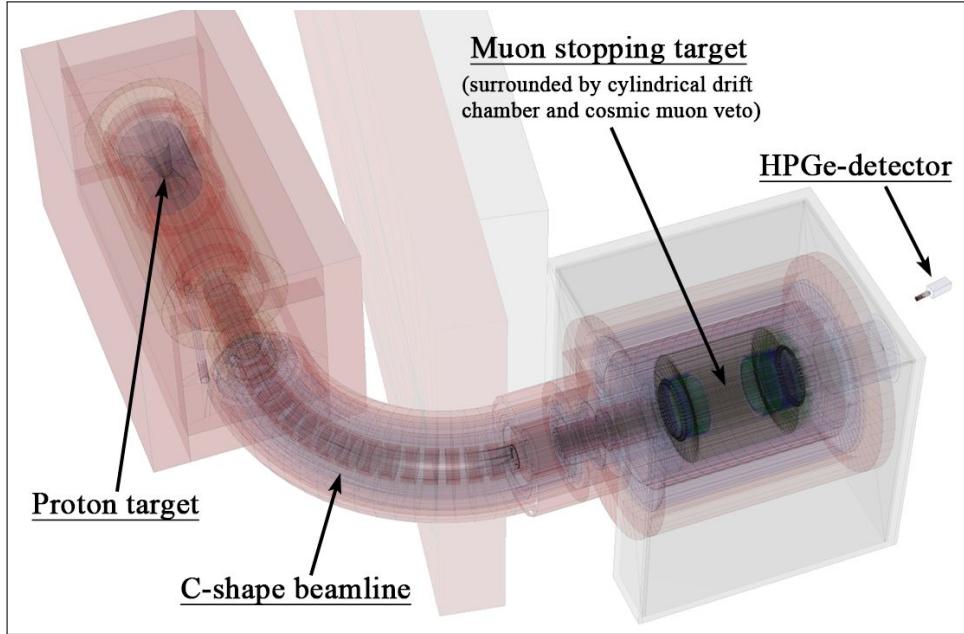


Figure 3.1.: COMET Phase-I geometry as used in ICEDUST. Although present in the simulation, the experimental hall is excluded in the picture to enhance clarity. Named components are described in detail throughout Section 2.

5×10^{-3} muons per POT successfully reaches this area. This translates to each muon requiring roughly 200 simulated protons or approximately 9 min of computing time per muon. Of all muons arriving at CyDet, only about 1/10 can be stopped within the muon-stopping target, effectively increasing the required time for stopped muon studies by an additional factor of 10. As a result, very low statistics of 10 000 muons at the detector region already require about 1500 computing hours (CPUh).

To solve this high CPUh problem, the simulation process is divided into two stages. During the first stage, the region around the pion production target is simulated, including the primary proton beam impinging on the target; the creation of pions, muons, and other secondaries; and the tracking of all particles through the experiment. The majority of secondaries remain within the vicinity of the pion production section, where they are absorbed by the beam dump, the ground, or the shielding surrounding the proton target. Only a small fraction of particles is collected by the pion capture solenoid and directed along the beamline toward the detector region. An even smaller fraction, primarily consisting of neutrons, can traverse the walls, floors, and shieldings to directly reach the detector region. Now, instead of finishing the simulation around CyDet, the particle tracks are terminated at the entry into the detector region, and their Particle Identification (PID), momentum, and position are recorded onto the disk.

In the second stage, the recorded data serves as input for a new ICEDUST run. Only events with particles reaching the CyDet were previously recorded. Thus, restarting the simulation from these input files, known as *resampling*, avoids the simulation of uninteresting events where no particle reached the CyDet. Moreover, for additional optimization, filters may be applied to exclusively restart specific particle types, like muons. Since the particles recorded in step one of this procedure have not yet traversed the detector region, one set of input files can be used for all simulations of the CyDet region, enabling fast testing of varying muon-stopping target or CyDet geometries. As a result, the time required to simulate 50 000 POT can be reduced from 37.5 h to 15 min, representing a significant reduction in CPUh by more than 99 %.

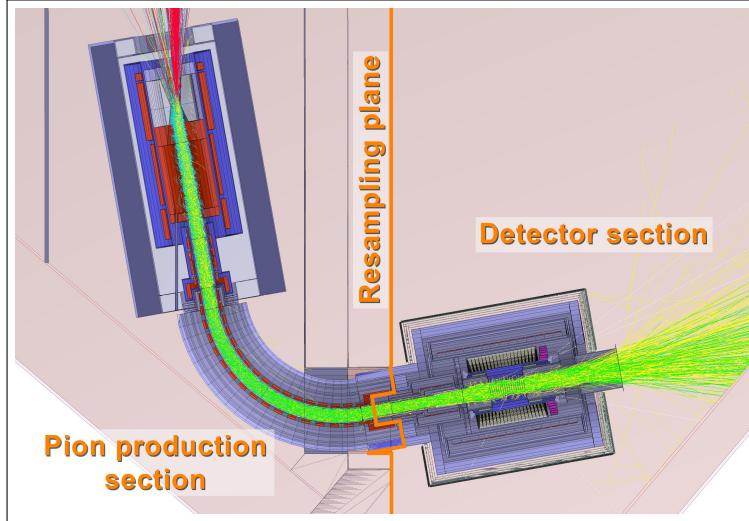


Figure 3.2.: Cutaway top-down view of the ICEDUST geometry. The bold orange line highlights the resampling plane, separating the pion-production (left) and detector (right) sections. Charged particle trajectories (green and yellow lines) corresponding to 1.6×10^5 events of MC5 are overlaid.

The described method was implemented in a joint effort with colleagues from the collaboration. Figure 3.2 presents a top-down cutaway view of the ICEDUST geometry, showcasing the employed *resampling plane* where particles are recorded or restarted. Of course, one cannot avoid simulating the initial proton beam to record interesting particles at least once. Thus, a mass simulation of particles in the upstream region containing the proton production target commenced in early 2020 [50]. This so-called Monte Carlo dataset 5 (MC5) large-scale data production was performed at the French National Institute of Nuclear and Particle Physics Computing Centre (CC-IN2P3) and resulted in data files corresponding to 1×10^9 POT or 62.5 bunches with 16×10^6 POT/bunch. As an example, 1.6×10^5 of those events (1 % of a bunch) are overlaid in Figure 3.2.

Most of the muon beam and background simulations throughout Sections 4, 5, and 6 are performed with MC5 files as input. It is important to note that the approach has two major limitations: First, the pion production target and beam transport system are fixed to one select configuration and can not be easily adjusted. A new large-scale data production, called MC6, including updates on these components, is currently underway, but such changes are not yet included in this thesis¹. Second, primary beam statistics are limited to only 62.5 bunches. In theory, each MC5 input file could be reused multiple times in the same run to artificially increase statistics. If the distance between a point of interest, usually a detector component or the muon-stopping target, and the resampled start position is large, the statistical nature of interactions along a particle track is assumed to be sufficient to randomize the data. However, for small distances or particles with low interaction probability, such as neutrons, resampling multiple times could lead to biases that are difficult to control. Consequently, this thesis refrained from resampling the same particle information multiple times.

3.2. Direct Simulation of Conversion Electrons and X-rays

Studies of secondary particles that directly depend on the primary proton beam, like muons stopped in target-disks or beam particles causing detector background, are performed using MC5 data files as input. However, using these MC5 files to study conversion electrons or muonic X-rays is impractical. Although ICEDUST includes customized muon-to-electron conversion and muonic X-ray physics to generate such secondaries when simulating the primary muon beam (and stops) from MC5 inputs (see Section 3.4), they both follow only at the very end of the simulation chain. Hence, a more practical approach is to further separate the simulation of muon-dependent secondaries to archive high statistics of conversion electrons or muonic X-rays without requiring a comparatively CPUh intensive re-simulation of the muon beam itself.

The method used to directly simulate muon-stopping dependent secondaries is similar to the resampling plane approach discussed in the last section. First, the distribution of muons stopped inside the muon-stopping target is recorded. Next, the position distribution is used as input to generate secondaries. Unlike before, however, the information if a muonic X-ray, DIO or conversion electron was emitted after the muon stopped in the initial ‘recording run’ is discarded. Instead, each recorded position is assumed to result in a particle of interest, drastically increasing the available statistics.

Each recorded starting position is assigned a randomized uniform momentum direction and paired with a starting energy. For conversion electrons, the starting energy is material-dependent but constant and, in the case of aluminium, lies at $E_{\mu e, Al} = 104.97$ MeV (see Section 2.3.3). Energies and branching ratios for muonic X-rays are taken from [66] and [67] and further details are discussed in Section 6. DIO electron energy distributions require input

¹Preliminary comparisons between MC5 and MC6 particle distributions show no major differences, thus no negative consequences are expected from this decision.

from theoretical calculations. Figure 1.1 shown as part of the introduction in Section 1.1.2, utilizes results taken from [17]. The performed DIO studies in Section 4.9 will use newer results presented in [16], where the DIO electron energy is resampled according to the function:

$$\frac{1}{\Gamma_0} \frac{d\Gamma}{dE_e} \Big|_{E_e \sim E'_{\text{end}}} = B \cdot E'_{\text{end}}^5 \left(1 - \frac{E_e}{E'_{\text{end}}}\right)^{5+\delta} \quad (3.1)$$

with the correction factor $\delta \approx 0.023$ stemming from soft photon radiation and a shifted endpoint energy, the material-dependent conversion electron endpoint energy E'_{end} , and the coefficient B , which depends on several overlap integrals over products of wavefunctions. Parameters E'_{end} and B are also presented and taken from [16]. The exact value of B depends on the assumed electric-charge distribution of nuclei. To obtain a conservative estimation of DIO background events, this study always assumed the largest provided B value for each element, which corresponds to the highest possible DIO background.

3.3. Details of the Geometry and Magnetic Field

COMET, and consequently ICEDUST, undergo constant development, leading to continuous adjustments of the simulated geometry to align with ongoing updates on components. Thus, special attention was paid to ensuring that all simulations were performed with an identical set of parameters for all parts throughout this thesis. Whenever necessary, older results from the start of this thesis were revised to facilitate inter-comparison between all values. This section presents details for select components of particular significance in this study. Not mentioned components can be presumed to follow the general discussion presented in Chapter 2. Note that the muon-stopping target, its holding structure, as well as the High Purity Germanium (HPGe) detector positioned at the DS side of CyDet are actively investigated in dedicated Sections 4, 5 and 6, respectively.

Starting with the general CyDet system, recent developments at the detector area focus on implementing additional shielding around its connection to the beamline and an improved design of the DS beam duct. Newer designs of this beam duct employ only a thin, flat sheet of steel or similar materials at the end of CyDet to reduce backscattering of beam particles. However, the design utilized in this thesis still consists of a tube-like structure, as shown on the right side in Figure 3.3a. The length of this steel tube was set to 950 mm, with an inner radius of 451 mm and a thickness of 10 mm. The end of this tube consists of a 2 mm thick steel plate with a radius of 451 mm.

The implementation of additional shielding at the US side required moving the entire CyDet system further DS. The change was implemented in ICEDUST as part of a major geometry update performed in early 2023, which also updated the surrounding solenoid, adjusted materials used, etc. However, these updates are not yet implemented for this thesis.

Most notably, the material of the cryostats around the entrance is still set to steel instead of aluminium. The center point of the CyDet geometry is located at (6410, 0, 7650) mm with respect to the pion production target, and no additional shielding for CDC or the cosmic ray veto is implemented yet. The gas inside CyDet surrounding the muon-stopping target is set to helium at exactly one atmosphere.

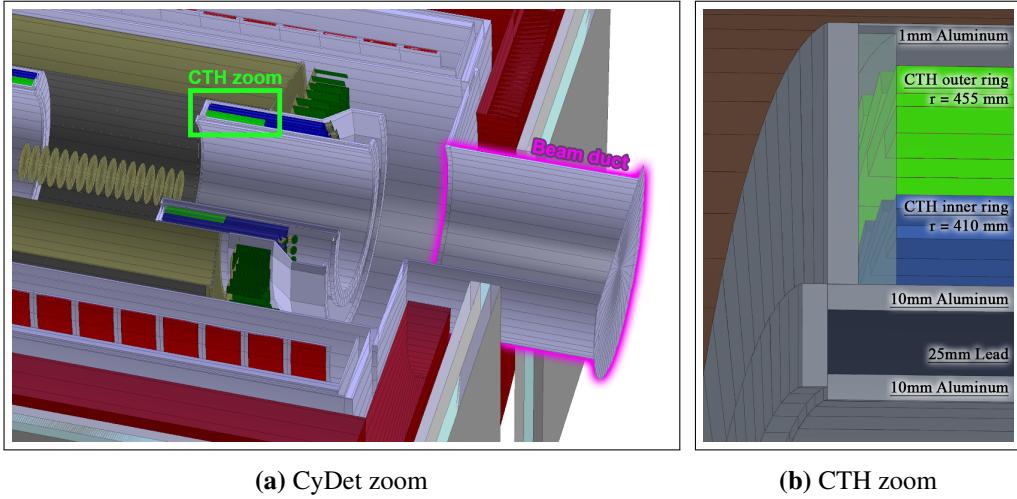


Figure 3.3.: Close up on select ICEDUST geometry elements to emphasize the configuration used throughout this thesis. Left: Cut open CyDet region with focus on the pipe-like beam duct geometry (pink) shown on the right side. For exact parameter values see text. Right: Zoom onto the CTH geometry (green square), highlighting the shielding surrounding the two scintillator counter rings.

The CDC detector was completed in early 2016; thus, no further updates on this component had to be implemented. However, the CTH detector underwent continuous updates to its design throughout the last years and over the lifetime of this thesis, significantly impacting the expected signal acceptance and background rates. The initial setup consisting of a combination of Cherenkov and scintillation detectors was overruled in favor of a scintillator-only configuration, and parameters regarding size, angle, and the number of scintillators are still under active development. The CTH configuration used for this thesis is pictured in Figure 3.3b. It consists of two rings with a total of 64 polyvinyltoluene scintillator counters. The inner ring is positioned at a radius of 410 mm, and utilizes counters with a length of 360 mm, a width of 75 mm, and a thickness of 5 mm. Counters in the outer ring are positioned at a radius of 455 mm, have a slightly shorter length of 330 mm, but an increased width of 85 mm and a thickness of 10 mm. To optimize trigger conditions (see discussion in Section 4.2.3), counters are tilted by 15° and 18.75° for the inner and outer rings, respectively. Each scintillator is wrapped in a 0.026 mm thick mylar foil to reflect stray light. The inner

shielding (smaller radius) of the counters consists of 25 mm of lead surrounded by 10 mm of aluminium on both sides to absorb low-energy photons and electrons originating from the muon-stopping target. The shielding at the outer side (large radius, adjacent to CDC) consists of a thin 1 mm layer to allow for unimpeded detection of signal electrons after they traversed the CDC.

Lastly, charged particles are subject to magnetic fields when traversing the experiment. The field strength at each space point is calculated separately from the simulation and stored in collaboration-wide shared “fieldmaps”. The fieldmaps for the pion production section and detector region were calculated by collaborators using the Opera 3D finite-element-analysis software. The fieldmap for the muon beamline dipoles was provided by the manufacturer, Toshiba².

3.4. Handling of Physics in ICEDUST

The physics in ICEDUST are built on top of the standard Geant4 procedure and extended where applicable. While all required particles and their associated processes³ can be specified separately, this approach is prone to errors and may introduce biases that are difficult to identify. Thus, Geant4 simplifies this process by providing a range of pre-defined reference “PhysicsLists”. These PhysicsLists define the physics environment of the MC simulation and are extensively tested and validated.

In ICEDUST, hadronic interactions like capture, fission, or scatter processes are handled by the QGSP_BERT_HP PhysicsList. The models contained in QGSP_BERT are recommended for collider and high-energy particle physics up to the order of 10 TeV [68]. To improve estimates of neutron backgrounds, the HP extension additionally enables the use of High Precision neutron models and cross-sections, which are based on the G4NDL database referring to the internationally evaluated JEFF-3.3 library [69]. Moreover, the extension activates G4RadioactiveDecay, used to simulate the decay of radioactive nuclei by α , β^+ and β^- emission and by electron capture based on data taken from the Evaluated Nuclear Structure Data File (ENSDF) [70].

Electromagnetic physics (EM) are handled by the G4EmStandardPhysics_option4 PhysicsList. This PhysicsList aims to deliver the most accurate EM physics, selecting the best combination of models per particle type and energy range with the drawback of a slight reduction in simulation speed. Additionally, the G4EmExtraPhysics PhysicsList is utilized to enable inelastic nuclear scattering for gammas and muons, as well as synchrotron radiation for electrons and positrons, as these processes are not included by default. Further meticulous description of each underlying model is outside the scope of this thesis. However, the interested reader is referred to the publications of the Geant4 collaboration regarding their physics model references [71].

²For internal reference, this thesis employs all fieldmaps in version ‘150630_Phase-I_opt’.

³Processes refer to and describe each particle’s interaction with matter. Processes are required to change particle properties, to produce secondaries, or deposit energy.

Lastly, a few select changes to muon physics are performed because of their elevated importance in COMET. First, the default Geant4 Nuclear Muon Capture (NMC) model is replaced with measured results from the AlCap experiment, which include the rate and spectrum of particles emitted after NMC on aluminium [72]. Second, the energy spectrum of electrons after DIO in aluminium is modified to account for nuclear recoil effects as calculated by Czarnecki et al. [17]. After this change, DIO electrons can have energies up to the conversion energy (see Section 1.1.2) and resemble a major source of physics background (see Section 2.3). Finally, this thesis will develop a special treatment for the muonic X-ray cascade created after a muon is stopped inside aluminium. The details of this are discussed in the dedicated Section 6.

Table 3.1.: Versions of the used external Geant4 datasets. Further information on these sources is provided in publication [60].

Dataset	Version
G4ABLA	3.0
G4EMLOW	6.41
G4ENSDFSTATE	1.0
G4NDL	4.5
G4NEUTRONXS	1.4
G4PII	1.3
G4SAIDDATA	1.1
PhotonEvaporation	3.1
RadioactiveDecay	4.2
RealSurface	1.0

The used version of Geant4, and thus all discussed PhysicsLists, was ‘10.01.p03’. The exact versions of the accompanying external libraries containing measured data for the underlying processes, such as radioactive decay or neutron cross-section, are listed in Table 3.1. At the time of writing, the most recent Geant4 version was updated to ‘11.1.1’, and extensive tests are put forward by the collaboration to verify the modified physics and cross-check the results between versions. However, this study refrained from updating Geant4 to the newest version to guarantee consistency with previous results⁴.

⁴Although newer Geant4 versions were already available at the beginning of this thesis, updating to them was not possible due to ICEDUST’s reliance on an outdated build system called CMT [73]. A significant part of the early stages of this thesis was the single-handed transition to a newer build system called CMake [74]. Despite the considerable time investment, documentation of this process is excluded from the thesis due to its minor relevance to the physics discussion. However, detailed explanations are provided in the collaboration internal documentation on GitLab.

3.5. General Optimization Strategy

MC simulations enable flexible customization and on-the-fly adjustments to geometry and detector. The difficulty regarding active detector development and a multitude of adjustments, however, is that future results may change past assumptions or parameters, rendering older results quickly outdated. Thus, throughout this work, many previously assumed-to-be-fixed results had to be overruled, looked at, and optimized again after incorporating new insights. To provide a red line through all these changes, later sections throughout this work are sorted by logic, topic, and results rather than the order of their investigation.

Within a MC simulation, there are not many restraints on what could be done. As further elaborated in Section 4.1, optimizations of the muon-stopping target, for example, only follow two basic principles, that is: *maximizing number of muons stopped* and *minimizing conversion electron energy loss*. In reality, not only the signal strength has to be maximized, but also the background rate in the detectors, along with more mundane factors such as material and production costs, must be considered. Nevertheless, the possibilities for optimizations are endless.

To obtain a tangible starting point, the initial configuration for this work builds on top of COMET’s 2014 TDR design. As a consequence, optimizations of the muon-stopping target are mainly centered around thin, circular stopping-disks. The justification to limit the study around this so-called *default* configuration is briefly discussed in Section 4.4. However, even such a seemingly simple design like several thin disks still provides a multitude of adjustable parameters, namely total number of disks; disk thickness; disk radius; disk positioning; and disk spacing. Each parameter will be explored in a separate section and combined to a greater picture in the end.

For all these optimizations, computing time plays the major limiting role, the biggest crux being the severe correlation between all studied parameters. Even if individual target-disk optimizations are disregarded, an exhaustive brute-force optimization study would need to scan through an at least four-dimensional parameter space. Ten changes of each parameter would translate to a total of 10 000 ICEDUST runs. Even after applying computing optimizations like MC5 resampling, each of those runs would take optimistically one day, and in total such a rather small study would require around 27 yr of CPUh – an impossible feat, even with modern computing clusters.

Instead of brute-force scanning through the parameter space, one might consider more modern optimization techniques like machine learning algorithms. However, machine learning suffers from a similar problem: the requirement of a large amount of input data to train the algorithm. Even worse, with the experiment and the simulation being in active development, muon-stopping target independent parameters of other experimental parts are regularly updated. Every change invokes the need for a new set of training data, amplifying the problem. Furthermore, a well-known downside of machine learning is the black-box-like behavior of the method. Parameters like the magnetic field strength might slightly differ from the archived real-life values. Adjusting machine learning results to later determined

real-life constraints may prove challenging, and one risks becoming overly dependent on simulation specifications without the flexibility of smaller adjustments.

In summary, computing time plays a major limiting role and, as a consequence, rules out a basic brute-force scanning or machine learning approach. Thus, some other form of intelligent scanning of the parameter space is necessary. Blindly changing a few parameters while disregarding others would introduce the risks of hard-to-control bias. Therefore, two approaches are adopted in this work:

- Heavily simplified stand-alone calculations/simulations
- Building intuition and an iterative approach

Simplified stand-alone calculations/simulations in this context refer to a rather short and highly specialized piece of software. Instead of following every created particle to its end and recording every detector hit, these simplified stand-alone calculations/simulations will focus on only one particle and detector part. As an example, the transportation of photons through some material is simplified by ignoring the production of secondary particles and energy depositions along the way. Instead, one only asks a single specific question: How many of N_0 initial photons reach the place Y after Z amount of material is traversed? This question can easily be answered with the Beer-Lambert law: $N(z) = N_0 \cdot \exp(-\lambda \cdot z)$ where z is the path length through the medium, and λ the material-dependent attenuation coefficient. In this example, the required time to study one photon is reduced from around 1 s to 0.6 μ s. While this effect might seem negligible, one can easily see how it influences simulations that require statistics in the order of $\mathcal{O}(10^{30})$.

Building intuition in this context refers to the difference between pure informatics and physics. While similar to informatics, much of the analysis is performed by looking at ones and zeroes, the outcome at the end can, should, and is interpreted from a physicist's perspective. The approach includes extrapolating results, testing assumptions before embarking on a full simulation campaign, or selectively emphasizing certain aspects of a parameter while disregarding others. Therefore, it is important to emphasize that the efforts within this study are not solely focused on optimization; rather, the primary focus lies in developing an in-depth comprehension of the underlying physics. The insights gained from this process are then utilized to re-evaluate earlier results and make further adjustments. This iterative process of knowledge acquisition and refinement is fundamental to the scientific approach taken in this work.

4. The Muon-Stopping Target

For coherent muon-to-electron conversion to occur, an integral ingredient is to stop muons within the Coulomb field of an atomic nucleus. This task is taken over by the muon-stopping target. It is evident that optimizations on such a central component can easily lead to large improvements on COMET's overall sensitivity.

This chapter will discuss various ideas for improvements and is divided into five main parts. As a starting point, Section 4.1 explores the current muon-stopping target design and presents the motivation for further investigations. Next, relevant parameters for optimization, their definitions; computation; and limitations are introduced in Section 4.2. Utilizing the discussed procedures, an in-depth analysis of each core parameter is presented throughout Sections 4.4 to 4.9. The results are combined and estimates for improved SES values are calculated in Section 4.10. Lastly, the production of muon-stopping targets is presented in Section 4.11.

4.1. Introduction: Status Quo and Motivation

Two competing parameters must be carefully balanced when optimizing the muon-stopping target:

1. Maximizing the number of stopped muons.
2. Minimizing the energy loss of conversion electrons.

Both parameters are crucial for observing a potential conversion electron signal, but they inversely depend on the total amount of muon-stopping target material positioned inside the muon beam.

More material is beneficial to increase the number of stopped muons. Before a muon can be stopped in an atomic orbit around the nucleus, it must be slowed down to near-thermal energies. The moderation of muons in mediums has been well studied [11, 75, 76], showing that they can be treated according to the conventional methods applicable to fast heavy particles. In particular, this means longer paths through the aluminium target-disks, which have a higher electron density than the surrounding helium gas, directly translates to higher energy loss and more stopped muons over the experiment runtime.

To minimize the energy loss of conversion electrons, less material should be employed. The conversion electrons' characteristic high energy is essential for separating and discarding lower-energy backgrounds. However, similar to muons, electrons lose more energy when

traversing materials with a high electron density. Hence, it is crucial that the conversion electron's path within the muon-stopping target (disks) is as short as possible.

The currently adopted design to balance these two competing parameters consists of 17 thin aluminium disks ordered in a tube-like structure, as shown in Figure 4.1. Its parameters are summarized in Table 4.1. Section 4.4 will provide the justification for this general shape, and the reasoning for selecting aluminium as the stopping material was provided in Section 2.3.3. The target-disks are placed in the center of CyDet at $r = 0$ mm. Along the beam, they are positioned over the full distance between US and DS CTH, which translates to a total length of $L = 800$ mm with a disk every 50 mm. The COMET Phase-I TDR [43] states that for every POT, about 4.7×10^{-4} muons can be successfully stopped. While this number may seem small without context, COMET will utilize 2.5×10^{12} protons/s, which translates to 1.2×10^9 muons/s, and with a total number of 3.2×10^{19} protons amounts to 1.5×10^{16} stopped muons.

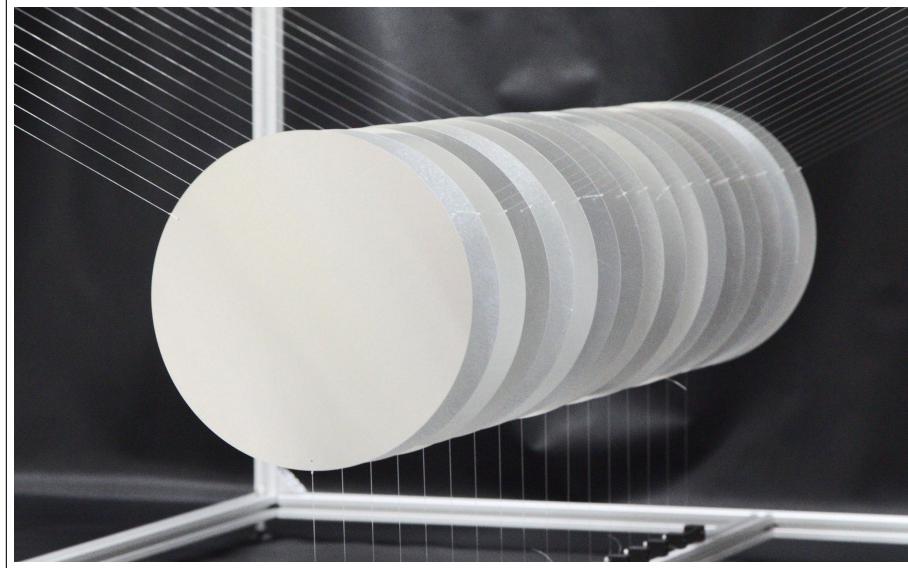


Figure 4.1.: Mock-up of the default muon-stopping target design produced in 2017, as taken from [77]. Grey circular aluminium target-disks are suspended on nylon strings within a temporary holding structure.

The motivation to further optimize the muon-stopping target is presented in Figure 4.2. The figure shows the momentum distribution of the total muon flux that reaches the muon-stopping target. Overlaid on this distribution is the fraction of successfully stopped muons, color-coded based on their respective stopping disks. The shown total muon flux peaks at about 60 MeV/c. In contrast, the distribution of the number of stopped muons for the 17×0.2 mm disk configuration peaks already at approximately 40 MeV/c and barely reaches up to 50 MeV/c.

Table 4.1.: Parameters of the default muon-stopping target design in COMET Phase-I, as taken from [43].

Item	
Material	Aluminium
Shape	Flat disk
Radius	100 mm
Thickness	200 μm
Number of disks	17
Disk spacing	50 mm

To put this into context, only approximately 15 % of all muons arriving at the muon-stopping target traverse enough material to be stopped. As an example, if it were possible to stop all muons with momenta between 50 MeV/c and 60 MeV/c, the total number of stopped muons would increase by approximately 112 %. The situation is somewhat better for muons with momenta less than 40 MeV/c. However, even for these low-momentum muons, about 22 % cannot be stopped, as indicated by the small grey band above the stopped-muon histogram stack.

In conclusion, dedicated design changes to the muon-stopping target have the potential to improve the number of stopped muons by approximately a factor of two. Since this parameter contributes linearly to the SES (see Formula 2.4 of Section 2.2.2), such improvements could double COMET’s sensitivity or, if the target sensitivity is kept constant, cut the necessary measuring time in half.

4.2. Methodology

The primary objective of COMET is to measure CLFV with maximum precision. In the event that no signal is detected, the experiment’s performance will be assessed based on its final limit, as determined from the SES. Consequently, when optimizing any component within COMET, the primary focus is directed toward enhancing the SES.

As introduced in Section 2.2.2, the SES consists of three key parameters:

1. The branching ratio for muon nuclear capture, denoted as $\mathcal{B}_{\text{capture}}$
2. The total number of stopped muons, represented as N_{μ}^{stopped}
3. The detector’s signal acceptance of conversion electrons, denoted as $A_{\mu \rightarrow e}$

The muon nuclear capture branching ratio is constant for a fixed choice of target material. However, the total number of stopped muons and the detector’s signal acceptance are strongly influenced by the exact configuration of the muon-stopping target.

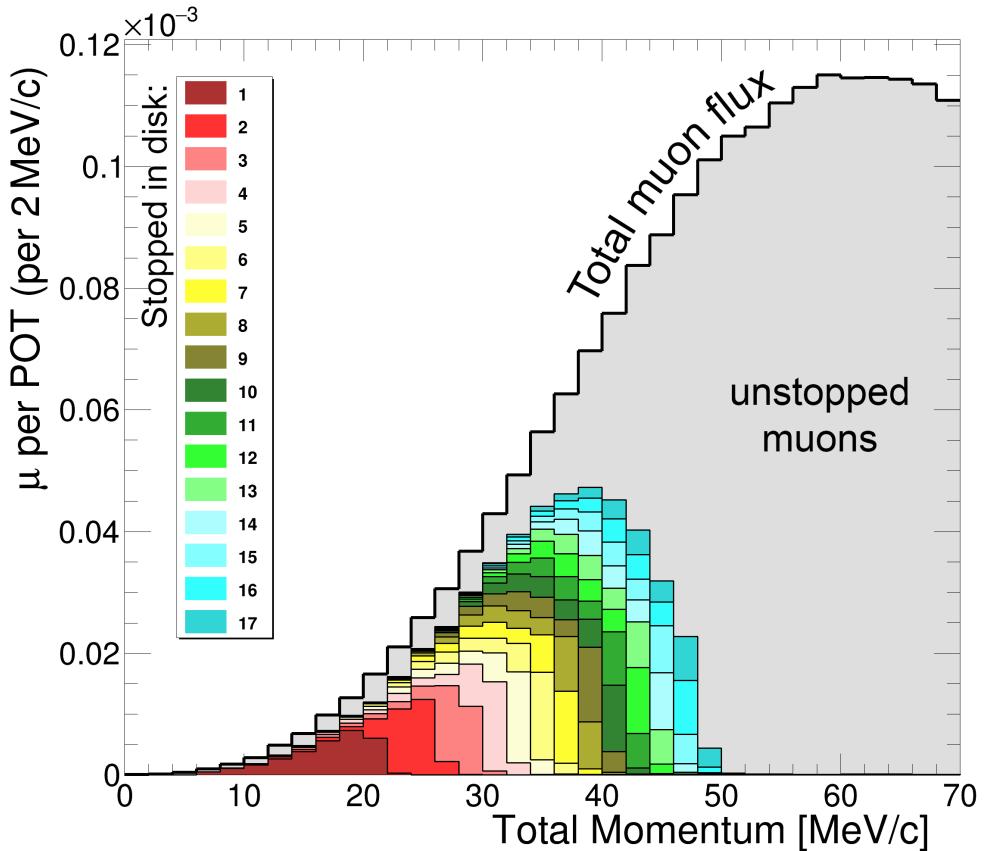


Figure 4.2.: Momentum distribution of muons reaching the muon-stopping target and the fraction of successfully stopped muons. Stopped muons are color-coded according to their stopping-disk. The data correspond to 60 bunches at 16×10^6 POT. As can be seen, only muons of up to 50 MeV/c can be stopped within the 17×0.2 mm target-disks and the default design falls short of covering the 60 MeV/c peak of the arriving muon flux. See text for further discussion.

As a consequence, within the scope of this thesis, N_{μ}^{stopped} and $A_{\mu \rightarrow e}$ will serve as placeholders for assessing COMET's performance. To prevent redundancy, this section provides a detailed definition of these parameters, highlights their dependencies, and introduces the analysis procedures used to derive them from the ICEDUST outputs. Subsequently, CDC and CTH detector backgrounds are introduced due to their limiting role throughout the optimizations. Particular emphasis is placed on addressing constraints imposed by the MC simulation. At the end, the treatment of uncertainties based on simulation results will be discussed.

4.2.1. Interpreting Plots: The Centered Coordinate System

Plots will play a fundamental role in guiding discussions throughout the subsequent sections. With the focus on the CyDet, or more precisely, the CDC, a centered coordinate system is defined around this region. Unless specified otherwise, this coordinate system will be consistently used for all plots in this thesis.

The coordinate system is illustrated in Figure 4.3 and Figure 4.4. Its point of origin corresponds to the center of the CyDet, and is denoted as $(x, y, z) = (0, 0, 0)$ mm. The position aligns with the origin of the middle disk (disk number 9) of the default muon-stopping target configuration. Similarly, the centers of the first disk (disk number 1, US side) and the last disks (disk number 17, DS side) are located at $(0, 0, -400)$ mm and $(0, 0, 400)$ mm, respectively. Both figures include an example of a muon trajectory and a monitor plane used to study particle x - y -distribution (see Section 4.2.2). Note that the z -axis of the coordinate system is chosen to align with the direction of the beam. In consequence, when examining x - y -distribution plots, the reader always looks from the US side down to the DS side of the experiment.

4.2.2. Investigating the Number of Muons Stopped

The definition of the number of stopped muons, denoted as N_{μ}^{stopped} , is straightforward and aligns with its name: Amount of muons that get successfully stopped inside the muon-stopping target (disks). This parameter can be directly read out from simulation data by following all muon tracks and identifying their end positions. Alternatively, target-disks can be configured to directly record muons stopped inside.

Additional Information: Utilizing Monitor Planes

While obtaining the number of stopped muons from a completed ICEDUST run is straightforward, drawing meaningful conclusions often requires additional information. In particular, information regarding particles that could *not* be stopped, momentum distributions, or x - y -profiles of arriving particles are of significant interest. To access this supplementary information, so-called monitor planes are used. These thin volumes are placed at key

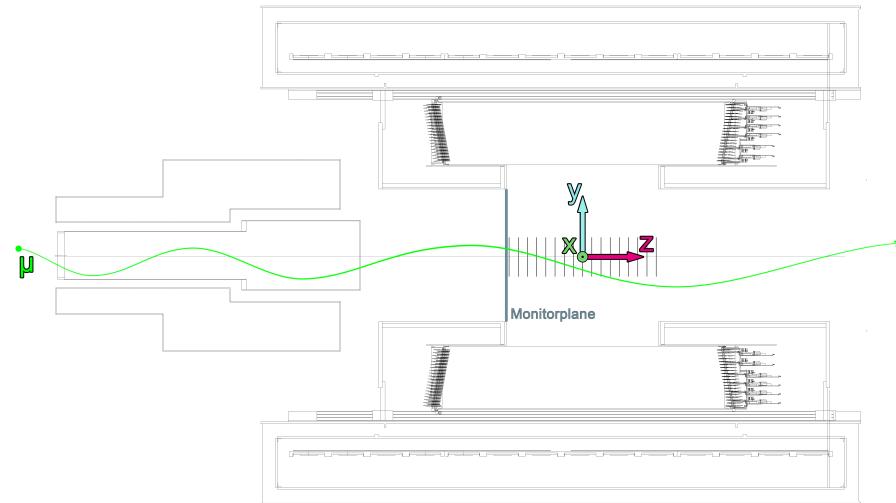


Figure 4.3.: Side view of the CyDet area. The technical drawing includes the defined centered coordinate system, the monitor plane utilized for particle x - y -distribution plots (grey blue), and an example muon trajectory (green line).

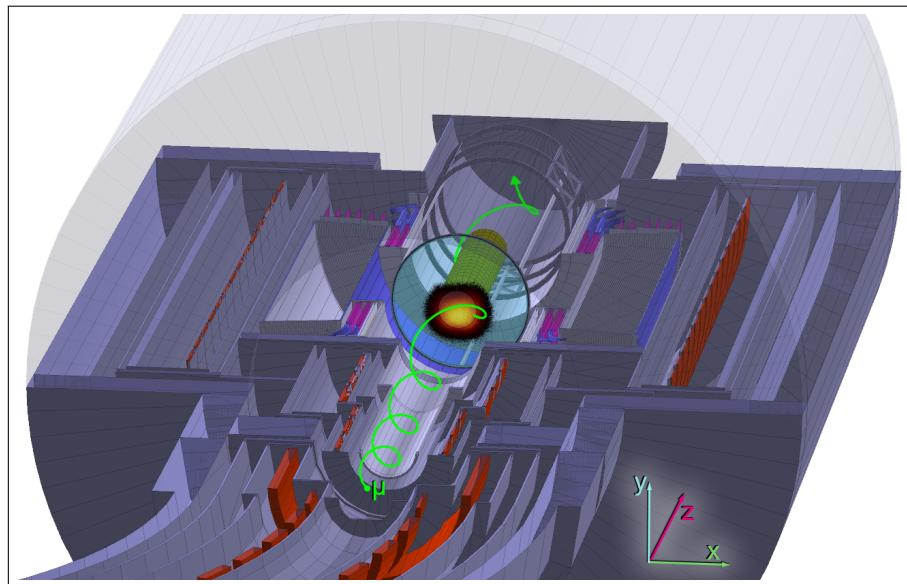


Figure 4.4.: Cut-open 3D view of CyDet area as implemented in ICEDUST. The defined centered coordinate system is shown on the bottom right. The monitor plane utilized for particle x - y -distribution plots (grey blue) and an example muon trajectory (green line) are included.

positions throughout the experiment and are configured to record particle information upon entry or exit. The material of monitor planes is identical to their surroundings, ensuring that they have no *physical* effect on particles passing through them.

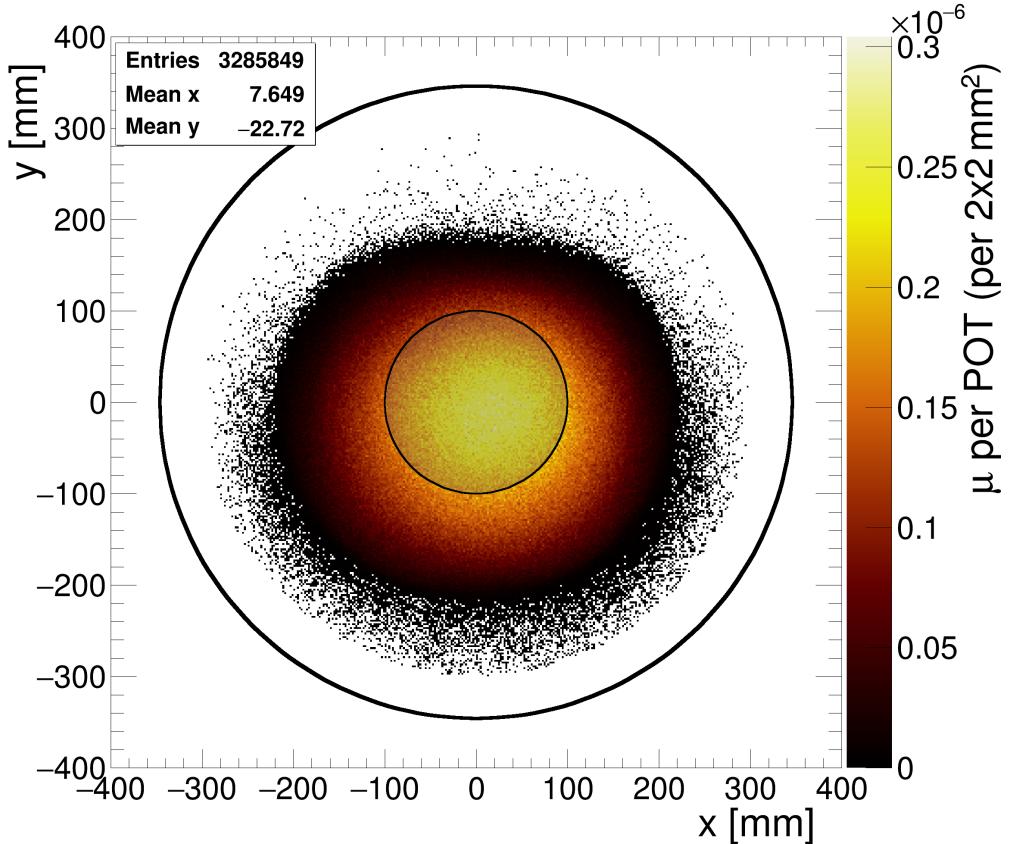


Figure 4.5.: Example of a muon x - y -profile recorded with the monitor plane in front of the muon-stopping target. The data corresponds to *all* muons arriving at the muon-stopping target for 60 bunches at 16×10^6 POT. Outer circle: CTH inner radius. Inner circle: Representation of the muon-stopping target (disks). A detailed discussion of the profile's features and their implications is given in Section 4.4.

An example of such a monitor plane positioned at $z = -405$ mm in front of the muon-stopping target was included in Figure 4.3 and Figure 4.4. A x - y -profile of all passing muons recorded with this monitor plane is shown in Figure 4.5. In contrast to directly using target-disks, the larger radius of the monitor plane and the removed limitation to only stopped muons enable the study of additional features. For example, the magnetic field configuration and the limited radius of the beam transport causing a nearly symmetrical but

slightly off-centered beam position. A more detailed discussion of these features and their implications is provided later in Section 4.4. For visual reference, such profile plots always feature a large circle at $r = 346$ mm, representing the inner radius of the CTH, and a circle with $r = 100$ mm, illustrating the muon-stopping target (disks). Furthermore, note that the shown profile has a sharp cutoff at 300 mm. This cutoff corresponds to the monitor plane's maximum radius, which was chosen to leave space for the muon-stopping target holding structure (see Section 5.1).

4.2.3. Investigating the Signal Acceptance

The signal acceptance in COMET, denoted as $A_{\mu \rightarrow e}$ and defined as the ratio of measured to occurred conversion events, is determined by the equation:

$$N_{\text{measured}} = N_{\mu \rightarrow e} \times A_{\mu \rightarrow e}. \quad (4.1)$$

Within ICEDUST, this acceptance is determined by isotropically generating conversion electrons from stopped muon positions and recording the fraction of electrons surviving a set of selection criteria (see Section 3.2).

Table 4.2.: Factors contributing to the $\mu - e$ conversion signal acceptance, as taken from the COMET Phase-I TDR [43].

Event selection	Value
Online event selection efficiency	0.9
DAQ efficiency	0.9
Track finding efficiency	0.99
Geometrical acceptance + Track quality cuts	0.18
Momentum window ($103.6 \text{ MeV}/c < P_e < 106.0 \text{ MeV}/c$)	0.93
Timing window ($700 \text{ ns} < t < 1170 \text{ ns}$)	0.3
Total	0.041

The signal acceptance comprises six efficiency factors, with previous values summarized in Table 4.2 [43]. However, not all of these efficiency factors are reevaluated in this thesis. For example, the identification of conversion electrons through *track finding and fitting* in the CDC detector is currently under active development by specialized COMET detector groups. To separate muon-stopping target optimizations from these ongoing developments, maintain consistency in results throughout this thesis, and limit the number of free parameters, the detector response is simplified. More precisely, no sophisticated detector response simulation is applied, and only MC truth values are taken into account. As a result, MC truth hit timings or conversion electron momenta are not reconstructed but recorded upon entry into the corresponding detector systems.

In summary, due to the constraint on MC truth values, *track finding*, and both hardware-related factors *Data AcQuisition (DAQ) efficiency* and *online event selection*, remain unchanged. For later calculations of the signal acceptance and the SES in Section 4.10, the values presented in Table 4.2 are used without modification. Definitions for the remaining efficiency factors, namely *geometric acceptance* (including a set of *track quality cuts*), *momentum window cut*, and *timing window cut* are discussed in the following.

Geometric acceptance To obtain a manageable data rate and suppress backgrounds, conversion electrons must first reach and travel through the CDC and subsequently hit the CTH. The geometric acceptance is defined as the ratio of conversion electrons fulfilling this condition to all conversion electrons emitted. Naturally, this geometric condition strongly depends on the conversion electron starting distributions and is significantly affected by changes of the muon-stopping target.

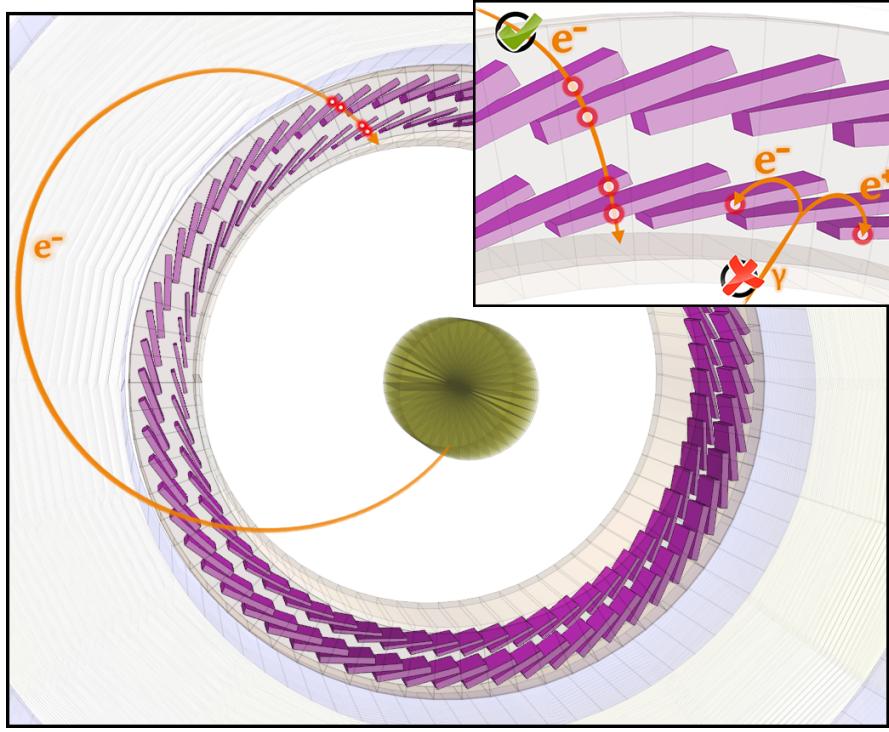


Figure 4.6.: Demonstration of 4-fold coincidence and its background suppression.

A conversion electron is emitted from the muon-stopping target (center yellow circles), traverses through the CDC layers (grey lines), and hits the CTH counters (purple squares). The inset demonstrates the required hit in four neighboring counters and the suppression of gamma-ray background.

Both CDC and CTH add further conditions. Regarding the CTH, a 4-fold coincidence is required. To fulfill 4-fold coincidence, four neighboring CTH counters (two in the inner and two in the outer layer) have to be hit in a time span of less than 10 ns. This procedure is illustrated in Figure 4.6 and is necessary to reduce fake triggers caused by gamma-rays, coming mostly from Bremsstrahlung of DIO electrons [61].

To further reduce low momentum background, an additional set of track quality cuts is applied for the CDC, as demonstrated in Figure 4.7. The five track quality cuts are defined as follows:

1. Tracks must reach the 5th sense layer (NL5)
2. At least one whole turn in CDC is required (NHIT)
3. Hits are required in more than three consecutive layers at both the entrance and exit points of the tracks (CL3)
4. *Track fitting*: The number of degrees of freedom must be greater than 30 (NDF30)
5. *Track fitting*: The normalized χ^2 must be less than two (χ^2)

Note that due to this thesis restricting CDC analysis to MC truth data, the track fitting procedure-dependent conditions 4 and 5 cannot be included in calculations of the geometric acceptance. However, they are accounted for separately when calculating the SES, as described in Section 4.10.

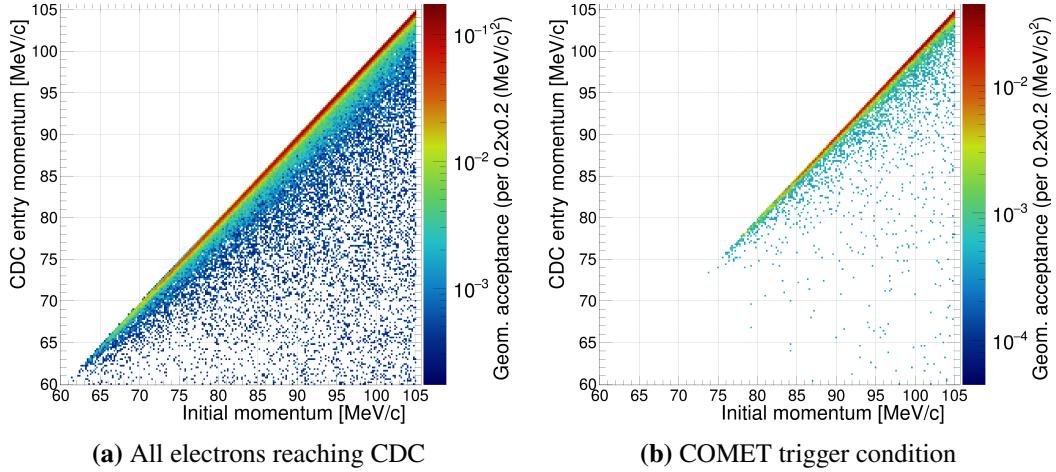


Figure 4.7.: Momentum transfer matrix for electrons originating from the muon-stopping target. Only electrons with an initial momentum larger than 60 MeV/c can reach the CDC. Requiring the COMET trigger condition (4-fold coincidence and CDC track quality cuts) removes most of the low-energy tail events and only electrons with momenta above 75 MeV/c are accepted, suppressing larger parts of the DIO background.

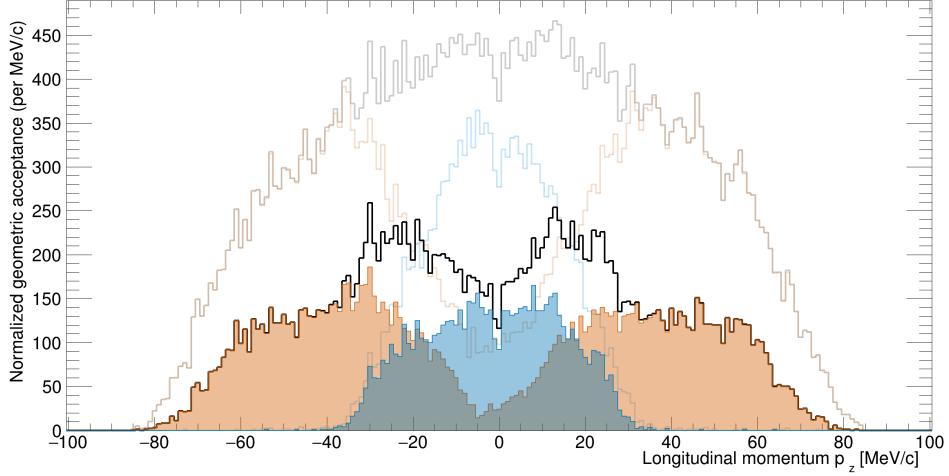


Figure 4.8.: Longitudinal momentum distribution of single- (orange) and multiple-turn (blue) tracks surviving the CDC track quality cuts. The shaded histograms in the foreground additionally require CTH 4-fold coincidence. Combined values are represented by the black line.

Table 4.3.: Geometric acceptance and breakdown of applied track quality cuts, as taken from the COMET Phase-I TDR [43]. Values are separated into single- and multiple-turn tracks and normalized to all signal tracks. Conversion electron signals were emitted isotropically from the muon-stopping target.

	single-turn tracks	multiple-turn tracks	combined
Geometric acceptance	0.21	0.13	0.34
incl. 4-fold coincidence	0.16	0.10	0.26
-----	-----	-----	-----
NL5	0.78	0.98	
NHIT + CL3 + NDF30 + χ^2	0.91	0.73	
-----	-----	-----	-----
Total	0.11	0.07	0.18

Previously estimated values for geometric acceptance and track quality cuts are presented in Table 4.3 [43]. Notably, they are separated into single- and multiple-turn tracks, referring to the conversion electrons' number of rotations through CDC before reaching the CTH. The number of rotations strongly depends on a conversion electron's longitudinal momentum, as shown in Figure 4.8. The ratio of single- to multiple-turn events is approximately 2/3 to 1/3. Due to a differing CDC response, this classification is particularly important for track fitting and finding procedures. However, it plays a minor role when assessing the performance of a muon-stopping target configuration, where only the combined total geometric acceptance is

required. Consequently, values for the geometric acceptance throughout this work are given as a combination of single- and multiple-turn tracks, which cause a 4-fold coincidence in CTH and satisfy track quality cut conditions 1 to 3 as listed. Electrons fulfilling these cut criteria are referred to as *fulfilling the COMET trigger condition*. The associated acceptance parameter is denoted as $A_{\mu \rightarrow e}^{\text{GM}^*}$, including a superscripted apostrophe highlighting the reduced set of track quality cuts.

Momentum Window Cut The momentum window cut rejects electrons outside a pre-defined momentum Region Of Interest (ROI). Energy losses due to Bremsstrahlung or ionization of material along the conversion electron's path will alter the intrinsically monoenergetic $E_{\mu e, \text{Al}} = 104.97 \text{ MeV}$ signal, leading to a low-momentum tail, as shown in Figure 4.9. To account for this tail, the momentum window in which electrons are accepted is chosen to be broad around their initial momentum, accepting all electrons within $P_e \in [103.6, 106.0] \text{ MeV}/c$, and resulting in a cut efficiency of $A_{\mu \rightarrow e}^{\text{ROI}} = 0.93$ [43].

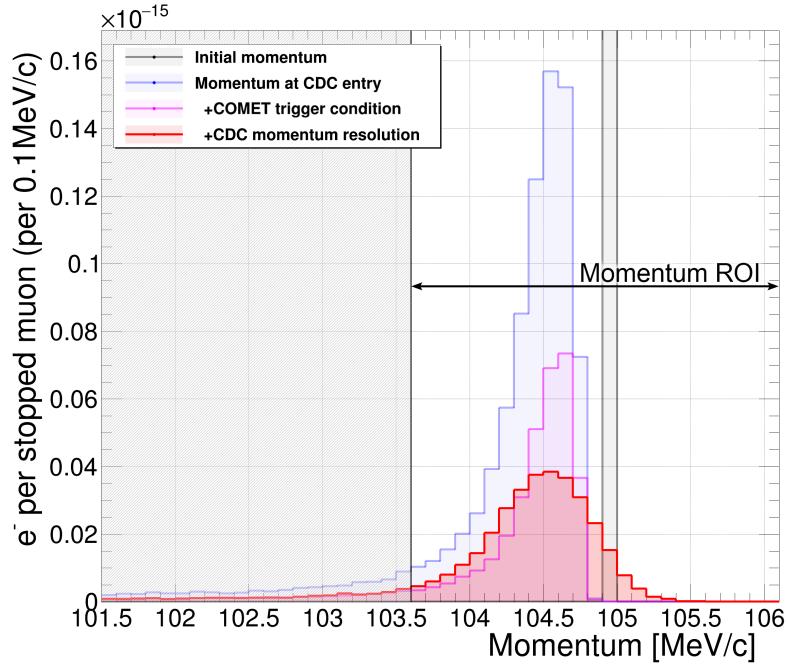


Figure 4.9.: Momentum distribution of conversion electrons when created (black) and after reaching the CDC (blue). Requiring the COMET trigger condition (purple) removes much of the low-energy tail but also reduces the overall signal by approximately 64 %. MC truth values are convolved with a Gaussian distribution with a standard deviation of 200 keV/c (red) to approximate the CDC momentum resolution.

Figure 4.9 includes resolution effects from the CDC. While the simulation is limited to MC truth data, investigations that involve the momentum window cut approximate the resolution by convolving MC truth values with a Gaussian distribution with a standard deviation of 200 keV/c, corresponding to COMET’s design value [43]. The combined acceptance parameter, fulfilling the COMET trigger and the ROI condition, will be cited as $A_{\mu \rightarrow e}^{\text{GM,ROI}}$.

The upper boundary of the ROI is selected to account for electrons with momenta higher than $E_{\mu e, \text{Al}}$, due to the CDC resolution effects. The lower boundary is determined by the acceptable level of background contamination. The cut is particularly effective in suppressing backgrounds from processes like DIO and Radiative Muon Capture (RMC) because their energy spectra drop off sharply toward the conversion energy. For the default muon-stopping target and $P_e \geq 103.6 \text{ MeV/c}$, contaminations were determined to be 1 % DIO and 0.19 % RMC [43].

Throughout most of this thesis, the ROI is held constant. A dedicated investigation of $A_{\mu \rightarrow e}^{\text{ROI}}$ and DIO background is contained in Section 4.9. Modifications of the muon-stopping target, particularly of the total amount of material, can amplify the low-momentum tail of the conversion electron’s distribution. However, the impact on the momentum window cut efficiency consistently stayed below a few percent and was therefore disregarded. Configurations involving a substantial increase in material, which might cause more significant changes in $A_{\mu \rightarrow e}^{\text{ROI}}$, are not possible due to excessively high detector backgrounds (see Figure 4.30 of Section 4.7).

Timing Window Cut The 864 ns lifetime of muonic aluminium is used to distinguish the conversion electron signal from backgrounds caused by the prompt beam flash by only accepting events occurring in a $700 \text{ ns} < t < 1170 \text{ ns}$ time window (compare Figure 2.4). The fraction of surviving conversion electron events defines the timing window cut efficiency. However, this thesis will not reassess this efficiency factor and instead uses the tabulated value in Table 4.2 for signal acceptance calculations. The lifetime of the muonic atom and, consequently, the efficiency of this cut depends solely on the target-disks’ unchanged material rather than their exact geometric configuration. Therefore, no significant changes in the timing window cut efficiency are expected in either case. The effects of this cut on background suppression are discussed in the next section.

4.2.4. Investigating Detector Background Rates

The background in COMET’s Phase-I detectors, namely CDC and CTH, will emerge as the major limiting factor for the design of the muon-stopping target. Positioned at the center of CyDet, the target-disks are hit by the majority of beam particles. Increasing the amount of muon-stopping material, for example by increasing the number of target-disks, inevitably raises detector background caused by Bremsstrahlung, scattered particles, or decay of excited nuclei. Background in the CDC and CTH must be treated quite differently. Hence, the

following two paragraphs explore both detector systems individually. To allow for the correct interpretation of later results, particular focus is placed on this study's limitations due to the reliance on MC truth data and the limited amount of MC5 statistics (see Section 3.1).

CDC Instead of looking at a total background rate, the severity of CDC background is estimated by inspecting the *occupancy* of its layers on a bunch-to-bunch basis. The layer occupancy is defined as the fraction of wires (or cells) in layer i where a hit occurred. It is calculated with:

$$O_i = \frac{N_{\text{hit}, i}}{N_{\text{wires}, i}}; \quad i \in [1, 18] \quad (4.2)$$

$$N_{\text{wires}, i} = (384 + 12 \cdot i) / 2.$$

Cell information is binary, and over the duration of one bunch each cell has or has not recorded a hit. The innermost ($i = 0$) and outmost ($i = 19$) layers are so-called guard layers (see Section 2.2.1) and are not considered in background estimates. The total CDC occupancy is defined as the fraction of all 4482 wires in layers 1 to 18, which recorded a hit.

Using ICEDUST and relying on MC truth data imposes three major limitations:

1. *Background suppression based on total energy deposition is not considered.*

Particles deposit energy in the CDC by ionizing the gas along their flight path. The ionization electrons drift toward the wire, are amplified, and recorded as an electrical signal. The total strength of this electronic signal differs between (conversion) electrons and backgrounds like protons, as shown in Figure 4.10. In short, approximately 68 % of background can be suppressed by rejecting cells with energy depositions larger than 5 keV [43].

However, energy depositions in ICEDUST are recorded as hits, separate for each particle (track). The formation of ionization electron avalanches and accumulation of same-cell energy depositions strongly depend on hardware parameters, particularly the timing response, and thus require a dedicated detector response simulation. Due to the restriction to MC truth data, background suppression based on a total energy deposition threshold is not included in this thesis.

2. *Background suppression based on hit timing is not considered.*

COMET utilizes a $700 \text{ ns} < t < 1170 \text{ ns}$ timing window to mitigate the high background rate from the beam flash. The potential strength of such a cut is shown in Figure 4.11, where the total CDC occupancy is reduced from 27.8 % down to 5.4 %. While hit times in ICEDUST are recorded instantaneously, in reality, the signal arrives over an extended period due to ionization electrons' drift times. Using isobutane ($\text{He:C}_4\text{H}_{10}$ (90:10)) as counting gas, the maximum drift time corresponding to a hit around the cell corners is 700 ns [43]. Hence, energy depositions from the beam flash could be smeared inside the delayed measurement timing window, or signal events could be

delayed enough to be superimposed by the beam flash of the subsequent bunch. In consequence, the CDC must be able to handle hits from an entire beam flash, and the simplified timing window shown in Figure 4.11 cannot be applied.

Naturally, not all hits occur simultaneously, and some background reduction based on hit times in the CDC will be utilized, especially in combination with information from the CTH. However, such timing cuts strongly depend on hardware parameters and cannot be applied to MC truth data without dedicated detector response studies. Thus, estimates of CDC occupancy in this thesis do not include timing cuts and instead consider *all* hits accumulated over an entire bunch of 16×10^6 POT.

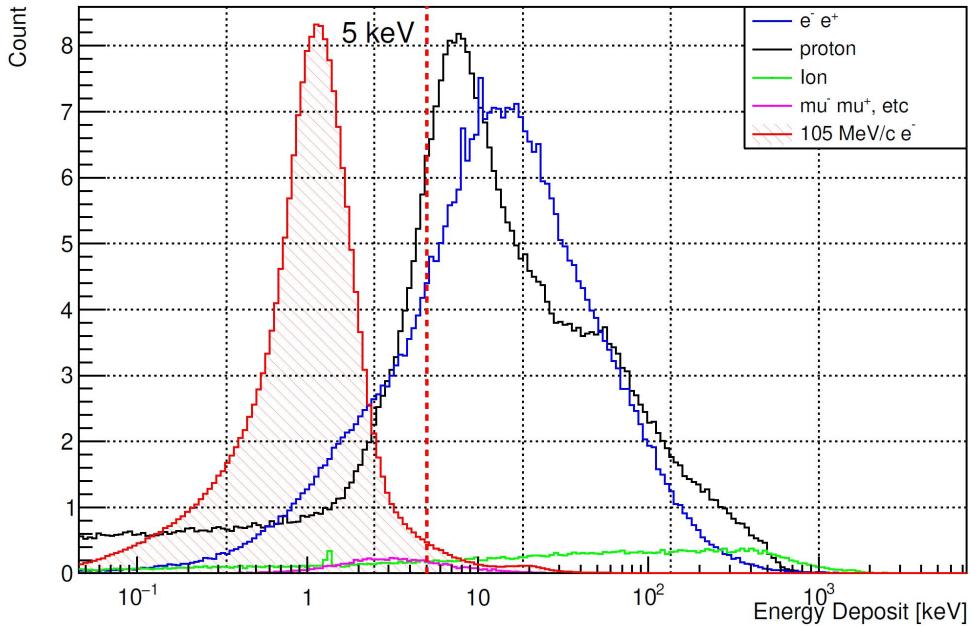


Figure 4.10.: Total energy deposits per cell for signal electrons and noise hits. The signal electrons are shown as red shaded histogram; noise hits from various sources are shown in other histograms. The 5 keV threshold line is also shown. Taken from the COMET Phase-I TDR [43].

3. Due to the limited statistics, no bunch pile-up can be considered.

Section 3.1 introduced the use of the MC5 dataset to ease the challenging demands on computing time. As a drawback, this also limits statistics to the size of the MC5 dataset, namely its 62.5 bunches at 16×10^6 POT. Given the experiment's small bunch separation of 1.17 μ s and the beforementioned CDC drift times, pile-up from delayed events superimposed on subsequent bunches is expected, leading to the build-up of a quasi-permanent background within CDC [61]. Due to the limited number of available bunches and the exclusion of CDC timings, such pile-up effects could not be

considered, and each bunch was treated independently. However, it should be noted that this effect is already partially accounted for because strongly delayed hits, usually contributing to pile-ups in follow-up bunches, are accounted for in their bunch of creation. The missing pile-up effects are estimated to amount to at most 2 to 3 % total CDC occupancy.

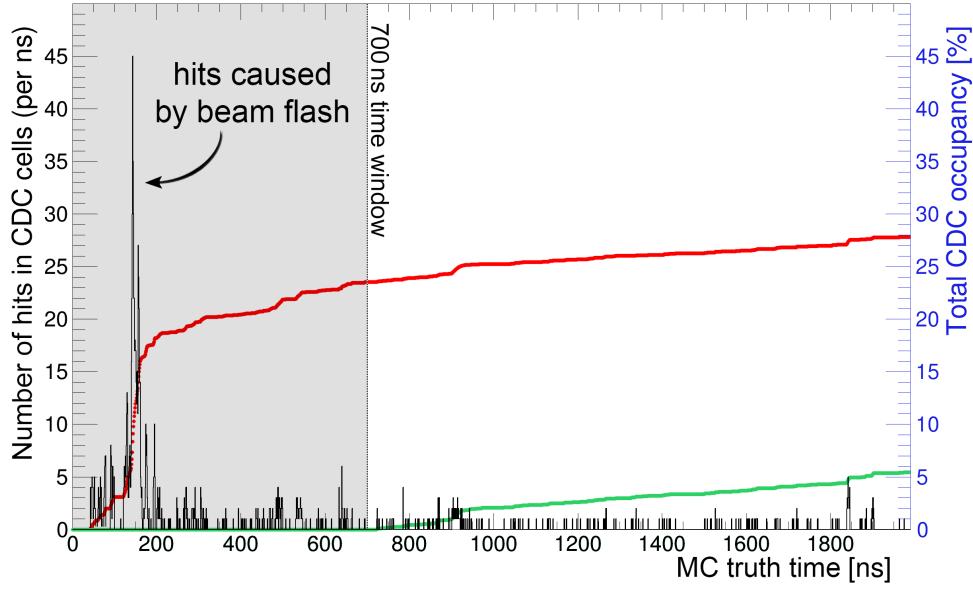


Figure 4.11.: MC truth timing versus number of hits in CDC cells. Overlayed in red (green) is the accumulated total CDC occupancy without (with) applied timing window cut (grey box). Data shown corresponds to a single isolated proton bunch at 16×10^6 POT.

In summary, quoted CDC occupancy levels within this thesis exclude any form of cuts on timing or energy of hits. Although bunch pile-up is not considered, presented values should therefore be understood as conservative *worst-case* limits. Comparisons to internal studies, which applied under-development online and offline hit classification using neural networks or gradient-boosted decision trees, suggest for MC truth-based analysis a maximum permissible CDC occupancy of 50 % before track fitting and finding gets impaired [61]. The performed optimizations will be guided by this value.

CTH The intensity of background in CTH is judged by the total amount of hits in the counters and the corresponding number of 4-fold coincidences. Unlike drift chambers, scintillators exhibit a timing response on the order of a few nanoseconds. Consequently, the differences between MC truth and actual measurement timings are sufficiently small to allow for timing considerations in the analysis [54].

The CTH background analysis performed on MC truth data is separated into three steps. First, the delayed measurement time window is applied, and only hits between 700 to 1170 ns are accepted. Second, energy depositions of all hits within the same counter and with a time difference smaller than 20 ns are grouped to prevent an artificial increase in the total number of CTH counter hits caused by the numerous low-energy secondaries inside the scintillators. Finally, the 4-fold coincidence condition between neighboring counters is checked using a 10 ns coincidence time window.

A more realistic detector response considering Birks' law, dead times, or a sophisticated waveform analysis for PID is still under active development by the COMET collaboration and has not been included yet. Therefore, the predicted background hit rate quoted in this work is many times higher than expected for actual counters. Due to such a high rate, each bunch contained an average of 25 4-fold coincidences. Consequently, the number of 4-fold coincidences, which must stay below one to work properly and remove background-only bunches, does not carry much meaning. To still assess the increase in CTH background, the total count rate will be provided as an arbitrary measure.

4.2.5. Treatment of Uncertainties

Uncertainties are essential for evaluating results in physics. However, this thesis will present most values obtained from MC simulations without associated uncertainties. The deliberate decision to omit uncertainties, particularly for performance overviews or a configuration's SES estimates, is not an oversight but is motivated by the desire to increase clarity. All investigations of the muon-stopping target parameters rely on ICEDUST, meaning the main source of uncertainty in these results is always this common source. Consequently, this section will summarize the uncertainties once to avoid repetitive and quasi-identical redundant information throughout the rest of this thesis.

Statistical uncertainties This discussion is divided into three parts, starting with contributions from statistical uncertainties. Results from the MC simulation are primarily based on counting. This counting of particles (or successes), such as the number of muons stopped or surviving signal electrons, follows Poisson statistics. In a single simulation run, the number of successes N is assumed to equal the mean, and a reasonable uncertainty estimate is given by the standard deviation defined as $\Delta N = \sigma_N = \sqrt{N}$. The relative (statistical) uncertainty is then calculated with $\Delta N/N = 1/\sqrt{N}$. As detailed in Section 3.1, the primary beam statistics correspond to *at least* 14 bunches at 16×10^6 POT, or 60 bunches for key results. This translates to at least $\mathcal{O}(10^5)$ stopped muons and a relative uncertainty of less than 0.3 %. For results based on counting secondaries, such as conversion electrons or X-rays (compare Section 3.2), a minimum of 100 000 started events was chosen. This number ensures that relative uncertainties always lie below 1 %.

Systematic uncertainties of MC components Accurately assessing systematic uncertainties of MC components is much more challenging. These uncertainties primarily include numerical inaccuracies from the implementation of geometries (size, positioning, material) or the digitized treatment of the magnetic field (compare Section 3.3). Smaller tests of individual subcomponents, such as Cosmic ray tests from the CDC [78, 79], agree well with results from ICEDUST and provide a general level of confidence. However, reliable constraints on uncertainties related to the interplay of multiple components or the muon beam (and deliberate fine-tuning of the simulation) will only be possible once COMET completes its construction and can conduct test measurements. In the absence of such data, a rough estimate can be made by varying key MC parameters. For example, adjusting the dipole field strength during beam transport by around 10 mT changes the rate of muon stops by approximately 1 %. Overall, these systematic uncertainties are conservatively estimated to lie around the 10 % level [49].

Systematic uncertainties of the MC physics model Systematic uncertainties also arise from inaccuracies in the MC physics model. As described in Section 3.4, ICEDUST relies on Geant4 PhysicsLists, which are extensively tested and validated against particle physics measurements. However, there is currently no comparison data available for a core aspect of COMET: the backward production of pions from 8 GeV proton collisions in the graphite pion production target. Different hadron models produce pion yields that vary by up to a factor of three [43]. The QGSP_BERT_HP model used in ICEDUST showed the second lowest yield, thus results based on the number of pions or muons can be understood as conservative estimates. Nevertheless, the systematic uncertainty of this chosen model is at least in the order of 10 %[†]. The staged approach adopted by COMET (see Section 2.1) will significantly help to limit (the impact of) these uncertainties. Additionally, results from a recently performed pre-measurement, called Phase- α [81], will be used to implement constraints for the Phase-I measurement.

Summary In summary, the relative statistical uncertainties for the number of stopped muons within target-disks are consistently below 0.3 %. Statistics from secondary simulations used to calculate acceptance factors were chosen high enough to ensure relative uncertainties below 1 %. However, as COMET is still in its design phase, these statistical uncertainties are dominated by significant hard-to-estimate systematic uncertainties of at least 10 %. A precise assessment of the contribution and correlation of these systematic uncertainties to individual results is challenging. To prevent cluttering later results with an always included greater than 10 % systematic uncertainty or misleadingly small statistical uncertainties, both were omitted when providing ICEDUST-based results.

[†]This uncertainty estimate aligns with findings from the Mu2e collaboration [44], which normalized backward pion production yields from various Geant4 models to experimental data (at higher proton energies) from the HARP experiment [80].

4.3. Performance Overview: Default Configuration

Since the publication of the Phase-I TDR in 2020, many updates and adjustments to the ICEDUST simulation have been performed to reflect the developments in COMET's design. To allow for comparison between investigated improvements and the default muon-stopping target, this section provides updated values for this default configuration obtained with an identical ICEDUST version and set of analysis tools as used throughout the later sections. The underlying geometrical details of ICEDUST, PhysicsLists, etc., were presented in Section 3. The definitions and especially the limitations of investigated main parameters were discussed in great detail throughout Section 4.2.

Table 4.4.: Updated performance overview of the 17×0.2 mm disk configuration.

The data is averaged over 60 MC5 bunches at 16×10^6 POT per bunch. Values are obtained following the procedure described in Section 4.2.

$N_{\mu/\text{bunch}}^{\text{stopped}}$ [per bunch]	$A_{\mu \rightarrow e}^{\text{GM,ROI}}$ [%]	normalized signal, K_{NS}	total CDC occupancy [%]	CTH rate [per bunch]
7833.94	17.79	1.00	32.07	1766.0

These so-called performance overviews will serve as a basis for comparisons between investigated designs. The data for the default muon-stopping target is presented in Table 4.4. Values are given on a *per-bunch* basis to allow comparisons between simulations with different underlying statistics (i.e. a different number of simulated bunches). In practice, this means parameters are first calculated individually for each bunch and then averaged over the total number of simulated bunches. To further facilitate comparisons between investigated configurations, the number of stopped muons per bunch $N_{\mu/\text{bunch}}^{\text{stopped}}$, and the signal acceptance factor $A_{\mu \rightarrow e}^{\text{GM,ROI}}$ are combined into a *normalized signal*, defined as:

$$K_{NS} = \frac{N_{\mu/\text{bunch}}^{\text{stopped}} \cdot A_{\mu \rightarrow e}^{\text{GM,ROI}}}{K_{NS}^{\text{default}}}. \quad (4.3)$$

The value is divided by the normalized signal of the default 17×0.2 mm disk configuration, denoted as K_{NS}^{default} . Thus, values of K_{NS} larger than one correspond to improvement.

Over 146 days of Phase-I runtime, COMET is expected to receive a total of 3.2×10^{19} protons, translating to 2×10^{12} bunches with 16×10^6 POT per bunch. Hence, the total number of muons stopped over this period is calculated by multiplying the per-bunch value with 2×10^{12} . In comparison to values from the COMET Phase-I TDR presented in the status-quo Section 4.1, this thesis obtains an increased value of 1.57×10^{16} total number of stopped muons, or 4.89×10^{-4} muons stopped per POT.

4.3.1. Additional Information: Single Layer Occupancy

The total CDC occupancy imposes a severe limitation for the muon-stopping target design. To further develop intuition, the separation of backgrounds into individual CDC layers and their origin will be utilized. Figure 4.12 visualizes this for the default muon-stopping target configuration. Muons interacting (stopped or passing) with the target-disks amount to approximately 35 % of all CDC background. An additional 25 % stems from muon interactions with parts of the beamline. The figure categorizes muon background into *stopped* inside the target-disks or *beam-related*. Unfortunately, due to the muons' core role in the experiment, this background is inevitable, and more optimized designs with higher muon stops per POT and stronger muon slowdown are expected to increase this contribution further.

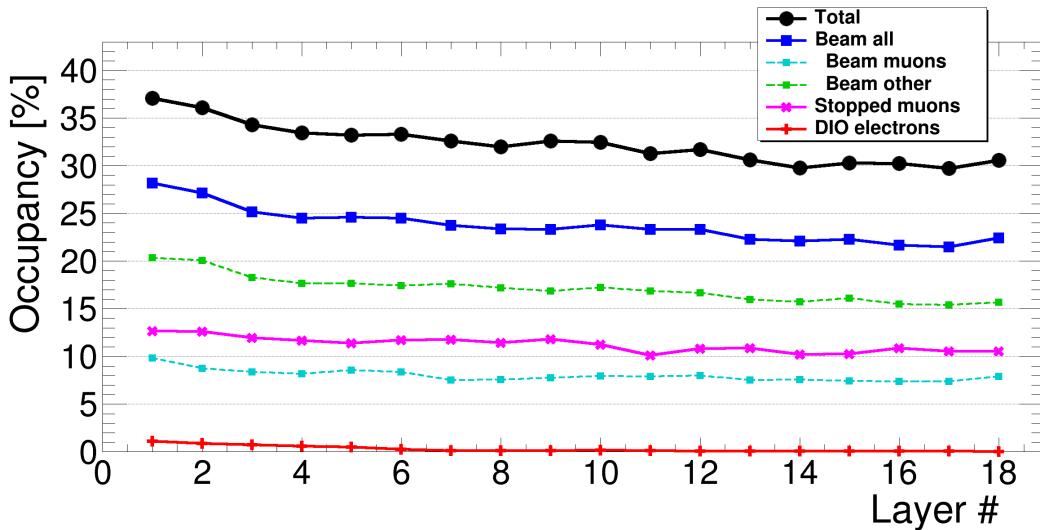


Figure 4.12.: CDC layer occupancy for the default muon-stopping target design. The data is averaged over 60 MC5 bunches. Statistical error bars are included but are too small to see. Background contributions are analyzed independently to avoid interference. Due to the binary treatment of cell hits (see Section 4.2.4), this procedure causes some cells to be double-counted for individual contributions. Consequently, beam all (sum of beam muons and beam other) and total (sum of beam all, stopped muons, and DIO) are 2 to 3 % smaller than the sum of their components. Smaller layer numbers correspond to smaller radii in the CDC. This not only translates to a smaller number of wires (see Formula 4.2), but also to an overall higher occupancy due to their closer position to the muon-stopping target. Further details of the *other beam* component are given in Table 4.5.

Table 4.5.: Particle separated contribution to total CDC occupancy for the default muon-stopping target design. Only primaries whose secondaries cause CDC hits are accounted for. The data is averaged over 60 MC5 bunches. See text for discussion.

Particle type	Number of primary traj. [# per bunch]	Total CDC occupancy [%]
Pion	39.12 ± 0.81	10.79 ± 0.06
Neutron	124.27 ± 1.44	4.82 ± 0.04
Electron	18.95 ± 0.56	1.49 ± 0.02
Positron	2.55 ± 0.21	0.15 ± 0.01

The remaining 40 % of CDC occupancy can be attributed to *other* particles, such as pions, neutrons, or electrons. The particle yield and momentum distributions after the transport solenoid were introduced in Figure 2.7 of Section 2.3.2. Further details of this component are summarized in Table 4.5. The term *primary trajectory* refers to the most important particle responsible for a background hit. Often, this corresponds to the direct parent of an ionizing electron inside the CDC. However, for more complex situations like a full-fledged electromagnetic shower, the full chain of reactions is analyzed, and the primary trajectory is selected based on the process or location of the interactions.

The largest part of the beam-related *other particle* background can be attributed to pions and neutrons. Neutrons arrive at the CyDet either by deflecting and scattering through the beamline or, due to their low interaction probability, pass through the walls and shielding and originate directly from the pion production target. They span a large momentum range, with most having around a few MeV/c and a long high-momentum tail up to 100 MeV/c. Because neutrons are electrically neutral particles they do not directly produce CDC signals. Their most intense background stems from protons emitted after neutron-proton scattering. Processes like slow neutron capture or neutron inelastic scattering are also possible but mainly produce gammas, which must undergo further interactions before causing CDC hits.

Negative pions primarily cause CDC background through hadronic capture at rest. Due to the pion's large total energy, the nucleus becomes highly excited after capture, and de-excitation occurs via emitting a multitude of neutrons, gammas, electrons, and protons. Suppose nuclear pion capture happens inside the muon-stopping target, outward-directed high energetic protons then travel through the whole radial length of CDC. Additionally, the reaction often produces up to four neutrons, accompanied by multiple gammas and electrons. Consequently, a single nuclear pion capture amounts to a significant number of hits in the counting gas. This translates to almost twice the total CDC occupancy of neutrons, even though only a third of primary pion trajectories reach the muon-stopping target region.

As Table 4.5 reveals, primary electrons play only a minor role for CDC background. As further discussed later (see Figure 4.24), especially low-momentum electrons ($< 20 \text{ MeV}/c$) have the correct energy-mass combination to survive through the beam transport. However, the momentum of such beam electrons is not large enough to directly reach and hit the CDC. Hence, beam electrons cause background by producing gammas through Bremsstrahlung inside the target-disks, which then require further reactions like pair production, photoeffect, or Compton scattering to cause hits in the CDC. Nevertheless, this background must not be disregarded, as it poses a major limitation for the muon-stopping monitor, as further described in Section 6.

4.4. Starting Point: The Tube-like Disk Structure

The starting point for the muon-stopping target investigations performed in this thesis is the design specified in the COMET 2014 TDR, as introduced in Section 4.1. The design consists of multiple circular disks positioned in a straight line after each other to form a tube-like structure. However, since its proposal, a lot of updates on the experiment were performed. Thus, before diving into details, this section will first verify once again the muon-stopping target's fundamental aspects to justify the chosen starting point, namely: the use of multiple, separate thin disks, their circular shape, and their aligned positioning to form the tube-shape.

4.4.1. Matching the Target Shape to the Beam Profile

Before reaching the muon-stopping target, muons must pass through the 90° circular transport solenoid. The transport shapes the arriving muon profile to resemble a round cylindrically symmetric Gaussian, as shown in Figure 4.13. The profiles possess two distinct features: First, a small dent in the upper part ($y \approx 180 \text{ mm}$) caused by the beam blocker inside the beam transport system. This cutoff is located far outside the intended target-disks, thus it does not affect the overall design. Second, a slightly off-center mean, which strongly depends on the muon momentum and is caused by dispersion effects introduced in Section 2.3.2. Because the consequences of this offset are more related to the overall muon-target's positioning than its shape they will be discussed separately in Section 4.4.2.

It is important to note that the shown profiles only represent a snapshot of the muon distribution recorded in front of the muon-stopping target (see Section 4.2.2). However, due to the approximately constant 1 T magnetic field immersing the CyDet, the Lorentz force forces muon trajectories to follow helical paths instead of straight lines. Details are discussed in Section 4.5.1. Consequently, the shown profiles should be understood as a probability distribution to find muons at a particular position. While one muon may rotate outside of a specific radius, another muon will take its place by rotating inside.

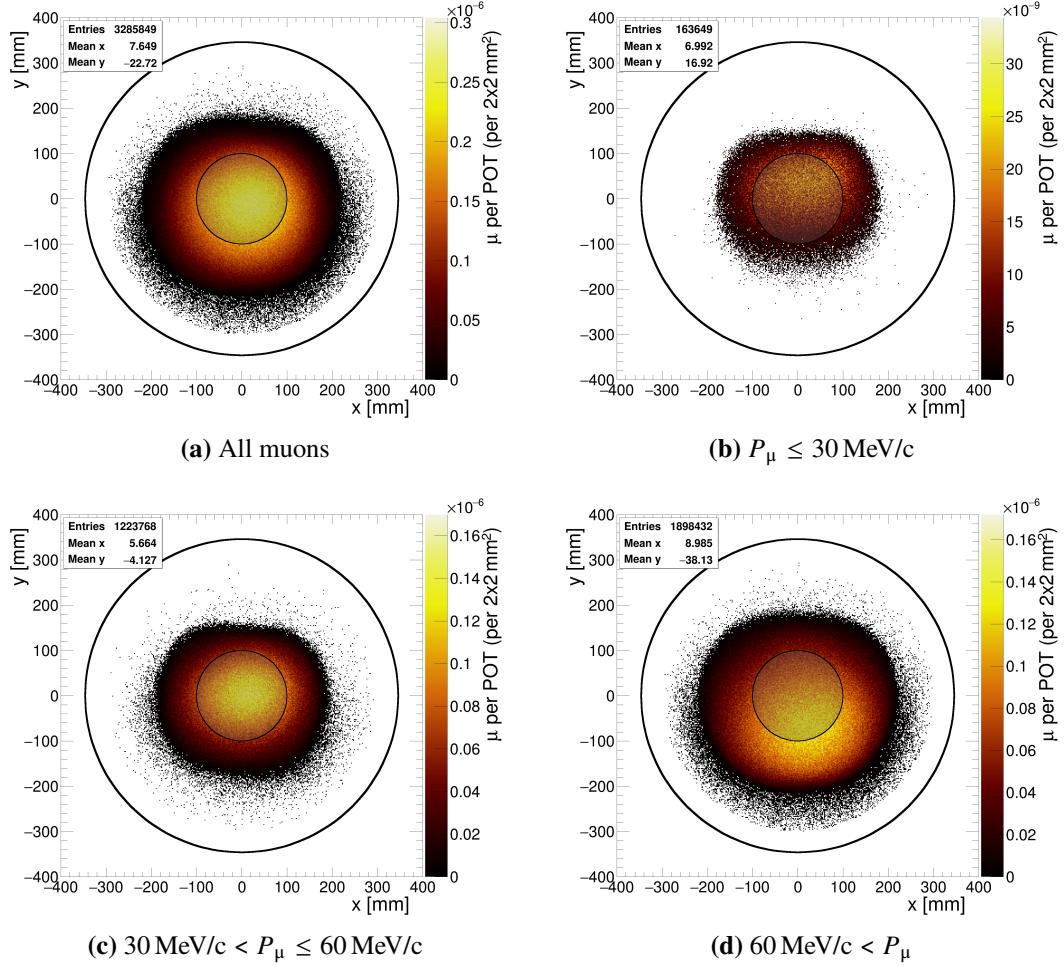


Figure 4.13.: Muon x - y -distribution in front of the muon-stopping target for different momentum ranges. Note that histogram z -axes are adjusted to account for different absolute numbers of muons (compare total muon flux shown in Figure 4.2). The data corresponds to 60 bunches at 16×10^6 POT. Overlaid circles at $r = 346$ mm and $r = 100$ mm illustrate the inner radius of CTH and the target-disks, respectively. The dip at the upper part of the profiles is caused by the beam blocker during beam transport but is located at a large enough radius not to impact the target design. For increasing momentum, a muon's most probable position in the beam drops lower due to dispersion effects discussed in Section 2.3.2. The magnetic field during beam transport is fine-tuned to center muons with momenta around 45 MeV/c.

To guarantee maximum muon-slowdown and number of stops, the x - y -geometry of the muon-stopping target should closely match the arriving muon beam. As shown, the expected muon profile resembles a cylindrically symmetric Gaussian, and consequently, the use of round target-disks is still well justified. It is also worth noting that the radially symmetric detector system for COMET Phase-I was designed around this circular muon profile, and the overall symmetry of the CyDet system will be actively used for signal tracking and background suppression.

4.4.2. Using Multiple Separate Target-Disks Forming a Tube

The decision to separate the circular muon-stopping target into multiple spaced-out disks is justified by revisiting its two competing constraints: maximizing the number of stopped muons and minimizing the energy loss of conversion electrons. Employing thin, multiple disks instead of, for example, a solid tube of material fulfills both requirements. On the one hand, using thin target-disks and a reasonably large spacing between them minimizes the conversion electrons traveled distance through stopping material, thus reduces the conversion electrons energy loss. On the contrary, using multiple target-disks placed in a row increases the total amount of material to guarantee sufficient slowdown of muons and increases the number of muons stopped.

Despite having multiple separate target-disks, changes in radius, thickness, spacing, or positioning were consistently applied to the whole muon-stopping target. Initially, this decision was necessary to reduce the number of free optimization parameters, which otherwise increase multiplicatively with the number of disks. However, basic tests conducted at later stages, implementing single target-disks with an enlarged radius; thickness; or adjusted positioning, also showed worse performance than comparable standard configurations.

Intuition for this result can be gained by examining COMET's key parameters. For the number of muons stopped, the shape of the muon beam profile is most important. The profile over the length of the muon-stopping target is shown in Figure 4.14. Variations of the profile center and width remain below a few millimeters – negligible compared to the target-disks' radius. Consequently, each target-disk plays an equal role in slowing down and stopping muons. For signal acceptance, the situation is more complicated but similar. Details are discussed later in Section 4.8, but small offsets at either end of the muon-stopping target lead to approximately constant signal acceptance for each target-disks. In summary, there is no apparent motivation to single out and change target-disks individually.

Target-disks are placed at identical x - and y -positions behind each other, forming a tube-like structure. The middle point of this tube is chosen to align with the center of CyDet to maintain its radial symmetry and facilitate CDC track-finding and -fitting. However, from Figure 4.13 and Figure 4.14, it becomes apparent that the muon beam is not yet perfectly centered on this tube.

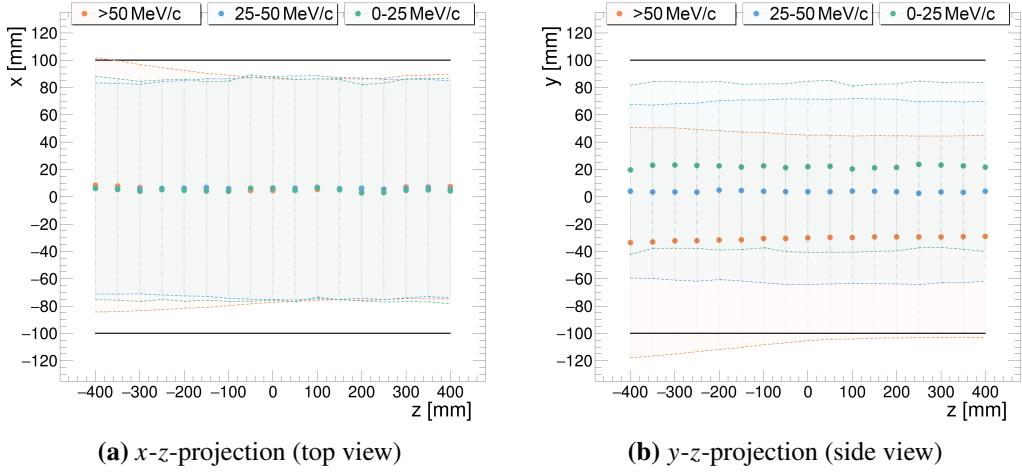


Figure 4.14.: Projected muon beam x - and y -distribution along the muon-stopping target. The muon profile was measured at each target-disk position, projected onto the x - (left) and y -axis (right), and then approximated with a Gaussian distribution. Larger dots represent the middle points, dashed lines the 68 % intervals. The bigger black lines correspond to the 100 mm default muon-stopping target radius. Muons are categorized into low-, medium- and high-momentum to separate momentum dispersion effects caused by the transport solenoid. For stoppable muons below 50 MeV/c, the profile remains approximately constant over the entire length and well confined within the muon-stopping target. The simulation was performed without target-disks to avoid bias from stopped muons, and the data corresponds to 60 bunches at 16×10^6 POT.

In x -direction, the most probable muon position is offset by approximately 6 mm. It can be adjusted by varying the position and tilt of the pion production target, and if corrected, the number of stopped muons could be increased by 0.5 % (compare Figure 5.7 in Section 5.4). However, no such corrections were performed to avoid an additional set of free parameters. Since the offset is small relative to the radius of the target-disks and remains identical across all configurations, the decision not to correct it does not have any immediate consequences for this study. However, after details of the pion production are finalized, proper fine-tuning and measurements should be conducted to guarantee a correct alignment.

The muon beam's y -position is adjusted by changing the compensating dipole field during beam transport (see Section 2.3.2). The discussion of its optimal configuration has to account for the strong momentum dispersion. Therefore, muons are divided into three categories following Figure 4.14, namely:

- *Low-momentum muons: $P_\mu \leq 25 \text{ MeV}/c$*

Muons of this category have a large y -offset of approximately 20 mm. Despite their off-center position, their 68 % interval is well confined within a 100 mm radius. Low-momentum muons do not require much slowdown, with most being stopped within the first or second traversed target-disks (see Figure 4.29 of Section 4.7). Consequently, no adjustments are necessary to stop muons of this category.

- *Medium-momentum muons: $25 \text{ MeV}/c < P_\mu \leq 50 \text{ MeV}/c$*

Muons in this momentum range will benefit the most from optimizing the number of passed through target-disks. Theoretically, $17 \times 0.2 \text{ mm}$ aluminium target-disks are sufficient to stop muons below $50 \text{ MeV}/c$. However, as highlighted throughout the motivation in Section 4.1, the fraction of stopped to available muons in this region drastically falls off. The reason lies with the muon helix trajectory, causing them to miss some target-disks, as further examined in the next section. Their most probable beam position should be focused on the middle to guarantee a maximum number of hits. As can be seen, appropriate adjustments to the magnetic field during beam transport have already been made. Consequently, no further changes are necessary.

- *High-momentum muons: $50 \text{ MeV}/c \leq P_\mu$*

Muons with momenta larger than $50 \text{ MeV}/c$ experience the strongest downward drift during beam transport. This results in a high y -offset of approximately -32 mm . Additionally, these muons have a very broad 68 % interval, which exceeds the muon-stopping target's default radius. However, the required material to sufficiently slow down such high-momentum muons far exceeds limits imposed by background concerns (see Section 4.7). Consequently, they are of no interest in COMET.

4.5. Changing the Target-Disk Radius

After establishing the general shape of the muon-stopping target as multiple thin disks, the significance of the target-disk's radius becomes apparent. Naturally, larger radii will lead to more muons stopped but a reduced signal acceptance and increased detector background. For smaller radii, the inverse is to be expected. To help develop a physicist's intuition, this section initially investigates the relationship between the target-disk radius and COMET's main parameters separately, disregarding other correlations. Thereafter, the acquired knowledge is combined to form a conclusion, and exact values are presented as a performance overview at the end.

4.5.1. The Effect on Total Number of Muons Stopped

While the target-disks' general shape can be justified by comparison to the arriving muon beam, maximizing the profile and target-disk overlap, the problem of using this method for in-detail radius adjustments is illustrated in Figure 4.15. Two effects are immediately

apparent: First, despite the 100 mm target-disk radius, muons as far outside as $r = 200$ mm can be stopped. Second, even muons at positions well within the target-disks might pass unstopped despite their material being sufficient to stop them. Both effects are consequences of the helical motion of muons due to Lorentz forces from the approximately constant 1 T magnetic field around CyDet. Therefore, it is now time to take a closer look at muon helices.

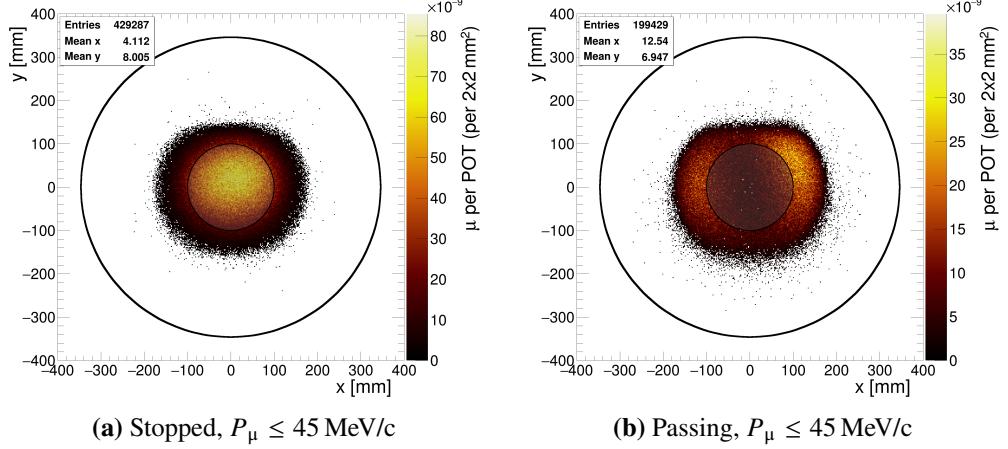


Figure 4.15.: Comparison of the x - y -distributions of (a) stopped and (b) passing muons. A filter is applied to only select stoppable muons with less than 45 MeV/c to prevent clutter from the large fraction of unstopped high-momentum muons. Overlaid circles at $r = 346$ mm and $r = 100$ mm illustrate the inner radius of CTH and the target-disks, respectively. The data corresponds to 60 bunches at 16×10^6 POT.

The calculation of a charged particle's helix radius has been established since the early 20th century. Accounting for the electric charge q , the magnetic field strength B , and the component of the relativistic particle momentum perpendicular to the magnetic field p_\perp , the Lamor radius (also known as Gyroradius) is defined as:

$$r_{\text{Lamor}} = \frac{\gamma \cdot m v_\perp}{|q| \cdot B} = \frac{p_\perp}{|q| \cdot B} = \frac{p_\perp}{1.602 \times 10^{-19} \text{ C} \cdot 1 \text{ T}}. \quad (4.4)$$

Therefore, in the limit of a purely perpendicular momentum, the maximum radius of a 25 MeV/c (50 MeV/c) particle corresponds to a Lamor radius of around 80 mm (160 mm).

In practice, however, the helix radii of muons traversing CyDet are much smaller. After the 90° transport solenoid, a significant fraction of the muon momentum is forward-directed and parallel to the magnetic field. Figure 4.16a shows the distribution of helix radii obtained from approximating muon tracks in ICEDUST. Compared to the above calculation, which assumed a purely perpendicular momentum, the most probable helix radius is just a quarter of that value, with the right tail of the distributions barely reaching half its size.

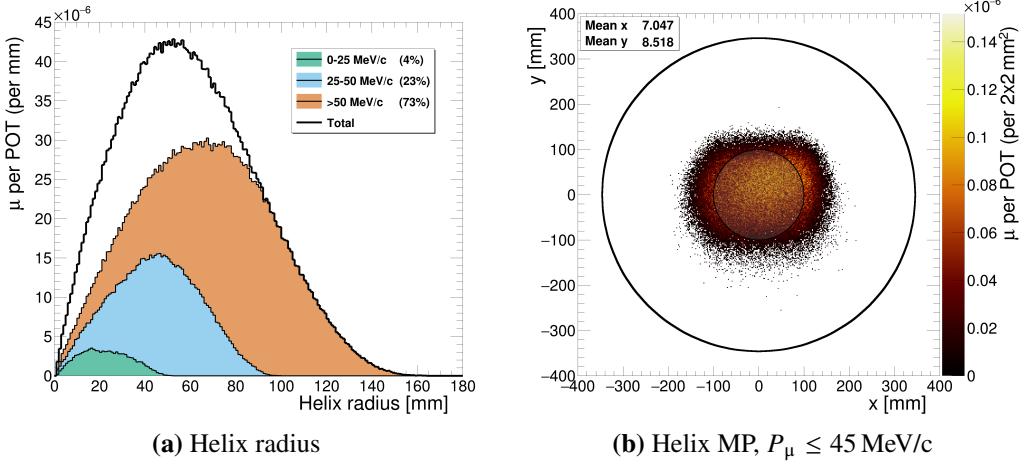


Figure 4.16.: Helix radius (a) and middle point (b) distribution of muons traversing CyDet. To avoid clutter of the middle point plot from unstopped high-momentum muons, only relevant muons with less than 45 MeV/c are included. Muon trajectory points were recorded at target-disk positions, projected onto the x - y -plane, and approximated with a custom circle fit using χ^2 optimization. The simulation was performed without target-disks to avoid bias from stopped muons, and the data corresponds to 60 bunches at 16×10^6 POT.

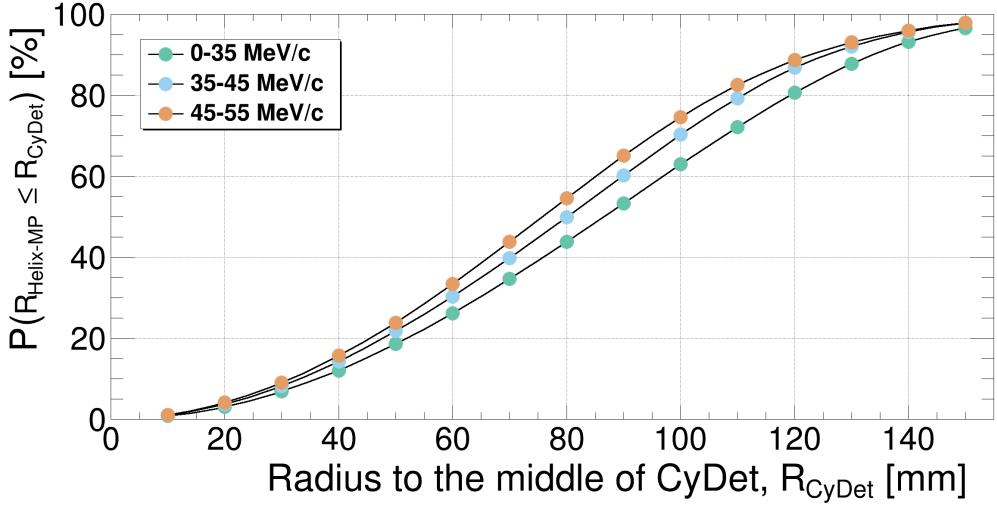


Figure 4.17.: Summary of Figure 4.16, showing the percentage of muon helix middle points confined within a given radius. As defined in Section 4.2.1, the CyDet center aligns with the center of the muon-stopping target.

The corresponding middle points of these helical muon tracks are shown in Figure 4.16b. Note the resemblance to the arriving muon profile shown in Figure 4.15. While the helix middle points are slightly more confined around the CyDet center, the general Gaussian profile and slight offsets discussed in the previous section can be rediscovered. The result is summarized in Figure 4.17. Assuming that over the entire muon-stopping target length of 800 mm, the shift of the helix center is negligible and each muon completes one full rotation, a few simplified conclusions can be drawn:

1. Approximately 70 % of muons with less than 45 MeV/c have helix middle points confined within the 100 mm default muon-stopping target radius. Consequently, this large fraction will traverse *at least* half of the employed target-disks. On the contrary, 30 % will at least miss half of them.
2. Approximately 25 % of muons with less than 55 MeV/c have a helix middle point smaller than 50 mm. The most probable helix radius in this momentum region is around 50 mm. Hence, these muons likely traverse all target-disks.
3. Approximately 20 % of very low-momentum (below 25 MeV/c) muons have helix middle points further outside than 120 mm. In combination with their helix radii peaking around only 20 mm, they are very likely to miss all target-disks.

Although the assumption of one complete rotation simplified the situation, the conclusions can be tested by comparing the momentum distributions of stopped muons from different muon-stopping targets. Figure 4.18 shows four cases: a muon-stopping target configuration with 17 or 33 total number of target-disks, with a radius of 100 mm or 300 mm¹. While the material of 17×0.2 mm target-disks is sufficient to stop muons up to approximately 50 MeV/c (compare Figure 4.18b), the limited 100 mm radius shifts the stopped muon peak downwards by approximately 10 MeV/c. The peak's position corresponds to muon momenta that can be stopped in half the material, and its height roughly matches those 70 % of muons from category one. Comparison with the 33 target-disk case yields identical results, whereby the peak shifts down from 60 MeV/c to the former maximum of 50 MeV/c. For both small-radius configurations, the tail toward higher momentum corresponds to muons in category two. Lastly, category three is visible as a light-grey band above the color-coded stacked histogram. In fact, in each momentum interval, about 20 % of muons cannot be stopped within 100 mm target-disks due to too-large helix middle points. On the contrary, the increase in the number of stopped muons, achievable *only* through larger target-disk radii, is also limited to this maximum of 20 %.

In summary, a 100 mm radius is sufficiently large that the majority of relevant low-momentum muons interact with at least *some* of the installed target-disks, thus have the potential to be stopped. The approach of a constant radius but thicker or additional target-disks will be explored in Section 4.6 and 4.7. Approximately 90 % of muons have

¹The larger radius corresponds to the inner radius of CTH minus a small length reserved for the muon-stopping target holding structure. It is large enough that muons below 60 MeV/c will traverse all employed target-disks.

helix middle points smaller than 120 mm. Combined with the fact that the majority of muon with momenta lower than 45 MeV/c have a helix radius smaller than 80 mm, the largest reasonable target-disk radius to stop such muons is around 200 mm. Of course, the improvements approaching this extreme will slowly plateau. However, judging solely from the muon-stopping perspective, target-disk radii up to 150 mm seem feasible.

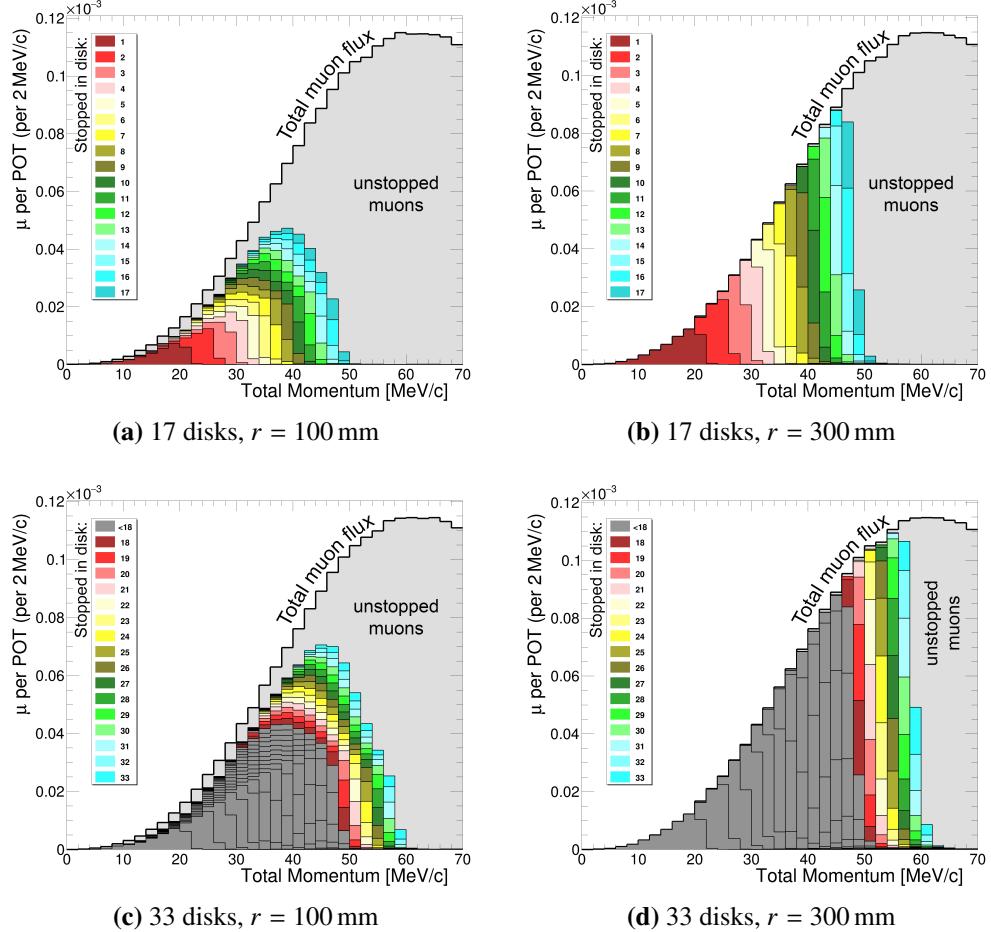


Figure 4.18.: Comparison of momentum distribution of stopped muons between four configurations. Stopped muons are color-coded according to their stopping-disk. Top-left equals the default configuration discussed in Section 4.1. The target-disks were uniformly spaced within the 800 mm muon-stopping target space of CDC, resulting in a disk-to-disk distance of 50 mm(25 mm) for the 17(33) configuration. Implications of a tighter disk spacing will be discussed in Section 4.8. Shown data correspond to 60 bunches at 16×10^6 POT. For discussion see text.

4.5.2. The Effect on Signal Acceptance

Changes in the target-disk radius affect the signal acceptance due to conversion electrons having to traverse more or less amount of material before reaching the CDC. Consequently, understanding electron energy loss and the relation between traversed material and target-disk radius is essential to form conclusions.

In contrast to heavy particles like muons, the energy loss of 100 MeV/c electrons is dominated by Bremsstrahlung due to their comparably small mass. The *mean* energy loss of electrons from radiation per interaction step can be calculated by the equation:

$$\begin{aligned} -\left(\frac{dE}{dx}\right)_{\text{rad}} &= NE_0\Phi_{\text{rad}}, \text{ with} \\ \Phi_{\text{rad}} &= \frac{1}{E_0} \int h\nu \cdot \frac{d\sigma}{dv}(E_0, \nu) dv. \end{aligned} \quad (4.5)$$

Here, E_0 is the initial total energy, N the number of atoms per cubic centimeter, σ the Bremsstrahlung cross-section, and $h\nu$ the Planck constant and frequency, representing the energy of the emitted photon [82, 83].

Although the range of electrons can be subsequently calculated by integrating Formula 4.5, measured values generally differ significantly due to electron multiple scattering, as well as the statistical nature of the number of collisions and transferred energy per collision. Especially for thin targets¹ the calculation of the number of collisions becomes further complicated. Typical energy loss distributions for thin targets are depicted in Figure 4.19. Note that rare high-energy-loss interactions in single collisions or from Bremsstrahlung form a long tail on the high-energy side of the distribution, shifting the mean energy loss toward higher energies and away from the most probable energy loss at the peak position.

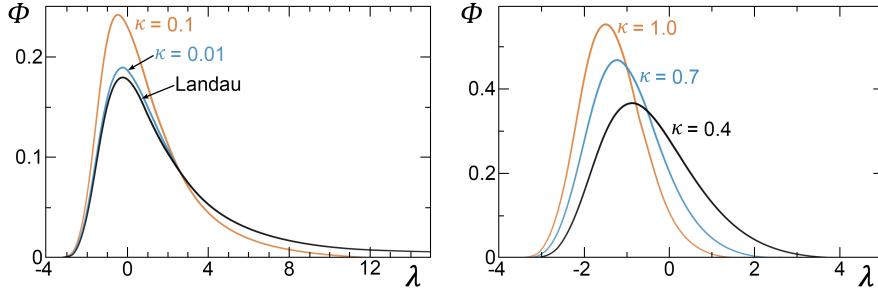


Figure 4.19.: Vavilov distributions for various κ . Larger κ equals thicker absorbers. For $\kappa \rightarrow 0$ ($\kappa \geq 1$), the distribution approaches a Landau (Gaussian) distribution. Adopted from Seltzer and Berger [84].

¹The parameter to distinguish different thickness regions is κ , as defined in [83] or [84]. The thin absorber region is generally considered to be $\kappa < 10$. For ultrarelativistic electrons and a 0.2 mm aluminium target-disk κ corresponds to less than 0.0001.

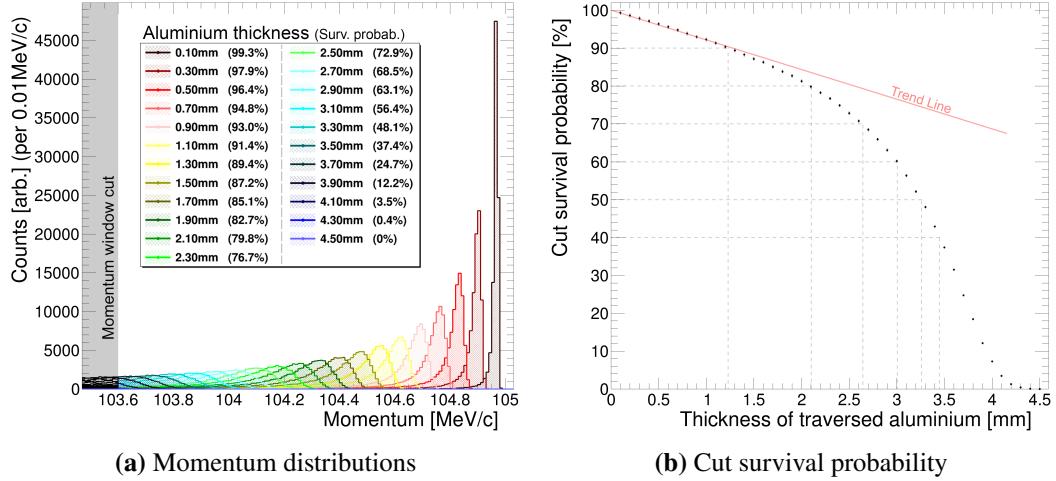


Figure 4.20.: Left: Conversion electron momentum distribution after traversal of X mm aluminium (color-coded). The lower edge of COMET’s momentum ROI at 103.6 MeV/c is highlighted in grey. The fraction of electrons above this threshold is quoted as the survival probability. Right: Summary of the survival probability over a larger range of thicknesses. A linear trend is observed for smaller aluminium layers, which breaks down at around 1.25 mm. Note that no detector resolution was applied for this study.

In consequence, the electron energy loss when traversing aluminium is determined through MC simulation rather than manual calculation. Results are shown in Figure 4.20a. Note how the remaining conversion electron momentum distributions correspond to mirrored Vavilov distributions. The conversion electron *survival probability*, defined as the fraction of electrons surviving COMET’s 103.6 MeV/c momentum window cut condition, is included in the legend and summarized in Figure 4.20b. After 1.25 mm of aluminium, 90 % of signal remains within the momentum ROI. The probability decreases to 80 % at 2.1 mm, after which it drops significantly due to larger parts of the distribution’s peak area now falling below the cut-off value.

To utilize the survival probability, a proper understanding of the total material a conversion electron has to traverse through target-disks is necessary. In accordance with Figure 4.21, it can be divided into three categories:

1. The electron has to leave its *disk of creation*.
2. The electron might move forward/backward enough to additionally hit *follow-up target-disks* before reaching a large enough radius to exit the tube.
3. In case of multiple rotations through CDC, the electron can hit *additional target-disks* during later turns.

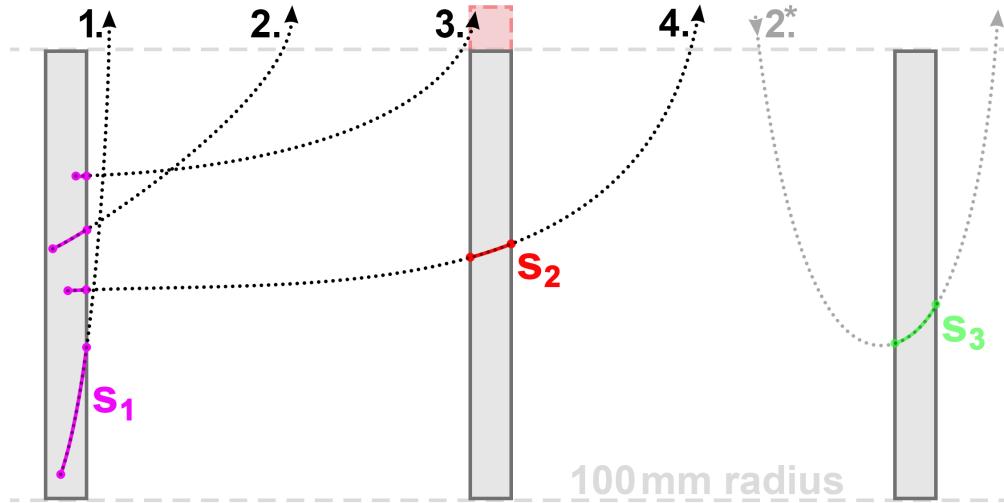


Figure 4.21.: Paths contributing to a conversion electron's traversed material through the muon-stopping target. The schematic drawing includes four example electron trajectories (black dotted lines) starting from within the leftmost target-disk (grey rectangle). An example multi-turn event is indicated as a grey dotted line (2*). Two grey dashed lines indicate the 100 mm tube containing all target-disks. A red rectangle extending the middle target-disk demonstrates the impact of an increased radius. The segments contributing to the total distance traversed through the muon-stopping target are color-coded. See text for further details and a discussion of each numbered segment.

For the sake of simplicity of this discussion, the effect of multiple-turns (category three) is disregarded. The electron momentum is recorded and checked against the ROI cut condition at its first entry in CDC. As stated in Section 4.2.3, two-thirds of all trajectories are single-turn and not affected by this simplification. For the other third, the decision implies that the electron's momentum can be sufficiently determined during its first rotation through CDC, and additional hits during later turns will not significantly affect the track-finding and -fitting performance.

The relation between the first and second categories is illustrated in Figure 4.22. First, conversion electrons with very small longitudinal momentum can potentially travel multiple millimeters within their disk of creation. However, to reach the CTH and fulfill COMET's trigger condition, a certain level of momentum in the forward or backward direction is required. Hence, the total path length of accepted electrons inside their disk of creation is limited to below 3 mm. That means relevant conversion electrons will quickly exit their initial target-disk through either side – an effect desired and facilitated by the thin target-disk design. Consequently, an increase in target-disk radius has no effect on this first category.

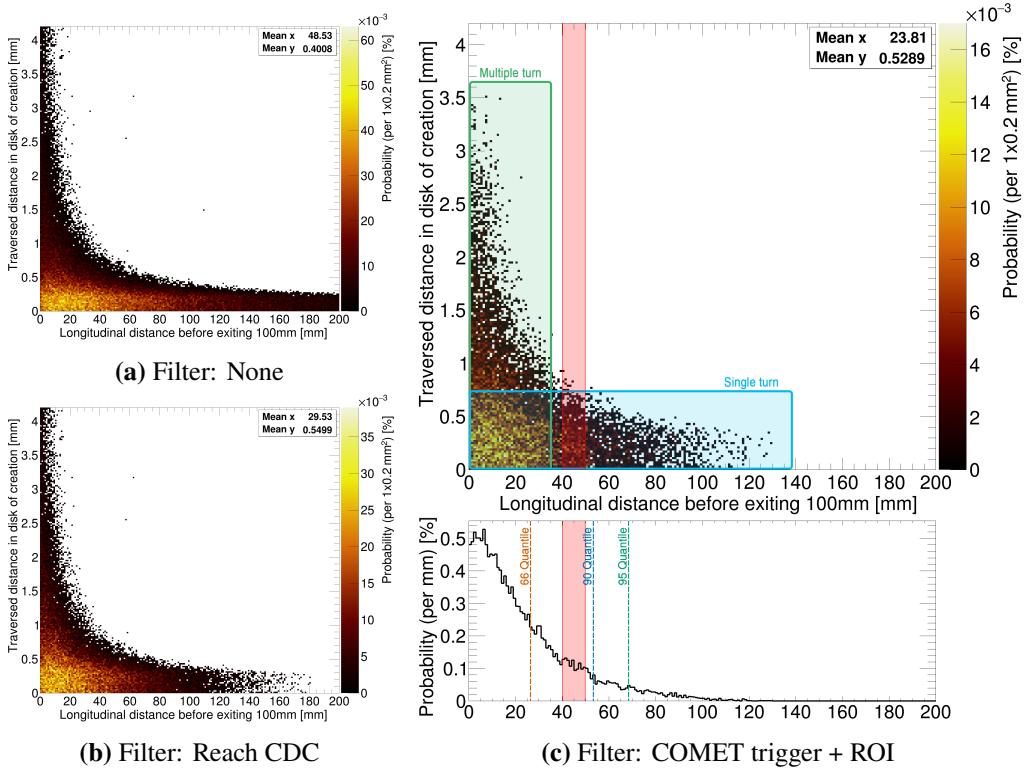


Figure 4.22.: Relation between a conversion electron's distance traveled through its disk of creation and the 100 mm tube containing all target-disks. Signal electrons were generated from prerecorded muon stopping distribution, assuming isotropic momentum directions (see Section 4.2.3). (a) Unfiltered result. (b) Required to reach the first CDC layer. (c) Fulfilled COMET trigger conditions and momentum window cut. For the latter, a projection of the 2D data onto the x-axis is included at the bottom. Conversion electrons with large transversal momentum components (green square) traverse multiple millimeters through their target-disk of creation but are unlikely to hit follow-up target-disks before reaching the CDC for the first time. However, due to their small longitudinal momentum, they often perform multiple rotations before triggering CTH. Single-turn events (blue square) have a larger longitudinal momentum component. However, still over 66 % exit the 100 mm radius before hitting follow-up target-disks. Electrons affected by a change in target-disk radius amount to approximately 2.5 % and are highlighted in a red area.

Second, conversion electrons with high longitudinal momentum can travel significant distances along the beam axis before reaching $r = 100$ mm. However, in cases of large forward or backward momentum, such electrons cannot reach the CDC (compare Figure 4.22a and 4.22b). Consequently, approximately 66 % of accepted electrons reach the 100 mm radius after a longitudinal distance of 30 mm (compare Figure 4.22c). Of the remaining 34 %, most will travel z -distances much larger than the 50 mm target-disk spacing and thus hit follow-up target-disks regardless of exact radius.

Finally, this leaves only the small fraction of electrons that reach $r = 100$ mm just before arriving at the follow-up disk to be affected by a radius change. This region is highlighted with a red band in Figure 4.22c. For the default 50 mm target-disk spacing, the exact number of follow-up target-disk hits using different target-disk radii are quantified in Table 4.6. Per each 10 mm increase in radius, the number of electrons hitting follow-up target-disk increases by approximately 2.5 %.

Comparing the 100 mm and 150 mm case in Table 4.6, the number of conversion electrons reaching CDC without passing through follow-up target-disks during their first turn decreases from 80.6 % to 65.2 %. However, still 89.2 % only traverse one additional target-disk, which adds about 0.2 mm of traversed aluminium. Given that most of these forward-directed conversion electrons travel very small distances through their disk of creation (compare Figure 4.21), their total material traversed is expected to remain below around 0.5 mm. At this value, the cut survival probability is still around 96 %. Consequently, the effects of an increased radius on signal acceptance are negligible compared to the significant gains in the number of muons stopped.

Table 4.6.: Number of follow-up target-disk hits for varying target-disk radii.

The total number of target-disks was fixed at 17, corresponding to a 50 mm distance. Only trajectories reaching at least the first CDC layer are taken into account. Values are presented in percentages and separated into follow-up hits from all turns (including multiple rotations) and only during the electron's first rotation.

Radius	During all turns [%]				Only 1 st turn [%]			
	0	1	2	3+	0	1	2	3+
60 mm	81.1	17.8	0.9	0.2	93.5	6.4	0.1	0.0
80 mm	70.1	26.9	2.7	0.3	86.9	12.0	1.0	0.0
100 mm	60.7	32.4	6.1	0.9	80.6	16.2	3.0	0.2
110 mm	55.7	34.6	8.2	1.5	77.3	18.1	4.1	0.5
120 mm	51.6	36.3	10.1	2.0	74.8	19.5	4.8	0.8
130 mm	47.3	37.3	12.4	3.1	70.9	21.7	5.9	1.5
140 mm	43.5	38.0	14.4	4.1	67.9	23.0	7.1	2.0
150 mm	40.1	37.9	16.6	5.4	65.2	24.0	8.1	2.7

4.5.3. The Effect on Detector Background

An increase of target-disks radius inevitably leads to a higher detector background due to an increase in the number of interactions taking place in the muon-stopping target. Especially CDC background from muon interactions, amounting to approximately one-third of all CDC hits, cannot be reduced as it is a direct consequence of the increased muon slowdown if one wants a higher number of muon stops. Exact values for CDC occupancy and CTH rate are provided in the performance overview (see Table 4.7). For now, two more unintuitive background sources will be discussed.

Bremsstrahlung from Beam Electrons The first background source to consider stems from electrons traveling alongside the beam's muons. Electrons traversing the disks can produce large amounts of high-energetic Bremsstrahlung. Bremsstrahlung photons have a comparably low impact on CDC occupancy and are efficiently suppressed by the CTH 4-fold coincidence requirement. However, as discussed in detail in Chapter 6, it has a particularly detrimental effect on the photon-sensitive HPGe detector used as a beam monitor and must not be disregarded.

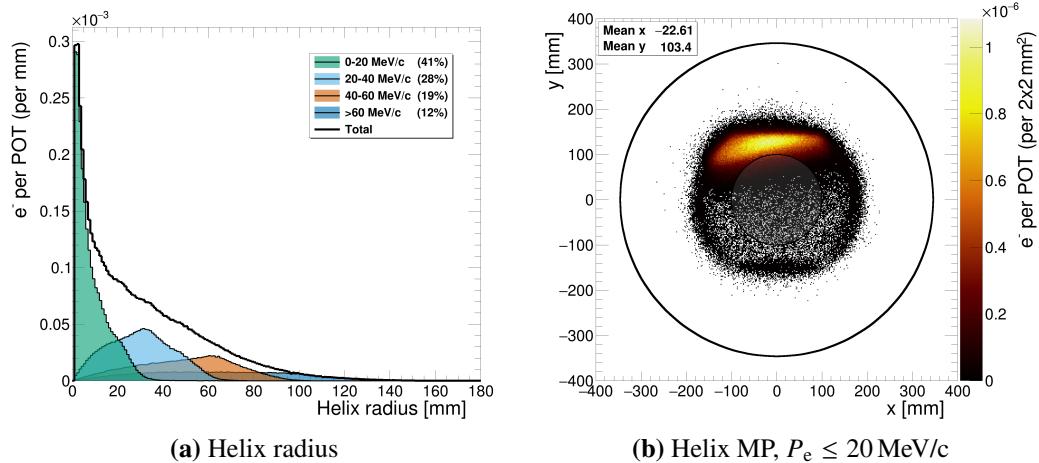


Figure 4.23.: Helix radius (a) and middle point (b) distribution of beam electrons traversing CyDet. The middle point plot filters for electrons with momenta less than 20 MeV/c to highlight the distribution's close resemblance to the monitor profile recorded in front of the muon-stopping target (see Figure 4.24a). Electron trajectory points were recorded at target-disk positions, projected onto the x - y -plane, and approximated with a custom circle fit using χ^2 optimization. The simulation was performed without target-disks to avoid bias from stopped or scattered electrons, and the data corresponds to 14 bunches at 16×10^6 POT.

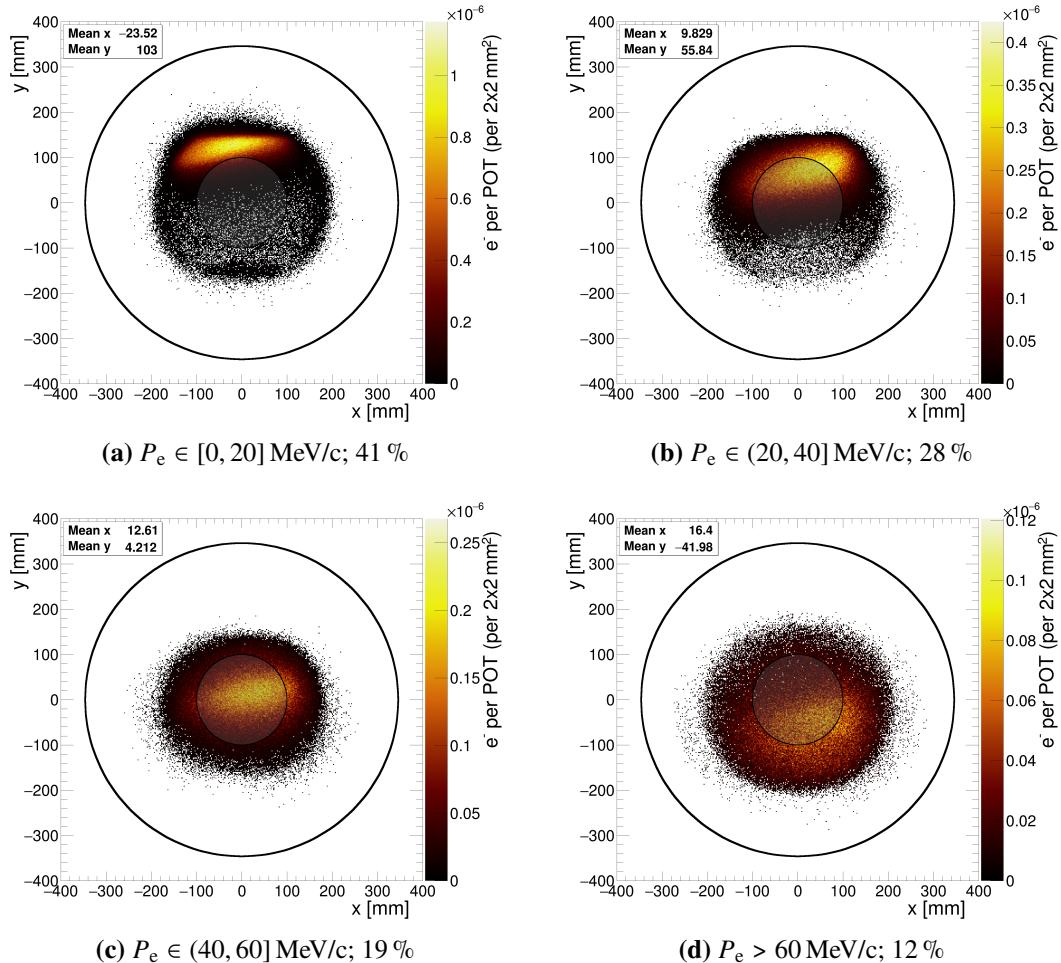


Figure 4.24.: Beam electron x - y -distribution in front of the muon-stopping target for different momentum ranges. Note that histogram z -axes are adjusted to account for different absolute numbers of electrons. The fraction of electrons in each momentum range is stated in percentage in the subcaptions. The data corresponds to 14 bunches at 16×10^6 POT. Overlaid circles at $r = 346$ mm and $r = 100$ mm illustrate the inner radius of CTH and the target-disks, respectively.

Like muons, beam electrons are subject to the CDC's magnetic field, thus move in helical trajectories because of the Lorentz force. In accordance with Figure 4.16 for muons, Figure 4.23 displays the helix radius and middle point distribution for electrons arriving at CyDet. Note that compared to muons, most electrons have an incredibly small helix radius because of their small momentum. Hence, in first approximation their helical movement can be disregarded and assumed to match simple straight lines.

Beam electron profiles recorded with the monitor plane in front of the muon-stopping target are shown in Figure 4.24. A strong momentum dispersion due to the beam transport can be observed. Beam electrons with higher momentum are comparatively spread out. Their most probable beam spot overlays the target-disks and a large fraction will hit the muon-stopping target independently of the exact target-disk radius. However, the largest fraction of beam electrons surviving the transport section, 41.4 %, has less than 20 MeV/c (see Figure 4.24a). Their downward drift during beam transport is small, and their most probable position arriving at the muon-stopping target lies around $y \approx 130$ mm. As these low-momentum muons move in approximately straight lines, they closely pass by the 100 mm target-disks. However, increasing the radius by just 10 mm would cause them to pass through the target disks, significantly raising the Bremsstrahlung background.

Direct Hits from Decay In Orbit Electrons Summarizing the information of the previous paragraph, 88 % of beam electrons have a momentum of less than 60 MeV/c. Their helix radii stay below 100 mm, and their helix middle points are closely confined around the CyDet center. Thus, these beam electrons can not directly hit the CDC.

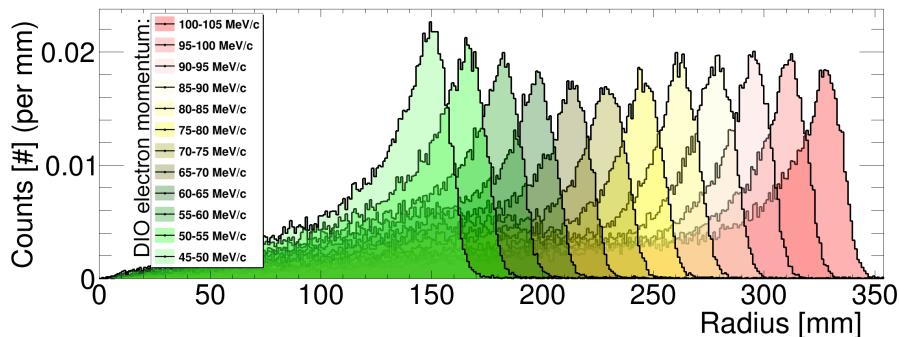


Figure 4.25.: Helix radius distributions from DIO electrons inside CyDet. Each 5 MeV/c momentum interval is normalized to one and corresponds to 62 500 simulated events. Electrons were generated inside a 100 mm radius around the CyDet center with an isotropic momentum direction. The distribution's right tail reaches up to the theoretical maximum Lamor radius (compare Formula 4.4), with the distribution's peak approximately 20 mm below. The long tail toward lower radii stems from electrons with low transversal momentum components.

The situation is different for DIO electrons. First, they do not rotate around but originate from the muon-stopping target. In consequence, their helix middle point is positioned between target-disk and CDC, making the helix diameter or twice the radius their parameter of interest. Second, the DIO electron momentum can reach up to approximately 105 MeV/c (see Section 1.1.2) and is isotropically distributed, allowing for comparatively large transversal momentum components. Figure 4.25 shows simulated radius distributions for DIO electrons of varying momentum. Note how the peak of the distributions lies very close to the calculated Lamor radii in the approximation of a purely perpendicular momentum. The peak position and maximum possible radius increase linearly with momentum because of its linear contribution in Formula 4.4. All distributions feature a long tail toward lower radii from electrons with high longitudinal momentum components. The tail-to-peak ratio amounts to roughly 1/3 to 2/3.

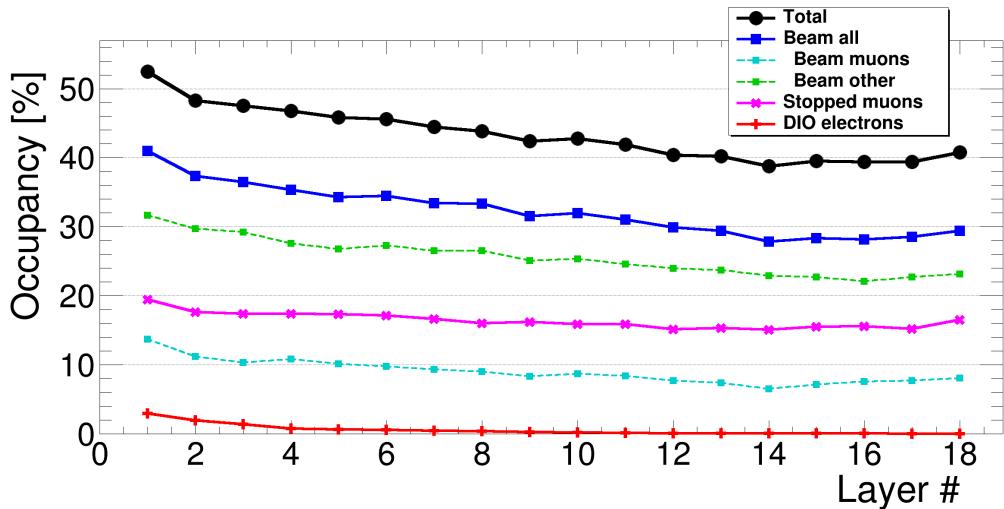


Figure 4.26.: CDC layer occupancy for enlarged $r = 200$ mm target-disks. The data is averaged over 60 MC5 bunches. Statistical error bars are included but are too small to see. Total CDC occupancy amounts to 43 %. Smaller layer numbers correspond to smaller radii in CDC, thus have higher occupancy due to a smaller number of wires and a closer position to the muon-stopping target. Compared to the default configuration (see Figure 4.12), the occupancy increases are: DIO hits (+1 % points), muon stops (+5 % points); increased interaction from unstopped high-momentum muons (+2 % points); additional pion interactions (+8 % points)

As discussed in Section Section 1.1.2, the DIO rate gradually increases up to around 45 MeV and drops rapidly afterwards. Hence, over 99 % of DIO electrons have a helix radius of less than 170 mm and, by design, they cannot reach the first layer of CDC located

at 496 mm. For target-disk radii larger than 150 mm, this intrinsic CyDet DIO suppression will become less effective. Figure 4.26 shows the single layer occupancy for target-disks with a radius of 200 mm. Compared to the default configuration (see Figure 4.12), the DIO occupancy in the first layers doubles. However, it remains surprisingly negligible compared to increases from muon interactions (+5 %) or the beam other component (mainly pions; +10 %). The reason is shown in Figure 4.27. Although the increased radius helps to slow down high-momentum muons spiraling to larger radii, most are stopped around a target-disk's center, roughly matching the muon beam profile shown previously in Figure 4.13.

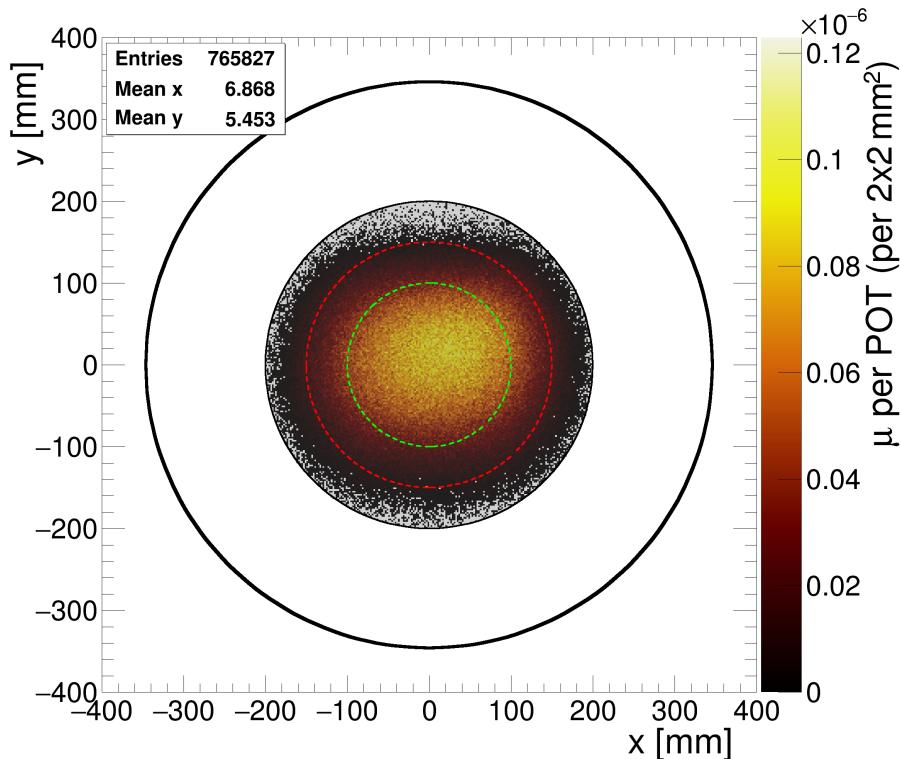


Figure 4.27.: Muon stopping distribution for enlarged $r = 200$ mm target-disks.

The data corresponds to all muons stopped over 60 bunches at 16×10^6 POT within the 17×0.2 mm target-disks, projected onto the x - y -plane. Outer black circle: CTH inner radius. Inner black circle: Representation of the muon-stopping target (disks). Dashed red circle: 150 mm radius. Dashed green circle: 100 mm radius. The profile closely matches the arriving muon profile (compare Figure 4.5)

4.5.4. Summary and Performance Overview

The performance of nine configurations with varying target-disk radius is summarized in Table 4.7. Increasing the target-disks radius up to 150 mm improves the overall available signal by 42 %. The increase is sufficient to stop a large fraction of low-momentum muons, which previously did not interact with enough target-disks due to their too far outside helix middle points. For every 10 mm increase in radius, approximately 10 % more muons can be stopped. The improvements become less prominent toward larger radii and start to stagnate after 200 mm.

The effect of larger radii on the signal acceptance is negligible. Conversion electrons quickly exit their target-disks of creation through its sides, and most reach large enough radii before hitting follow-up target-disks. Even for enlarged radii up to 150 mm, the total amount of material a conversion electron traverses through target-disks stays below 1.25 mm aluminium, thus 90 % survive the [103.6, 106.0] MeV/c momentum cut.

Total CDC occupancy gradually increases but remains below its maximum value of 50 %. Despite this manageable background and the large improvements in available signal, larger radii are strongly disfavored due to their severe consequences for the muon-stopping monitor. Even for a small radius increase by only 10 mm, a significant fraction of low-momentum beam electrons will start to hit target-disks, leading to a steep increase in Bremsstrahlung background. Fortunately, only 20 % of arriving muons cannot be stopped within a default 100 mm radius. The remaining 80 % interact at least with some of the target-disks. Consequently, the following sections will explore increasing the number or thickness of target-disks while leaving the radius unchanged.

Table 4.7.: Performance overview of configurations with varying target-disk radii.

Total number of target-disks and thickness was kept constant at 17×0.2 mm.

The data is averaged over 14 MC5 bunches at 16×10^6 POT per bunch. Values are calculated following the procedure described in Section 4.2.

Disk radius	$N_{\mu/\text{bunch}}^{\text{stopped}}$ [per bunch]	$A_{\mu \rightarrow e}^{\text{GM*ROI}}$ [%]	normalized signal, K_{NS}	total CDC occupancy [%]	CTH rate [per bunch]
60 mm	3375.49	18.59	0.45	21.97	1463.4
80 mm	5569.37	17.98	0.72	26.66	1605.3
100 mm	7838.68	17.79	1.00	32.07	1736.2
110 mm	8907.65	17.48	1.12	34.65	1857.3
120 mm	9832.96	17.33	1.22	36.04	1885.1
130 mm	10675.13	16.87	1.29	38.84	1926.6
140 mm	11343.25	17.00	1.38	39.45	1967.1
150 mm	11858.31	16.71	1.42	40.74	2013.6
200 mm	12755.73	15.98	1.46	42.92	2096.6

4.6. Material Budget: Target-Disk Thickness

Up to this point, motivation has been provided for the general shape of the muon-stopping target, namely: multiple, thin target-disks aligned in a tube at the center of CyDet. Larger radii for the target-disks were discouraged due to a dangerous rise in Bremsstrahlung from beam electrons, which would overwhelm the muon-stopping monitor discussed in Chapter 6. Alternatively, one can increase the material within the path of the muon beam by increasing the target-disk number or thickness. This section aims to identify the more effective approach between these two options.

Throughout this section, the total amount of aluminium is kept constant at 6 mm. Since muons stop uniformly along the thickness of a target-disk and the muon profile remains constant over the length of the muon-stopping target, the number of stopped muons is identical for a solid block of aluminium and multiple spaced-out blocks adding up to the same thickness. Concrete numbers verifying this statement will be presented in the performance overview in the summary (see Table 4.9). For now, this limitation helps to focus on the main difference between thicker and thinner target-disks, specifically the implications for conversion electrons and the impact on signal acceptance.

As discussed in Section 4.5.2, the survival chance of conversion electrons, and thus signal acceptance, directly depends on the total amount of material traversed before reaching CDC. It was broadly categorized into *disk of creation* and *follow-up target-disk* hits, as shown in Figure 4.21. Thicker target-disks will cause conversion electrons to traverse longer distances before exiting their disk of creation. However, if each target-disk contains more material, fewer total target-disks are required. Thus, larger spacing between them can be employed, and the probability of follow-up hits decreases. Naturally, the opposite is true for thinner target-disks.

Figure 4.28 provides a comparison of traversed target material for six different configurations of number of target-disk times thickness. The top shows traveled distances through an electron's disk of creation. Accounting for all emitted conversion electrons (left), the distributions have a plateau at low distances approximately corresponding to the target-disk thickness. If electrons are required to reach CDC (right), highly forward or backward-directed electrons are removed. Consequently, the relative amount of electrons with high transverse momentum increases, and the low-distance plateau gets enlarged by roughly 50 %. Additionally, the peak-to-tail ratio worsens, and 90th percentiles increase significantly. The effect is particularly detrimental for larger target-disk thicknesses. Not only is the shift here more prominent, but a significant fraction of conversion electrons now traverse distances larger than 2 mm of aluminium, where the survival probability starts to drop off rapidly (compare Figure 4.20 of Section 4.5.2).

The traversed material through follow-up disks is shown in the middle of Figure 4.28. When hitting follow-up target-disks, most electrons traverse their full thickness, translating to a distinct rising edge at the corresponding thickness of each configuration. The long tail toward the right corresponds to electrons traveling through target-disks at an angle, resulting

4. The Muon-Stopping Target

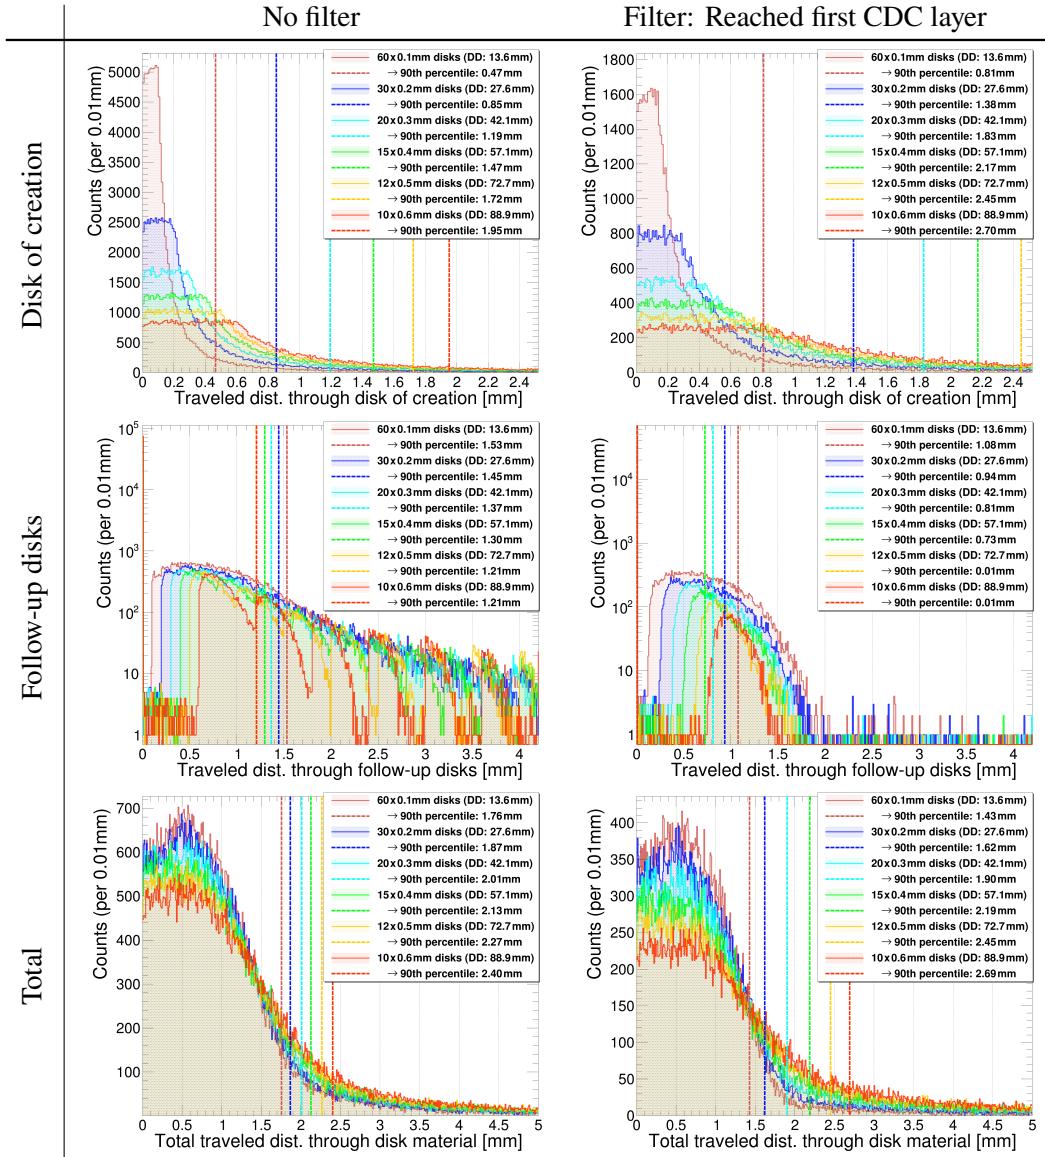


Figure 4.28.: Amount of traversed target material for varying number and thickness of target-disk. Left: Without filter. Right: Filtered for conversion electrons reaching the CDC. Each configuration's number of target-disks, thickness and disk distance (DD) are stated in the legend. The 90th percentiles are overlaid in matching colors. See text for further details.

in slightly more material traversed. The structure is followed by another rising edge when the next target-disk is hit. All these features are particularly prominent for thicker target-disk configurations and become blurred for thinner ones. As before, requiring electrons to reach the CDC removes those with predominantly forward-directed momentum and a large number of follow-up target-disk hits, limiting the large-distance tail and causing a cut-off around 1.8 mm. Note that for configurations with distances larger than 70 mm, 90 % of electrons encounter zero follow-up target-disks. The exact numbers are quantified in Table 4.8.

Table 4.8.: Number of follow-up target-disk hits for varying total number of target-disks. Spacing was uniform along the 800 mm inside CDC, and distances were adjusted accordingly as provided in the left row. To keep the total amount of material constant at 6 mm, target-disk thickness was varied from 0.1 mm to 0.6 mm. Based on Section 4.5, also target-disk radius was kept constant at 100 mm. Only trajectories reaching at least the first CDC layer are considered. Values are presented in percentages and separated into follow-up hits from all turns (including multiple rotations) and only during the electron's first rotation.

Number of disks (disk distance)	During all turns [%]				Only 1 st turn [%]			
	0	1	2	3+	0	1	2	3+
60 (13.6 mm)	19.2	19.9	17.6	43.0	38.9	21.6	13.5	25.9
30 (27.6 mm)	37.3	33.3	18.8	10.5	60.9	23.0	10.0	6.1
20 (42.1 mm)	53.6	34.7	9.6	2.0	75.4	19.0	4.8	0.8
15 (57.1 mm)	65.6	29.9	3.9	0.6	84.5	13.9	1.6	0.0
12 (72.7 mm)	74.5	23.3	1.7	0.5	90.7	9.0	0.3	0.0
10 (88.9 mm)	80.3	18.1	1.1	0.4	94.7	5.3	0.0	0.0

Both datasets are combined for the conclusion shown at the bottom of Figure 4.28. For thick target-disks, the distribution is dominated by the traveled distance through their disk of creation because of their negligible amount of follow-up hits. In the case of thinner target-disks, the pronounced peak at low distances gets significantly broadened after accounting for follow-up hits. However, compared to thicker target-disk configurations, it stays slightly more prominent, translating to overall smaller 90th percentiles. Consequently, the distributions of total material traversed seem similar, but thinner target-disks show better performance.

4.6.1. Summary and Performance Overview

The performance overview for the tested configurations discussed in regards to Figure 4.28 is provided in Table 4.9. First, the assumption of a constant number of muons stopped is verified within statistical uncertainties. Second, also the total CDC occupancy shows no significant variation and stays approximately constant around 44 %. Lastly, the values of

the normalized signal confirm the conclusion of the preceding discussion: A large number of very thin target-disks perform significantly better than fewer but thicker target-disks. In numbers, the geometric acceptance including the ROI cut drops by approximately 5 % for each 0.1 mm increase in target-disk thickness.

Table 4.9.: Performance overview of configurations with varying number and thickness of target-disks. The total material in the beam was kept constant at 6 mm. Target-disks were uniformly spaced within the 800 mm muon-stopping target space of CDC, and their radius was fixed at 100 mm. The data is averaged over 14 MC5 bunches at 16×10^6 POT per bunch. Values are calculated following the procedure described in Section 4.2.

Disk thickness (number)	$N_{\mu/\text{bunch}}^{\text{stopped}}$ [per bunch]	$A_{\mu \rightarrow e}^{\text{GM*ROI}}$ [%]	normalized signal, K_{NS}	total CDC occupancy [%]	CTH rate [per bunch]
0.1 mm (60)	12960.00	18.31	1.70	45.44	2044.3
0.2 mm (30)	13020.87	17.37	1.62	43.71	2071.5
0.3 mm (20)	13036.80	16.56	1.55	44.81	2055.4
0.4 mm (15)	13035.17	15.50	1.45	43.16	2073.2
0.5 mm (12)	13046.60	14.52	1.36	44.91	2067.9
0.6 mm (10)	13057.36	13.62	1.28	42.92	2094.5

4.7. Material Budget: Total Number of Target-Disks

The last section demonstrated that many thin target-disks are preferred over fewer but thicker ones. The next step is to investigate the maximum feasible amount of material before losses due to the ROI cut or detector background become too severe. The study is performed using 0.2 mm thick target-disks, but can be converted to thinner ones by increasing their numbers accordingly. As before, COMET's main parameters are first examined individually, then combined, summarized, and presented as a performance overview at the end.

4.7.1. The Effect on Total Number of Muons Stopped

Ideally, the amount of employed material should be sufficient to at least reach the peak of the arriving muon flux around 60 MeV/c (compare Figure 4.2 of Section 4.1). Hence, Figure 4.29 shows the required number of target-disks to stop muons of varying momentum intervals. To facilitate the discussion and distinguish effects from the helical muon trajectory (see Section 4.5.1), the figure includes a comparison between the default 100 mm and an exaggerated 300 mm radius configuration.

Beginning with the large radius case, all muons of a given momentum interval are stopped within a comparatively confined range of disk numbers. For example, muons between 50 MeV/c and 55 MeV/c require at least 20 and at most 30 target-disks. The small difference arises from the statistical nature of muon interactions and the size of the momentum interval itself. If each muon were to traverse all employed target-disks, stopping 65 MeV/c (70 MeV/c) muons would require approximately $50(60) \times 0.2$ mm target-disk, corresponding to a total thickness of 10 mm (12 mm) aluminium.

However, in the case of a limited 100 mm radius, approximately 30 % of low- and 50 % of high-momentum muons require significantly more target-disks to be stopped. This high-stopping-disk-tail is a consequence of the helical muon trajectory, discussed in Section 4.5.1. Hence, while in theory only 40×0.2 mm target-disks are required, about twice as much (approximately 16 mm aluminium) will be necessary to stop muons up to 60 MeV/c if the default radius is left unchanged.

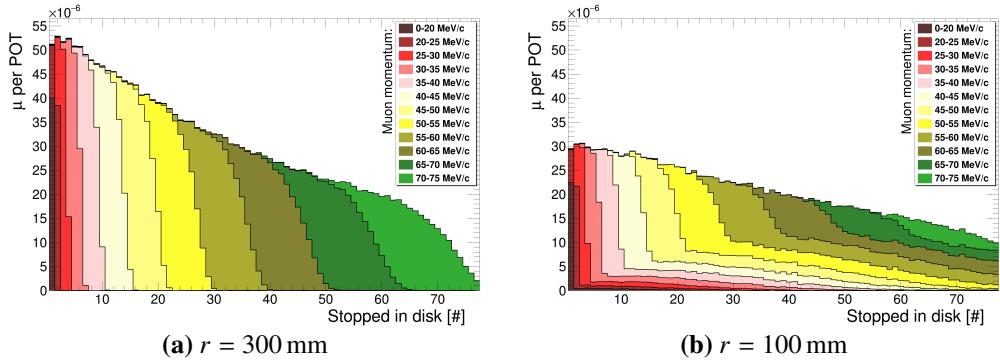


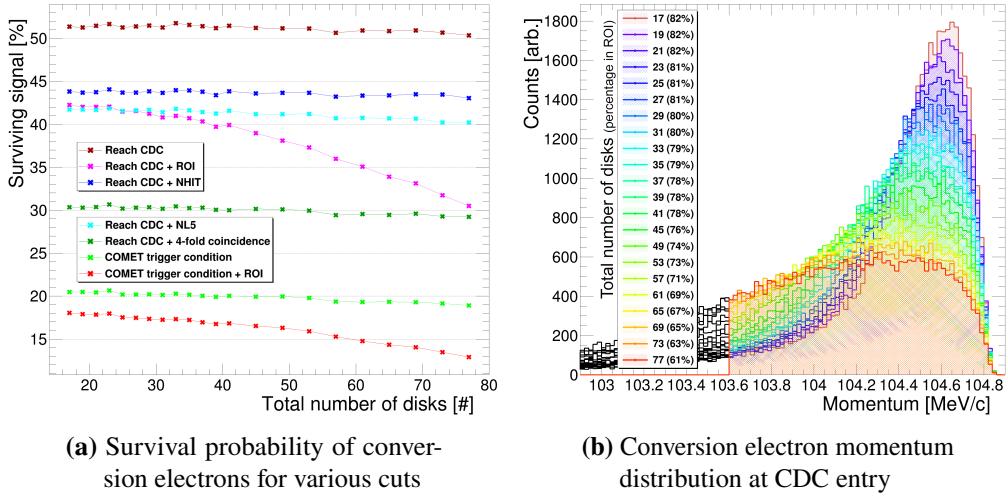
Figure 4.29.: Number of 0.2 mm thick aluminium target-disks required to stop muons of varying momenta. Muons below 20 MeV/c stop in the very first encountered target-disk and are grouped separately. Notably, the total area of each segment is different because the data is based on COMET’s expected muon flux and normalized to arriving muons per POT. Comparing the 300 mm (a) and 100 mm (b) radius case, approximately 30 % of low- and 50 % of high-momentum muons are smeared toward higher stopping-disk numbers due to the helical muon trajectory.

4.7.2. The Effect on Signal Acceptance

Losses in signal acceptance constrain improvements from an increase in the number of target-disk. Configurations with additional target-disks in the limited space of $L = 800$ mm within the CDC require tighter spacing. Figure 4.30a shows the effect such tighter spacing has on the applied track quality cuts defined in Section 4.2.3. Expectedly, the conversion electron’s momentum window cut survival probability drops significantly due to an increased energy loss from additional hits in follow-up target-disks. However, the efficiency loss of

other cut criteria caused by the scattering of electrons, after which they can not reach the detector anymore, remains between only 1-2 percentage points.

Corresponding momentum distributions at CDC entry for different numbers of target-disks are shown in Figure 4.30b. The fraction of electrons remaining inside the ROI in relation to the electrons reaching the CDC is indicated in the legend. For less than 30 target-disks, the distributions have a clear peak around 104.6 MeV/c. For a higher number of target-disks, increasingly larger fractions of the peak get smeared and fall below the 103.6 MeV/c ROI threshold. However, the momentum loss is small enough that even for the tightest shown target-disk spacing of around 10 mm (corresponding to 77 target-disks), 61 % of the signal can still be retained.



(a) Survival probability of conversion electrons for various cuts

(b) Conversion electron momentum distribution at CDC entry

Figure 4.30.: Effect of an increased total number of target-disks on CDC cut survival probability. Apart from an increased number and adjusted spacing, default parameters for radius and thickness of target-disks were applied. Left: Various combinations of CDC analysis cuts as defined in Section 4.2.3. Right: Conversion electron momentum distribution at CDC entry. Note that probabilities in the left plot are given in regards to all emitted conversion electrons, while the ROI survival chance stated in the legend of the right plot only accounts for electrons reaching the CDC.

The results from this discussion and the previous section are combined in Figure 4.31. A linear trend is noticeable for configurations with up to 40×0.2 mm target-disks for both the number of stopped muons and the percentage of surviving conversion electrons. Beyond this point, returns start to gradually diminish. However, the maximum normalized signal is not reached until approximately 60×0.2 mm target-disks, after which it begins to plateau. In conclusion, for target-disks with a 100 mm radius, configurations containing between 50 and 60×0.2 mm target-disks, or between 10 and 12 mm of aluminium, yield the best results.

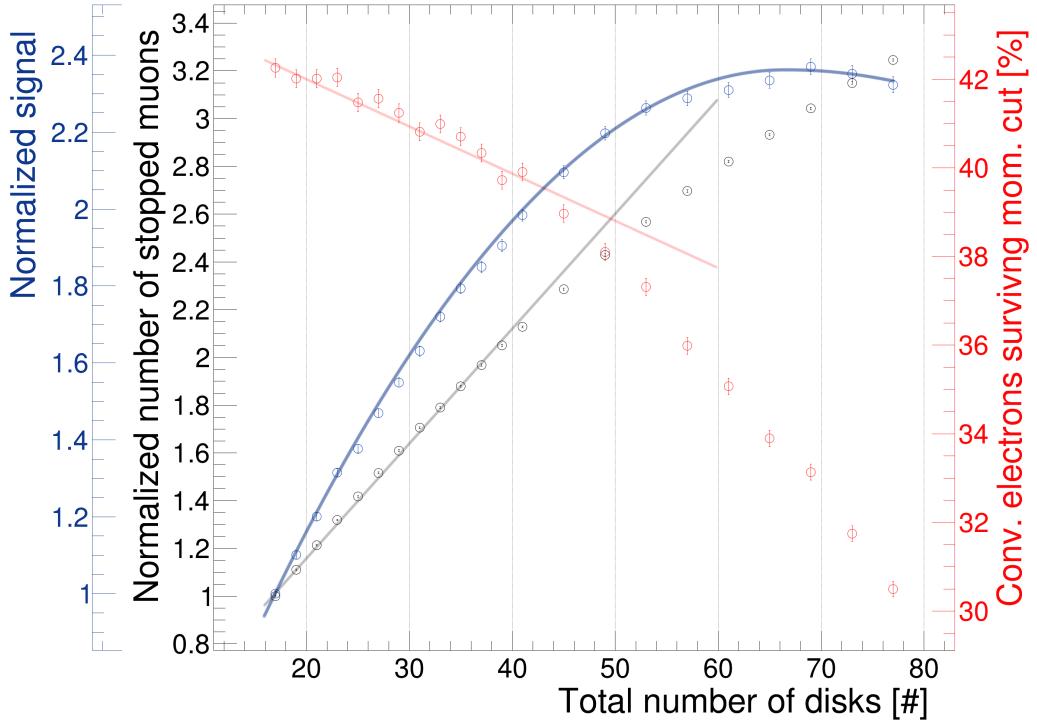


Figure 4.31.: Summary for number of stopped muons (black), momentum cut survival probability (red), and the resulting normalized signal (blue) for configurations with varying total number of target-disks. Thickness and radius was kept constant at 0.2 mm and 100 mm, respectively. Overlaid trendlines show a linear behavior up to 40, while the maximum normalized signal corresponds to approximately 60 target-disks.

4.7.3. The Effect on Detector Background

Not only the number of stopped muons and available signal but also detector backgrounds increase with additional material within the beam. The relationship between total CDC (layer) occupancy and the number of employed target-disks is shown in Figure 4.32. For each additional target-disk, the total CDC occupancy increases approximately by an absolute value of 1 %. CTH hit rate follows a similar trend as the total CDC occupancy and is included in the figure for reference. As discussed in Section 4.2.4, the maximum permissible total CDC occupancy for the employed MC truth analysis is limited to 50 %. The value is indicated in Figure 4.32b and corresponds to 33×0.2 mm target-disks or 6.6 mm of aluminium – merely half the optimum for the available signal.

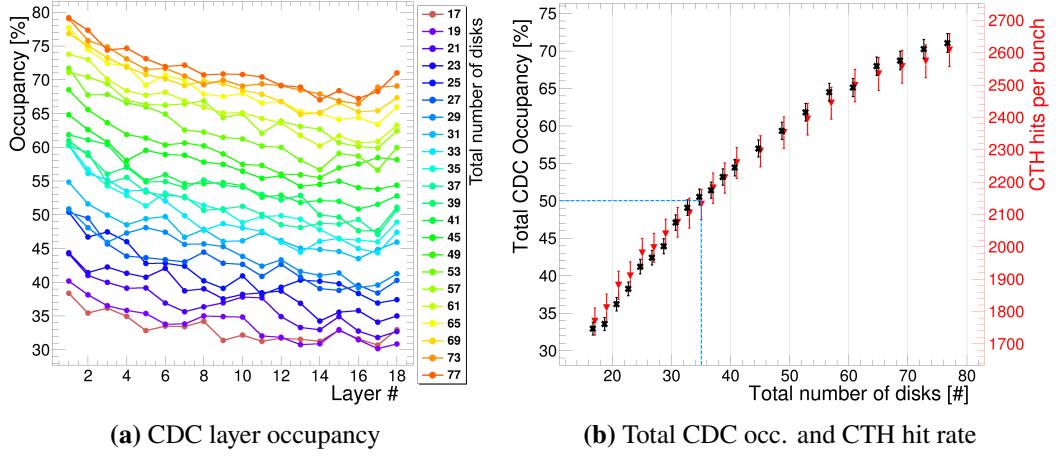


Figure 4.32.: CDC and CTH background summary for varying total number of target-disks. Target-disks were uniformly spaced within the 800 mm muon-stopping target space of CDC, and their radius was fixed at 100 mm. The data is averaged over 14 MC5 bunches. The 50 % maximum permissible total CDC occupancy, corresponding to 33×0.2 mm target-disks, is indicated by a blue dashed line.

A breakdown of CDC layer occupancy for the 33×0.2 mm target-disks configuration is shown in Figure 4.33. While in comparison to the default configuration (see Figure 4.12 in Section 4.3.1) the number of stopped muons almost doubles, the muon-associated background roughly increases by manageable 50 %. However, the *other beam* occupancy also doubles. Its details are quantified in Table 4.10. Background from neutrons remains constant within statistic uncertainties. The neutron interaction probably with the muon-stopping target is small, and most cause hits directly in CDC or after interactions with more massive surrounding beam line parts. Background from electrons and positrons also slightly increases, but their contribution to CDC occupancy remains small due to their limited initial relevance. By far the most crucial increase comes from pions. Because not-yet-decayed pions now traverse more target-disks, their probability of being stopped and captured in them increases. The number of background-inducing primary pion trajectories per bunch more than doubles. As most produce a multitude of other background particles (see discussion in Section 4.3.1), the associated CDC occupancy almost triples. Unfortunately, the arriving pion profile closely matches that of muons, preventing muon-target-specific adjustments for pion suppression. Consequently, the number of feasible target-disks is heavily limited by CDC occupancy.

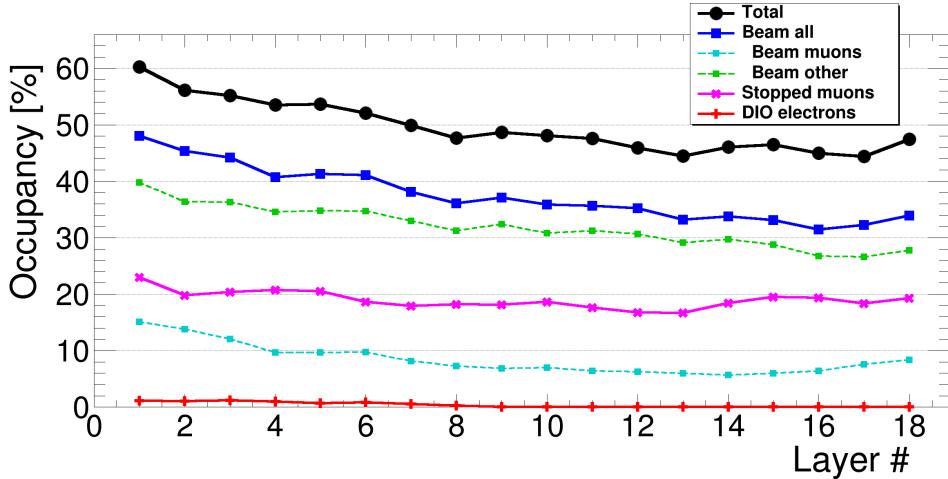


Figure 4.33.: CDC layer occupancy for a muon-stopping target configuration with 33×0.2 mm target-disks. The data is averaged over 14 MC5 bunches. Total CDC occupancy amounts to 49 %. Further details of the *other beam* component are given in Table 4.10.

Table 4.10.: Particle separated contribution to total CDC occupancy for a muon-stopping target configuration with 33 target-disks. Only primaries whose secondaries cause CDC hits are accounted for. The data is averaged over 14 MC5 bunches. While electron, positron, and neutron occupancy remain comparable to the default configuration (see Table 4.5), the number of primary pion trajectories causing background doubles and their contribution to total CDC occupancy almost triples.

Particle type	Number of primary traj. [# per bunch]	Total CDC occupancy [%]
Pion	80.79 ± 2.40	26.28 ± 0.20
Neutron	122.71 ± 2.96	4.70 ± 0.09
Electron	25.50 ± 1.35	1.92 ± 0.06
Positron	2.14 ± 0.39	0.09 ± 0.01

4.7.4. Summary and Performance Overview

Up to 60 target-disks, the gains from an increased number of stopped muons outweigh the signal losses resulting from a stronger spread in conversion electron momentum. However, the total CDC occupancy also rises by roughly 1 % for each additional target-disk. Exact numbers for configurations with up to 77 target-disks are summarized in the performance overview shown in Table 4.11.

The total CDC occupancy within the limitations of this MC truth-based analysis is set to a maximum of 50 %. In consequence, the number of target-disks must not exceed 33×0.2 mm, equivalent to 6.6 mm of aluminium. Such a configuration would increase the available signal by 74 %. The result is expected to improve further when combined with the previous finding that thinner target-disks result in better performance. Discussion and values with this optimization applied are presented in Section 4.10.

Table 4.11.: Performance overview of configurations with varying total numbers of target-disks. Target-disks thickness was kept constant at 0.2 mm. The data is averaged over 14 MC5 bunches at 16×10^6 POT per bunch. Values are calculated following the procedure described in Section 4.2.

Number of disks	$N_{\mu/\text{bunch}}^{\text{stopped}}$ [per bunch]	$A_{\mu \rightarrow e}^{\text{GM*,ROI}}$ [%]	normalized signal, K_{NS}	total CDC occupancy [%]	CTH rate [per bunch]
17	7838.68	17.79	1.00	32.07	1766.0
19	8697.53	17.89	1.12	33.56	1811.9
21	9506.79	17.85	1.22	36.21	1881.6
23	10336.20	17.97	1.33	38.24	1910.1
25	11106.04	17.52	1.40	41.21	1981.6
27	11878.06	17.48	1.49	42.40	1997.4
29	12609.23	17.36	1.57	43.94	2040.4
31	13357.55	17.25	1.65	47.09	2076.1
33	14020.63	17.33	1.74	49.08	2104.6
35	14719.05	17.22	1.82	50.51	2131.4
37	15412.29	16.95	1.87	51.41	2181.9
39	16062.77	16.75	1.93	53.18	2213.2
41	16666.21	16.82	2.01	54.44	2259.8
45	17906.86	16.53	2.12	57.00	2295.9
49	19025.36	16.31	2.23	59.34	2353.4
53	20109.96	15.90	2.29	61.83	2395.1
57	21123.13	15.30	2.32	64.49	2444.4
61	22093.55	14.77	2.34	65.12	2499.0
65	22973.34	14.36	2.37	67.98	2534.7
69	23833.87	14.05	2.40	68.73	2556.9
73	24672.68	13.47	2.38	70.27	2574.1
77	25426.72	12.91	2.35	71.07	2608.6

4.8. Disk Positioning Along the Beam Axis

With the physical characteristics and number of target-disks decided, this section now focuses on their positioning. Section 4.4.2 motivated the decision to exclude individual x - and y -positions. Thus, only their z -position along the beam axis is left as a free parameter. The total number of stopped muons and detector background are also assumed to stay constant because the amount of the material interacting with the beam is kept constant. Consequently, the sole parameter of interest will be the signal acceptance.

As a short reminder, to be accepted by the COMET trigger a conversion electron must perform at least one rotation through CDC before hitting the CTH (see Figure 2.3). Thus, conversion electrons originating from target-disk located very close to either US or DS CTH must have very low longitudinal momentum to fulfill this condition. Alternatively, they require a substantial longitudinal momentum component to travel all the way to the far-sided CTH counters, potentially traversing multiple follow-up target-disks along their path.

Figure 4.34 explores the consequences of this problem by visualizing the probabilities of conversion electrons originating from different z -positions to satisfy the CDC trigger (Figure 4.34a), produce a 4-fold coincidence in CTH (Figure 4.34b), or fulfill the combined COMET trigger condition (Figure 4.34c). Six different configurations are overlaid to verify that the results are independent of the total number of employed target-disks. The x -binsize for each configuration is adjusted accordingly, and 4th order polynomial fits are superimposed on the data as trendlines. The 103.6 MeV/c momentum window cut condition is excluded. Accounting for the ROI condition only reduces the overall histogram's height for configurations with a larger total number of target-disks, but the general shape and conclusion would remain unchanged.

Figure 4.34 confirms the initial assumption, namely: the COMET trigger acceptance is at its maximum for central target-disks within the plateau region between $-250 \text{ mm} < z < 250 \text{ mm}$. It falls off steeply when approaching the CTH counters located at approximately $\pm 410 \text{ mm}$. In numbers, the acceptance approaching the outmost target-disks drops from 22 % to below 10 %, equal to a reduction of approximately 40 %. Consequently, moving these target-disks closer to the center should be beneficial. The drawback of this shift is a tighter target-disk spacing, thus more hits on subsequent target-disks and increased signal loss due to the momentum window cut.

The optimal balance between z -offset and the spacing of target-disks is shown in Figure 4.35. 2nd order polynomial fits are superimposed on each dataset to aid visualization. Note that in comparison, the signal acceptance for the default configuration is higher due to less employed material. However, the effect is dominated by the (not shown) significantly reduced amount of stopped muons (see Section 4.7), which results in less available signal. Offsets at US and DS sides were varied separately to account for a higher number of stopped muons in the first (US side) encountered target-disks, which decreases linearly by approximately 1 percent for later ones. However, effects from this difference appear negligible within the scope of the statistical uncertainties, and switching US and DS offsets yields identical results.

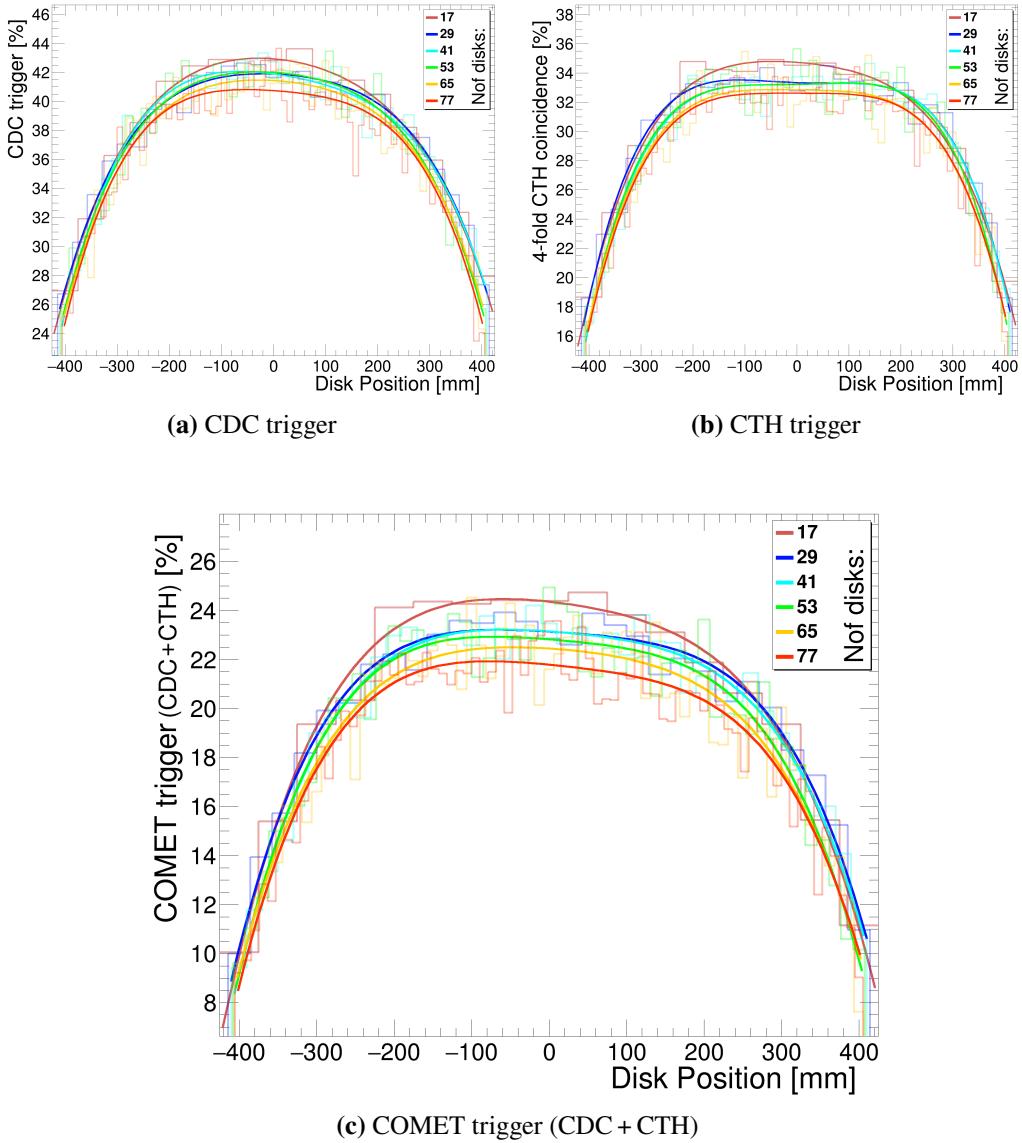


Figure 4.34.: Position dependent probability of conversion electrons to fulfill the CDC, CTH or both trigger conditions. The 103.6 MeV/c momentum window cut is excluded in all plots. The binsize of the x -axis is adjusted to correspond to exactly one target-disk. 4th order polynomial fits are overlaid as trendlines. The data is averaged over 14 MC5 bunches.

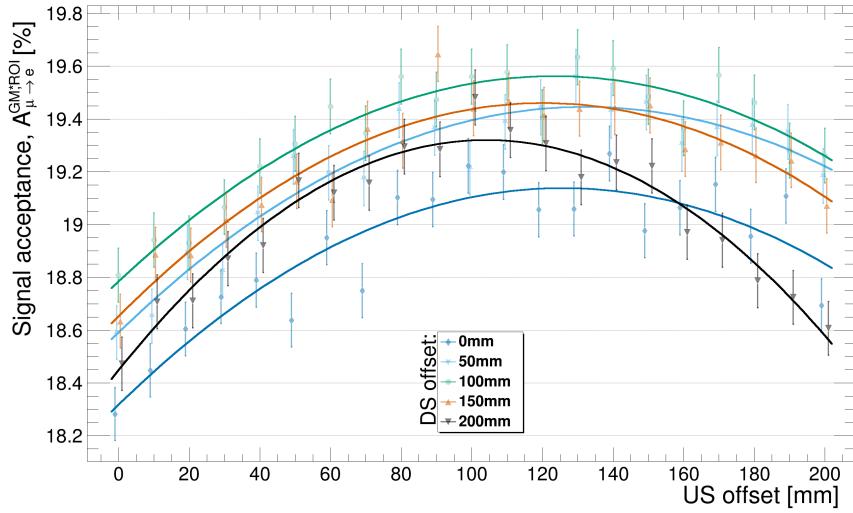
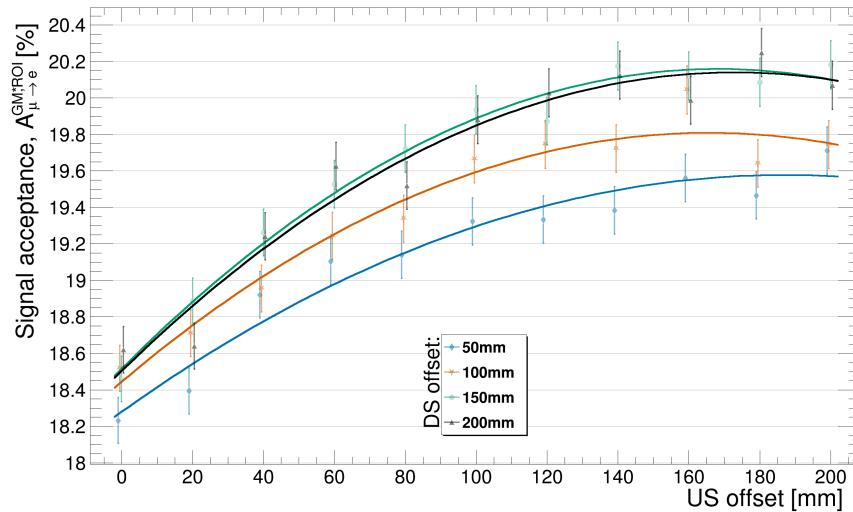

 (a) 60×0.1 mm target-disk configuration

 (b) 17×0.2 mm (default) target-disk configuration

Figure 4.35.: Signal acceptance for varying US / DS z -offsets. The data points of different DS offsets are slightly shifted to increase visibility but refer to their integer value. 2nd order polynomial fits are overlaid as trendlines. Shown error bars include only statistical uncertainties, and statistics correspond to 200 000 restarted conversion electrons. Due to more target-disk, hence tighter target-disk spacing, the optimum offset of the 60×0.1 mm configuration (a) is smaller than for the default configuration (b). Compared to no offsets, the signal acceptance improves by 1.2 and 2.4 percentage points, respectively.

For the configuration with 60×0.1 mm target-disks, the best performance corresponds to an 100 mm US and DS offset, improving the total signal acceptance by approximately 1.2 percentage points compared to zero offsets. For the default configuration, larger offsets are possible without reducing their spacing too heavily as less target-disk are used. The default configuration's optimum lies around a combined offset of 300 to 400 mm, equal to roughly 175 mm at the US and DS side. Compared to the zero-offset case, the signal acceptance improves by approximately 2.4 percentage points.

4.9. Signal Background from Muon Decay In Orbit

An introduction to the background from DIO was provided in Section 2.3 and previous estimates from the COMET Phase-I TDR [43] indicate a fraction of 0.01 DIO events for a SES of 3×10^{-15} and a momentum window cut of $103.6 \text{ MeV}/c < P_e < 106.0 \text{ MeV}/c$.

Incorporated updates on the CyDet geometry within ICEDUST and improved theoretical predictions for the DIO endpoint spectrum [16] make it worthwhile to reevaluate those values. The MC procedure for this study was described in Section 3.2. Notably, the energy of the DIO electrons is resampled according to the function provided in Formula 3.1, as taken from [16]. The first part of this section applies the improved spectrum to update DIO background predictions for the default muon-stopping target. The second part then estimates DIO background for a set of improved configurations. While previous DIO calculations focused primarily on key μ -e conversion elements such as aluminium, [16] includes parameters for a wide range of isotopes. Therefore, the final part of this section will use these predictions to investigate, for the first time, the severity of COMET's DIO background from impurities in the target-disks or the surrounding helium atmosphere.

4.9.1. Updated Results for the Default Configuration

Example DIO endpoint spectra for a selection of interesting isotopes are shown in Figure 4.36. The figure includes a comparison between results from [17] and the new DIO spectra obtained from [16], which improves the previous result by implementing leading order radiative corrections. Results are given as Taylor expansions around $E_{\mu e}$, the energy of the conversion electrons. A comparison to the ‘easier-to-use’ polynomial from [17] is also included. Differences between the old and new result for the aluminium nucleus (Al, $Z = 13$) amount to approximately 1 % at 103 MeV/c and a 2 keV shift in the DIO endpoint energy.

The estimation of the COMET DIO background, accounting for detector acceptance and resolution, is conducted with ICEDUST. The starting positions of DIO electrons are generated based on prerecorded muon-stopping positions. The energies of the DIO electrons are resampled from the corresponding energy spectra of each element. Care must be taken to select an appropriate energy range for resampling due to the spectra's steep decline at higher energies. Choosing an unsuitable large range would result in most of the generated

DIO electrons being outside of the ROI, rendering the statistics of relevant DIO electrons with momenta higher than 103.6 MeV/c insufficient. Therefore, the energy resampling range for the DIO spectra was chosen to start from 103.00 MeV up to 105.64 MeV. Within this range, a total of 400 000 DIO electrons were started, resulting in a reasonable statistic of over 1000 remaining DIO events within COMET's default ROI *after* applying COMET's trigger condition for all investigated elements. A second energy resampling range, starting from 100.00 MeV, was utilized for figures throughout this section. All shown results represent a combination of both ranges after their respective normalizations.

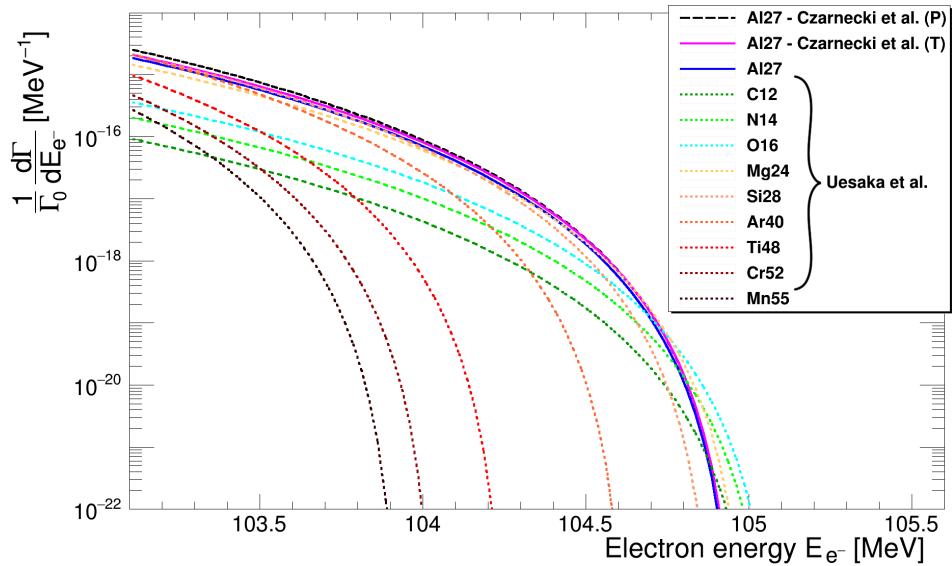


Figure 4.36.: Comparison of DIO endpoint spectra for various isotopes. Older results for aluminium from [17] are included as references. Both the easy-to-use polynomial (dashed black) and a Taylor expansion around the spectra's endpoint energy (solid magenta) from the paper are included for comparison. Newer results from [16] utilize a Taylor expansion based on Formula 3.1. Parameters were chosen conservatively to always yield the highest possible DIO background.

Signal conversion electrons were generated in a similar manner, utilizing a fixed starting energy of $E_{\mu e, \text{Al}} = 104.97$ MeV. However, since DIO and conversion results originate from separate MC campaigns, they require individual normalization before comparison. Both arise from muons stopped inside the target-disk, making the total number of stopped muons, here denoted as $N_{\mu\text{-stopped}}$, a suitable normalization factor.

Calculating $N_{\mu\text{-stopped}}$ for the normalization requires the assumption of an *a priori* branching ratio. Conventionally, $\mu\text{-}e$ conversion branching ratios are expressed with respect

to the standard-model muon capture rate² as defined in Formula 2.3. From this partial branching ratio $\mathcal{B.R.}$, the full conversion branching ratio (= probability for μ -e conversion), denoted as \mathcal{B} (μ -e conversion), is calculated with:

$$\mathcal{B} (\mu\text{-e conversion}) = \mathcal{B.R.} \times \mathcal{B} (\text{nuclear capture}) \quad (4.6)$$

and the normalization factor for the signal MC can be obtained as:

$$N_{\mu\text{-stopped}}^{\text{conv.}} = N_{\text{simulated}}^{\text{conv.}} / \mathcal{B} (\mu\text{-e conversion}). \quad (4.7)$$

Here, $N_{\text{simulated}}^{\text{conv.}} = 100000$ corresponds to the total number of simulated conversion electrons. The analysis assumed a partial branching ratio of $\mathcal{B.R.} = 3.0 \times 10^{-15}$, equal to the predicted SES of the COMET default configuration as detailed in the next Section 4.10. The total muon capture probability was taken from [55] as \mathcal{B} (nuclear capture) = 60.9 %.

To normalize the corresponding DIO spectra, it is essential to consider both the DIO branching ratio and the limited energy range used in the resampling procedure. The muon DIO branching ratio is calculated using the relation:

$$\mathcal{B} (\text{DIO}) = \frac{Q \Gamma_{\text{free-decay}}}{\Gamma_{\text{total}}} \quad (4.8)$$

with: $\Gamma_{\text{total}} = 1/\tau_{\mu^-}$
and: $\Gamma_{\text{free-decay}} = 1/\tau_{\mu^+}$.

Values for the free muon lifetime τ_{μ^+} , the lifetime of the bound muon τ_{μ^-} used to calculate the total decay rate Γ_{total} , and the Huff factor Q are taken from [14, 85]. The small correction factor Q , introduced in Section 1.1.2, accounts for a slight reduction of the free muon decay rate for bound muons, as binding of the muon slightly reduces the available energy for the decay [18]. For aluminium, the DIO branching ratio is found to be $\mathcal{B} (\text{DIO}) = 39.1$ %.

The reduced resampling range is incorporated by numerically integrating the given DIO energy spectra. In combination, the normalization factor for the DIO MC can be obtained with:

$$N_{\mu\text{-stopped}}^{\text{DIO}} = \frac{N_{\text{simulated}}^{\text{DIO}}}{P(103 \text{ MeV} \leq E_e^{\text{DIO}} \leq E_{\mu e})} / \mathcal{B} (\text{DIO}) \quad (4.9)$$

where:

$$P(103 \text{ MeV} \leq E_e^{\text{DIO}} \leq E_{\mu e}) = \left[\frac{1}{\Gamma_0} \int_{103 \text{ MeV}}^{E_{\mu e}} \frac{d\Gamma}{dE_e} dE_e \right] / \left[\frac{1}{\Gamma_0} \int_{0 \text{ MeV}}^{E_{\mu e}} \frac{d\Gamma}{dE_e} dE_e \right] \quad (4.10)$$

and: $\frac{1}{\Gamma_0} \int_{0 \text{ MeV}}^{E_{\mu e}} \frac{d\Gamma}{dE_e} dE_e = Q$.

²The normalization to the rate of muon capture is convenient, as the ratio cancels many details of the nuclear wavefunction in theoretical calculations.

Figure 4.37 presents the results of the procedure for the 17×0.2 mm target-disk default configuration. A combination of DIO and conversion signal spectra are overlaid as dashed black lines for four different branching ratio assumptions. Two important points should be noted. First, apart from the geometric acceptance and track-quality cuts, no additional acceptance factors (see Table 4.2) have been applied. However, assuming these factors affect both DIO and conversion signal alike, they could be included by a simple multiplication to the shown result. If the factors are applied to the plot, the area below the red signal line within COMET’s ROI would then equal one event over the full COMET runtime.

Second, the result utilizes COMET’s design CDC momentum resolution of 200 keV/c, as introduced in Section 4.2.3. This resolution will be utilized throughout this entire section, including for the discussion of the momentum ROI cuts. While previous works [50] applied similar simplified Gaussian smearing, the realistic detector response is expected to differ. Preliminary studies hint at a more asymmetric resolution function featuring a tail toward lower momentum, but these details are still under investigation and have not yet been applied. While both DIO and signal are affected by changes in the momentum resolution, the impact on the precision of DIO estimates should be reassessed once all details are finalized.

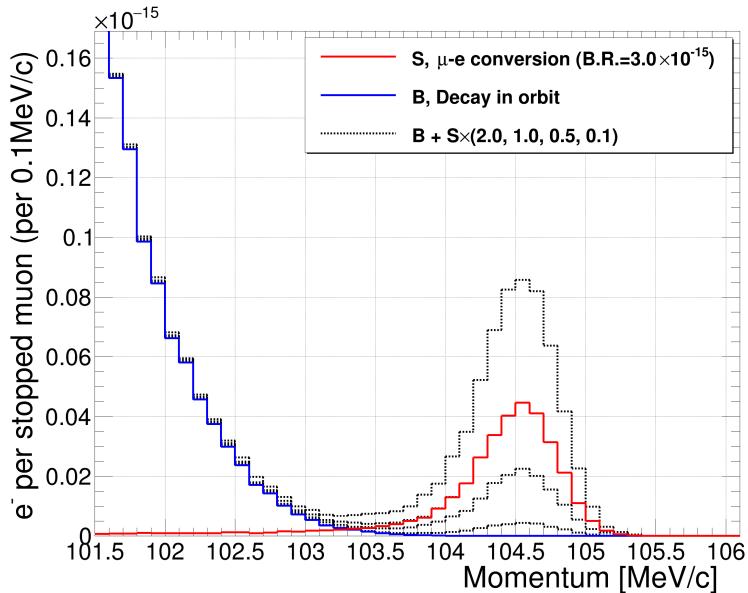


Figure 4.37.: The spectrum of DIO (blue) and μ -e conversion (red) electrons as seen by CDC after applying the COMET trigger condition. The CDC momentum resolution is modeled by convolving MC truth values with a Gaussian function with a standard deviation of 200 keV/c. A partial branching ratio of $\mathcal{B.R.} = 3.0 \times 10^{-15}$ is assumed. Combined electron distributions (black dashed lines) for the assumed conversion rate, double, half, and one-tenth of that rate are included.

Adjustments to the Momentum Region of Interest (ROI) Previous sections of this thesis utilized the default momentum ROI as defined in the COMET Phase-I TDR [43] with $P_e \in [103.6, 106.0] \text{ MeV}/c$. This range is primarily determined by the acceptable level of DIO background. Therefore, it is necessary to include a cross-check of the momentum ROI window in this section.

Figure 4.38 illustrates the significant impact of the lower bound of this ROI. Displayed is the remaining amount of DIO and conversion signal events in relation to the lower bound of the integration range. The integration was performed over the histograms presented in Figure 4.37. The upper bound of the ROI was held constant at 106.0 MeV/c. The graph's right axis demonstrates the momentum cut efficiency, calculated by forming the fraction of remaining to total number of events. The y-axes range is adjusted to match 100 % signal, corresponding to one event over COMET's full runtime for a partial branching ratio of $\mathcal{B.R.} = 3.0 \times 10^{-15}$. An indicator for a reasonable lower momentum cut is given as the signal divided by the square root of the signal and background³.

Several key points should be noted regarding Figure 4.38. First, in comparison to the COMET Phase-I TDR, the momentum signal acceptance for the default ROI range reduces from 0.93 down to 0.85. This reduction presumably arises from the improved electron resampling using prerecorded muon-stopping positions, which results in more electrons originating from the centers of target-disks rather than their edges.

Second, a similar reduction is observed for the expected DIO background. The more realistic DIO electron starting positions, in combination with the updated aluminium DIO energy spectrum predicting a slightly smaller DIO rate around $E_{\mu e, Al}$, leads to a decrease in the expected fraction of DIO events for an ROI lower bound of 103.6 MeV/c from 0.01 down to 0.003. Notably, the partial conversion branching ratio $\mathcal{B.R.}$ used for the normalization was chosen to match with the expected SES for the default configuration. Consequently, over the entire COMET Phase-I runtime, the expected number of stopped muons is such that the area under the signal curve above 103.6 MeV equals exactly one event. Hence, the 0.3 % DIO contamination equals an expectation of 0.003 DIO events.

Third, it is important to note the steep increase of DIO background and the quasi-plateau of the signal curve toward lower momenta. Although the sensitivity function $S/\sqrt{S + B}$ peaks around 103.4 MeV/c, changes in the lower bound of the ROI primarily impact the DIO rate. For instance, shifting the threshold down by 0.1 MeV/c improves the signal acceptance by merely 1 % but doubles the fraction of DIO events.

In summary, the previously selected ROI range of $P_e \in [103.6, 106.0] \text{ MeV}/c$ is shown to remain suitable for the default configuration. A small reduction of the ROI's lower bound by 0.2 MeV/c could be feasible. However, considering that the high-momentum DIO spectrum has yet to be experimentally confirmed and that this analysis is based on the simplified assumption of a Gaussian CDC momentum resolution, such changes were disregarded but could be revisited in future analysis.

³The sensitivity function $S/\sqrt{S + B}$ was not actively used to draw conclusions throughout this thesis but was included to facilitate comparison with a similar Phase-II study presented in [59].

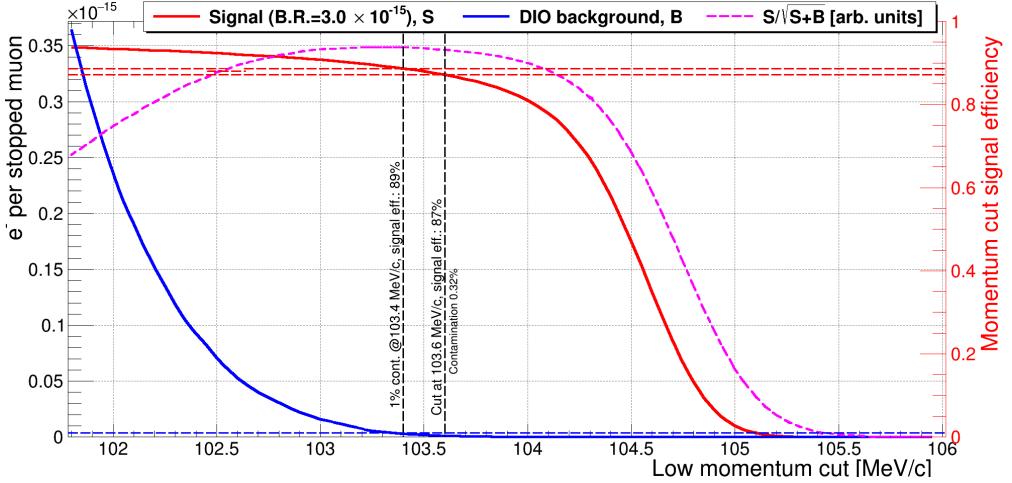


Figure 4.38.: Signal electron momentum cut acceptance and DIO background as a function of the ROI low-momentum threshold for the default muon-stopping target configuration. The upper bound of the ROI was held constant at 106.0 MeV/c and a partial branching ratio of $\mathcal{B.R.} = 3.0 \times 10^{-15}$ is assumed. Signal over the square root of signal plus background, shown as a magenta line, represents an indicator of a reasonable cut value.

4.9.2. Results for New Configurations

The expected DIO background and momentum window cut efficiency are significantly influenced by the chosen configuration of the muon-stopping target due to varied energy loss of electrons as they move toward the CDC. The following provides the momentum cut acceptance and expected DIO background for a subset of improved configurations. A comprehensive summary, including the final SES calculations for these configurations, is presented in Section 4.10.

Precise numerical values for momentum cut efficiency and DIO background are provided in Table 4.12. To guide understanding, details and differences between configurations are discussed using the three select examples shown in Figure 4.39. The simulation and analysis procedures for all results follow the same methodology as for the default configuration outlined in the previous section. A few things can be noted in regard to Figure 4.39:

- First, increasing the amount of the beam material, for example when increasing target-disk thickness from 33×0.1 mm (Figure 4.39a) to 33×0.2 mm (Figure 4.39b), results in more pronounced smearing of the high-momentum rising edge of the signal curve toward the plateau at lower momenta. Expectedly, this leads to reduced momentum cut efficiency for these configurations. This effect can be partly mitigated by utilizing thinner target-disks (Figure 4.39c).

- Second, the y-axes in all plots of Figure 4.39 are normalized to the number of muons stopped, and their range is adjusted to match 100 % momentum cut signal efficiency. Although configurations with less material yield more accepted conversion electrons per stopped muon (as seen at the intersection of dashed red lines with the vertical axis) and thus might seem to perform better, this effect is dominated by the overall increased number of stopped muons in configurations with a larger total number of target-disk.
- Third, while the geometric acceptance is noticeably enhanced by including a US and DS offset (compare Section 4.8), the impact on the momentum cut signal acceptance is minor and becomes noticeable only for high total target-disk configurations. However, due to the steep decline of the energy spectrum of DIO toward their endpoint energies, even a small increase of electron momentum losses from tighter target-disk spacing results in a noticeable decrease in the expected DIO rate within the ROI. Consequently, the proposed US and DS offset not only increases the available signal but also concurrently reduces DIO backgrounds.
- Lastly, the DIO contamination is expressed as fraction of the signal. For the default configuration, the area under the signal curve above 103.6 MeV corresponds to one event over the entire Phase-I runtime. However, assuming a constant runtime, the number of stopped muons is expected to nearly double for the 33×0.2 mm or 66×0.1 mm target-disk configurations. Consequently, for the selected partial conversion branching ratio $\mathcal{B.R.} = 3.0 \times 10^{-15}$, the signal area in these cases equals approximately two events. Likewise, the absolute DIO background is doubled. Nevertheless, the total count of DIO events over the entire Phase-I runtime remains below the previously predicted value of 0.01.

Table 4.12.: Momentum cut acceptance A^{ROI} and expected DIO background within the ROI for a subset of improved configurations. SES calculations for the configurations are included in Section 4.10. For discussion see text.

Configuration (Nof disks \times thickness)	A^{ROI} [%]	Fraction of DIO background in ROI [%]
Default, 17×0.2 mm	0.87	0.32
+175 mm US / DS offset	0.87	0.26
<hr/>	<hr/>	<hr/>
33×0.1 mm	0.91	0.38
+100 mm US / DS offset	0.91	0.32
<hr/>	<hr/>	<hr/>
33×0.2 mm	0.85	0.23
+100 mm US / DS offset	0.83	0.19
<hr/>	<hr/>	<hr/>
66×0.1 mm	0.88	0.26
+100 mm US / DS offset	0.86	0.14

4. The Muon-Stopping Target

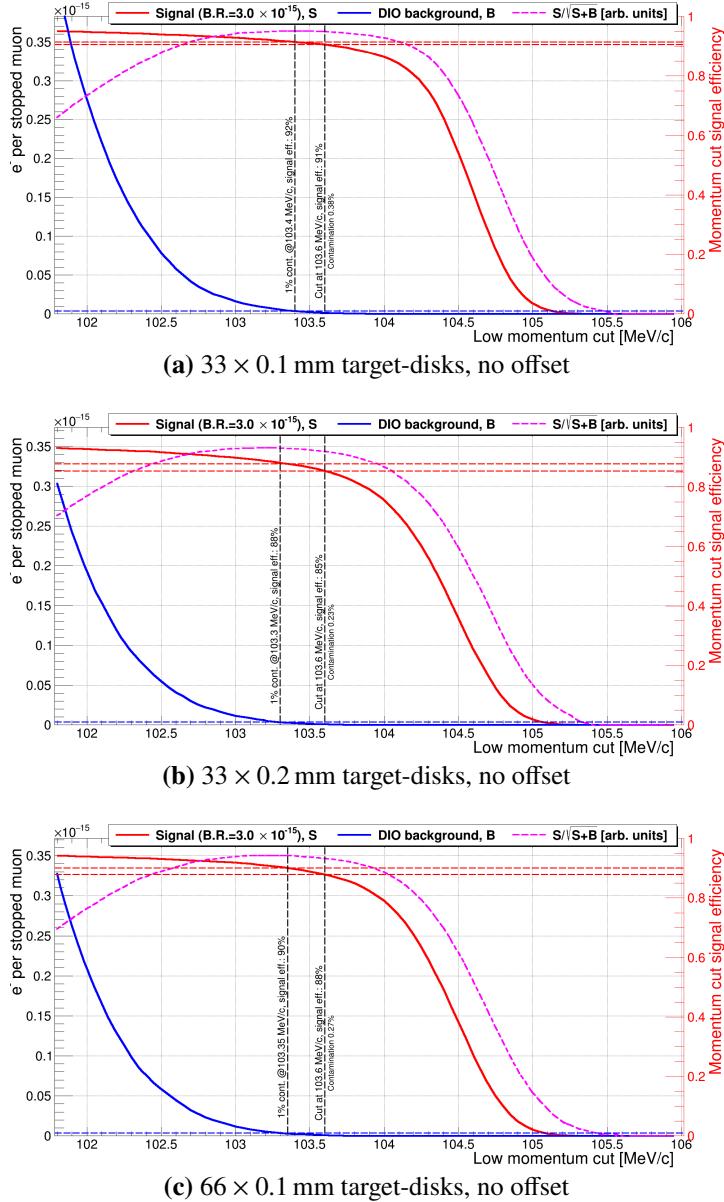


Figure 4.39.: Signal electron momentum cut acceptance and DIO background as a function of the ROI low-momentum threshold for a subset of improved configurations. The upper bound of the ROI is held constant at 106.0 MeV/c, and a partial branching ratio of $\mathcal{B.R.} = 3.0 \times 10^{-15}$ is assumed. The magenta line denotes the signal over the square root of signal plus background and serves as an indicator for a reasonable low-momentum cut value. Numerical values are summarized in Table 4.12.

4.9.3. DIO Background from Other Isotopes

The endpoint energy of DIO depends on the atomic number Z of the element (compare Figure 4.36 or Figure 1.1 in the Introduction). Particularly, for elements with Z values ranging from 5 to 12, prominently Nitrogen, Oxygen, or Carbon, the DIO endpoint energy exceeds $E_{\mu e, Al}$, thus making these elements potentially problematic sources of background. This contribution has not been taken into consideration thus far because precise theoretical calculations of DIO spectra around $E_{\mu e}$ have been absent. As a general approach, potentially problematic elements were avoided around the muon-stopping target. However, using new predictions from [16], this section and Section 4.9.4 will address inevitable contributions from impurities in the aluminium target-disks and the surrounding helium atmosphere, respectively.

The newly available results from [16] were illustrated in Figure 4.36. Although the DIO endpoint energies for elements between $Z = 5$ and $Z = 12$ exceed $E_{\mu e, Al}$, their distributions turn out to be much shallower compared to that of aluminium. Additionally, elements with atomic numbers slightly higher than aluminium, such as silicon (Si, $Z = 14$), were investigated. While their DIO endpoint energies remain below $E_{\mu e, Al}$, they still reach inside the experiment's ROI. Hence, due to the CDC's finite resolution, these elements could also contribute to the DIO background rate. Fortunately, the impact of this contribution is constrained by the steep decline in the DIO rate when approaching each element's respective endpoint energy, with a dependence approximately proportional to the fifth power of $(E_{\mu e} - E_e)$.

Results for a selected set of relevant elements are outlined in Table 4.13. Their total decay rate within the bounds of the ROI can be computed through numerical integration of the DIO spectra:

$$\Gamma_{\text{DIO}}^{\text{ROI}} = \frac{1}{\Gamma_0} \int_{103.6 \text{ MeV}}^{106.0 \text{ MeV}} \frac{d\Gamma}{dE_e} dE_e. \quad (4.11)$$

Due to an increasingly larger overlap of the muon and nucleus wavefunction, the nuclear capture branching ratio rises for elements with higher atomic numbers Z (see Section 2.3.3). Conversely, muonic atoms of lighter elements show comparably higher DIO rates. Values for the DIO branching ratios \mathcal{B} (DIO) are calculated following Formula 4.8, using parameters taken from [85]. The product of these two parameters yields the probability $P(E_e^{\text{DIO}} \in \text{ROI})$ of generating a DIO electron within the ROI. Assuming 1.57×10^{16} stopped muons over a runtime of approximately 146 days during COMET Phase-I yields the total number of DIO inside the ROI $N_{\text{DIO}}^{\text{ROI}}$.

Notably, $N_{\text{DIO}}^{\text{ROI}}$ refers to DIO electrons created at the muon-stopping target. The number is further reduced after accounting for acceptance factors. However, the detector acceptance, especially the efficiency of the momentum cut, changes depending on the element's exact DIO energy distribution. Consequently, final values for the total DIO backgrounds need to be derived from dedicated ICEDUST MC campaigns. As an example, five DIO spectra accounting for CDC resolution and the COMET trigger are shown in Figure 4.40. While the

rates for all DIO contributions get smaller, their order of importance before accounting for COMET trigger and ROI remains unchanged.

In summary, the intermediate estimate of $N_{\text{DIO}}^{\text{ROI}}$ is already sufficient to form a conclusions: the most severe DIO background arises from aluminium itself, rendering contributions from impurities inside the target-disks or muon stops in surrounding materials negligible. Note that values presented in Table 4.13 assume that all stopped muons during Phase-I runtime are within the respective element. Given that, for example, aluminium impurities are at most around the 1 % level, impurity-induced DIO rates are even further minuscule.

Table 4.13.: Probability for DIO, the total decay rate and probability to generate a DIO electron within the bounds of the ROI, and the resulting number of DIO background events emitted from the muon-stopping target for a variety of dangerous elements. Note that detector acceptance factors from Table 4.2 are not yet included, and the estimate assumed all 1.57×10^{16} stopped muons over the COMET Phase-I runtime of 146 days were within each respective element. Consequently, actual background numbers from impurities are further suppressed by a few orders of magnitude. See text for details.

Element	\mathcal{B} (DIO) [%]	$\Gamma_{\text{DIO}}^{\text{ROI}}$ $[\times 10^{-17}]$	$P(E_e^{\text{DIO}} \in \text{ROI})$ $[\times 10^{-15}]$ [%]	$N_{\text{DIO}}^{\text{ROI}}$ (over 146 d)
^{13}Al [16]	39.1	8.99	3.54	0.55
^{13}Al [17], (P)	39.1	11.51	4.53	0.71
^{13}Al [17], (T)	39.1	10.32	4.06	0.64
^{12}C	92.2	0.52	0.48	0.08
^{14}N	86.8	1.19	1.04	0.16
^{16}O	81.6	2.14	1.75	0.27
^{24}Mg	48.3	7.46	3.62	0.57
^{28}Si	34.1	8.94	3.07	0.48
^{40}Ar	24.4	5.15	1.27	0.20
^{48}Ti	14.7	0.65	0.10	0.02
^{52}Cr	11.4	0.08	0.01	0.00
^{55}Mn	10.3	0.02	0.00	0.00
^{56}Fe	9.1	0.00	0.00	0.00
^{63}Cu	7.1	0.00	0.00	0.00
^{64}Zn	7.0	0.00	0.00	0.00

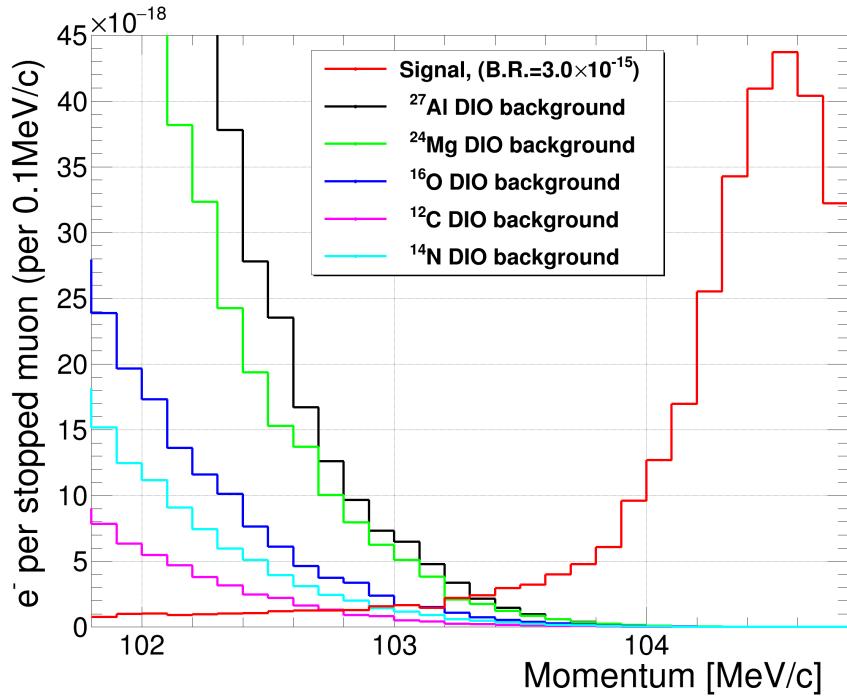


Figure 4.40.: Normalized DIO spectra for various elements as seen by CDC after applying the COMET trigger. The μ -e conversion signal (red line) assumes a partial branching ratio of $\mathcal{B.R.} = 3.0 \times 10^{-15}$. The CDC momentum resolution is modeled by convolving MC truth values with a Gaussian function with a standard deviation of 200 keV/c. Numbers are summarized in Table 4.13.

4.9.4. DIO Background from the Surrounding Helium Gas

Besides DIO background from target-disk impurities, muon stops inside the surrounding helium gas were feared to produce a non-negligible background rate. From [16], the helium DIO endpoint energy is taken as $E_{\mu e, He} = 104.15$ MeV/c and thus smaller than $E_{\mu e, Al}$. However, similar to heavier elements, helium DIO electrons still could contribute to the background due to the 103.6 MeV/c lower ROI bound and the limited CDC resolution. Furthermore, helium DIO electrons experience less momentum loss before reaching CDC because they do not have to first escape from the aluminium target-disks and can originate from larger radii.

The severity of this background primarily depends on the amount of muon stops in the gas. Numbers are obtained from ICEDUST and are pictured in Figure 4.41. The data is averaged over 60 MC5 bunches and yields about 0.58 muons stopped per millimeter per bunch. Over the entire the length of CyDet ($-1000 \text{ mm} \leq z \leq 1300 \text{ mm}$), this translates

to approximately 1200 stopped muons per bunch. However, as motivated in Figure 4.34c of Section 4.8, only z -positions larger (smaller) than -420 mm (420 mm) have a non-zero geometric acceptance. Accounting only for muon stops inside this critical region yields a total of 480 per bunch. The value could be further reduced to 350 muon stops per bunch if track-finding and -fitting algorithms limit the radius from where conversion electrons are accepted to $r \leq 110$ mm, but for this study such a cut-condition is not considered.

In conclusion, the number of muons stopped in the helium gas surrounding the muon-stopping target is at least 16 times smaller than the number of stops inside the target-disks. Additionally, following the same procedure as described in the previous section, one can calculate the probability of generating a DIO electron within the ROI. Despite 99.9 % of muonic helium atoms undergoing decay [85], the significantly reduced DIO endpoint energy of helium yields only $P(E_e^{\text{DIO}} \in \text{ROI}) \approx \mathcal{O}(10^{-22})$. Combining these results, it is evident that the DIO background originating from the surrounding helium gas can be safely disregarded.

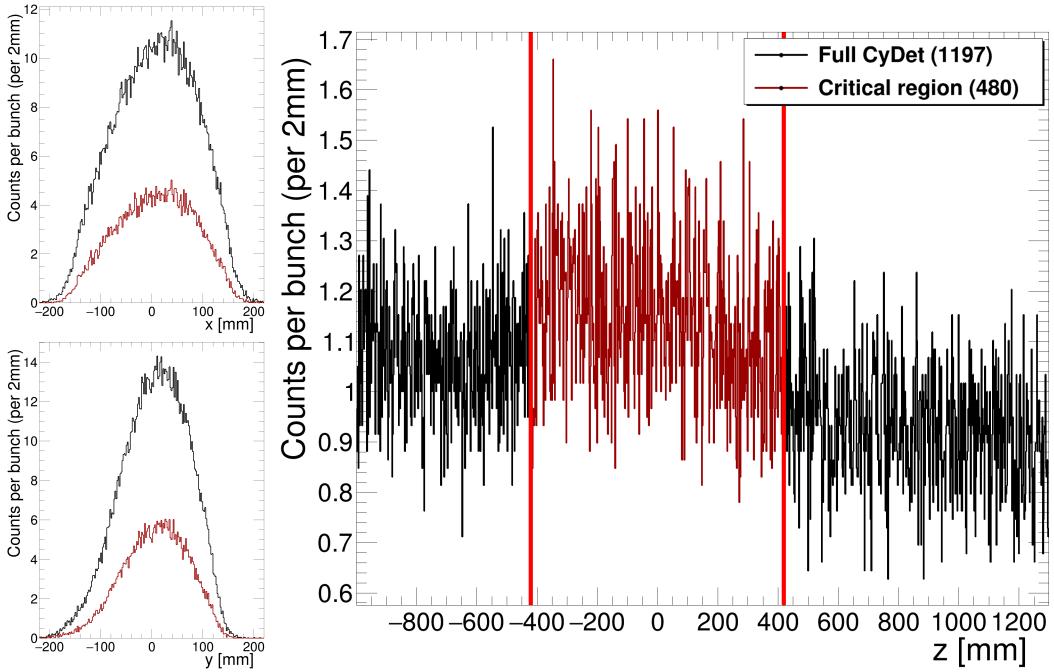


Figure 4.41.: Position of muon stops inside the helium gas surrounding the muon-stopping target. The critical region is limited to $-420 \text{ mm} \leq z \leq 420 \text{ mm}$, as geometric acceptance for larger or smaller z equals zero. The data was averaged over 60 MC5 bunches.

4.10. Single Event Sensitivity for New Configurations

The preceding sections explored various key parameters of the muon-stopping target individually. This section now aims to consolidate the gathered information and provide a concise overview. Given that the determination of an *optimal configuration* heavily depends on the maximum permissible CDC occupancy, it includes an overview of multiple options.

Up to this point, the performance of a configuration was evaluated based on its normalized signal K_{NS} as defined in Formula 4.3, along with its two necessary factors: the number of stopped muons per bunch $N_{\mu/\text{bunch}}^{\text{stopped}}$ and a subset of conversion electron acceptance factors $A_{\mu \rightarrow e}^{\text{GM,ROI}}$. To facilitate the comparisons with other publications, the first part of this section will demonstrate the calculation of the more universally applicable SES from these parameters. The second part will apply this procedure and provides a comparison and summary of proposed enhancements for the muon-stopping.

4.10.1. Calculating the SES

The SES as introduced in Section 2.2.2 is calculated as

$$\text{SES} = \frac{1}{(I_p/e) t_{\text{run}} R_{\mu/p} \cdot A_{\mu \rightarrow e} \cdot \mathcal{B}_{\text{capture}}}. \quad (4.12)$$

The equation expands the previously used number of stopped muons N_{μ} in Formula 2.4 to $(I_p/e) t_{\text{run}} R_{\mu/p}$. The beam current $I_p = 0.4 \mu\text{A}$, the elementary charge $e = 1.6 \times 10^{-19} \text{ C}$, and the run time $t_{\text{run}} = 146 \text{ d}$ yield the anticipated total number of POT for COMET Phase-I as $N_{\text{POT}} = 3.2 \times 10^{19}$. The constants are multiplied by the muon stopping rate per POT $R_{\mu/p}$, which directly depends on the muon-stopping target configuration and is derived from previously tabulated numbers of stopped muons per bunch using

$$R_{\mu/p} = \frac{N_{\mu/\text{bunch}}^{\text{stopped}}}{16 \times 10^6 \text{ POT/bunch}}. \quad (4.13)$$

As explained in Section 2.2.2, the coherence factor f_{coherent} is assumed to be one and thus omitted in Formula 4.12. The muon nuclear capture branching ratio $\mathcal{B}_{\text{capture}}$ is taken as 0.61 from [55].

The remaining parameter is the detector's signal acceptance $A_{\mu \rightarrow e}$. Calculation of $A_{\mu \rightarrow e}$, specifically its contribution from the geometric acceptance, involves certain assumptions. As detailed in Section 4.2.3, some track-quality cuts, namely NDF30 and χ^2 , could not be reassessed due to the study's reliance on MC truth values. They are assumed to remain comparable to results from the COMET Phase-I TDR, and potential correlations with other track-quality cuts are disregarded. Conservative estimates of 0.9 and 0.7 are used for single- and multi-turn tracks, respectively.

Table 4.14.: Updated geometric acceptance and breakdown of applied tracking quality cuts. Comparable values from the COMET Phase-I TDR are given in Table 4.3. The data is given for the default configuration with 17×0.2 mm target-disks. Unlike the older study, the presented results utilized a realistic muon-stopping distribution instead of emitting conversion electrons isotropically from the muon-stopping target. Values of NDF30 and χ^2 are conservatively estimated in conjunction with the previous results.

	single-turn tracks	multiple-turn tracks	combined
Geometric acceptance	0.26	0.08	0.34
incl. 4-fold coincidence	0.18	0.07	0.26

NL5 + NHIT + CL3	0.72	0.98	
$\text{NDF30} + \chi^2$ (<i>estimated</i>)	0.90	0.70	

Total	0.12	0.05	0.17

An updated summary of all track-quality cuts (compare to Table 4.3 of Section 4.2.3) is presented in Table 4.14. Unlike before, a realistic muon-stopping distribution is used as input for the conversion electrons' starting positions. This change causes a shift in acceptance from multiple- to single-turn tracks due to more starting positions located around the center of CyDet rather than the edges of the CTH. Furthermore, more starting positions are closer to the middle of target-disks instead of on their edges. Consequently, the NL5 cut efficiency, requiring conversion electrons to reach at least the 5th CDC layer, is reduced. In combination with the conservative estimates for NDF30 and χ^2 , the total geometric acceptance in this study decreases slightly to 0.17 from 0.18.

4.10.2. Summary and Interim Results

Table 4.15 presents results for the SES, including its contributing factors. The provided values for the detector's *total* signal acceptance incorporate the unchanged efficiency factors for *DAQ*, *online event selection*, *track finding*, and *timing window*, as taken from Table 4.2. Contributing values of *geometric acceptance* $A^{\text{Geom.}}$ and the *momentum window cut efficiency* A^{ROI} are listed separately, given their significant dependence on the muon-stopping target configuration. The calculation of the geometric acceptance follows the procedure described in the last section. The values for the momentum window cut efficiency factor are taken from Section 4.9.1, assuming a momentum window of $P_e \in [103.6, 106.0]$ MeV/c. $R_{\mu/p}$ is derived according to Formula 4.13, and calculation of numbers of stopped muons per bunch $N_{\mu/\text{bunch}}^{\text{stopped}}$ and the total CDC occupancy follows the established procedure described throughout Sections 4.2.2 and 4.2.4.

Table 4.15.: Updated results for the COMET Phase-I SES of various configurations for a total number of 3.2×10^{19} POT over a runtime of 146 days.

Configuration (N of disks \times thickness)	$R_{\mu/p}$ [$\times 10^{-4}$]	$A_{\mu \rightarrow e}$			SES [$\times 10^{-15}$]	total CDC occ. [%]
		A_{ROI}	$A_{Geom.}$	Total		
Default, 17×0.2 mm	4.89	0.87	0.17	0.035	2.99	32.07
+175 mm US / DS offset	4.83	0.87	0.19	0.041	2.59	31.08
<hr/>	<hr/>	<hr/>	<hr/>	<hr/>	<hr/>	<hr/>
33×0.1 mm	4.76	0.91	0.17	0.038	2.83	32.80
+100 mm US / DS offset	4.73	0.91	0.19	0.042	2.58	32.76
<hr/>	<hr/>	<hr/>	<hr/>	<hr/>	<hr/>	<hr/>
33×0.2 mm	8.79	0.85	0.17	0.035	1.67	48.75
+100 mm US / DS offset	8.61	0.83	0.19	0.037	1.61	47.87
<hr/>	<hr/>	<hr/>	<hr/>	<hr/>	<hr/>	<hr/>
66×0.1 mm	8.73	0.88	0.17	0.036	1.63	49.51
+100 mm US / DS offset	8.60	0.86	0.19	0.039	1.53	49.71

Expectedly, the increase in the total number of disks has the biggest impact on the SES. As motivated in Section 4.7, the maximum allowed total CDC occupancy of 50 % is reached at approximately 6.6 mm of aluminium, allowing an increase of up to 33×0.2 mm or 66×0.1 mm target-disks. Doubling the amount of material increases the number of muons stopped per POT by a factor of 1.8. Losses in momentum window cut efficiency due to the increased amount of material can be counteracted by employing thinner (but more) target-disks. Furthermore, independent of the exact configuration, a slight improvement of the geometric acceptance can be achieved by implementing small offsets at both ends of the muon-stopping target.

The proposed best configuration of this work is 66×0.1 mm target-disks with 100 mm US and DS offsets. For this configuration, the SES can be improved by 50 %. Alternatively, the runtime to achieve COMET’s design goal of 3×10^{-15} could be halved.

4.11. Production of Target-Disks

To test the stability of thin aluminium disks and later the installation procedure inside their holding structure (see Section 5.1), a set of prototype target-disks was produced. ^{27}Al is a stable isotope, it possesses a natural abundance of 100 %, and it is not poisonous, thus the material is affordable and easy to work with. Because aluminium is commonly used for machinery, a multitude of options in regard to alloys, mechanical stability, purity, etc. exists. Although Section 4.9.3 demonstrated that DIO background due to material impurities can be neglected, a high purity still guarantees a maximum number of muonic *aluminium* atoms, hence is desirable.

Pure wrought aluminium is commercially available with up to 99.85 % purity (DIN: EN AW-1085). Theoretically, purities of up to 99.999 % are possible but are not feasible for multiple thin 0.2 mm sheets due to their low mechanical stability and exceedingly high cost. In general, including traces of iron and silicon is common to increase the aluminium's strength and hardness. To further complicate things, strength and hardness are significantly defined by the material's temper. It indicates the general sequence of operations during production to achieve the desired properties. For this thesis, only strain-hardened aluminium is used as it provides the best material stability. This temper designation is described by the letter H, and is followed by one or two digits indicating the level of hardening. Full-hard aluminium of this temper is described by H18.

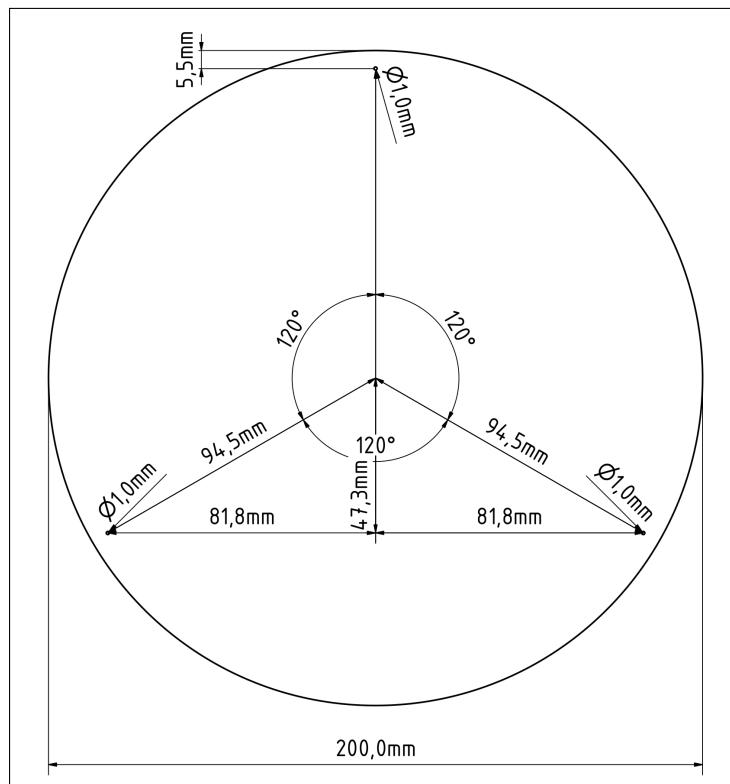


Figure 4.42.: Technical drawing of the aluminium target-disk. It includes three 1 mm diameter holes spaced 120° apart and positioned at a distance of 5.5 mm from the edge, used to suspend the target-disks within the holding structure (see Section 5.1).

The aluminium was ordered as metal sheets and precision laser-cut into its disk shape. A technical drawing of the target-disks is provided in Figure 4.42. The design includes three 1 mm diameter holes spaced 120° apart and positioned at a distance of 5.5 mm from the edge,

Table 4.16.: Maximum impurities of aluminium according to the DIN standard “EN AW-1050A” and “EN AW-1200”. The specified mass percentage applies to impurities as a maximum or, in the case of aluminium, as a minimum threshold. Other impurities refer to elements that must each have an individual mass percentage below the specified limit.

Chemical element	Atomic number	Mass percentage [%]	
		EN AW-1050A	EN AW-1200
Magnesium (Mg)	12	0,05	0,3
Silicon (Si)	14	0,25	Si+Fe: 1,0
Iron (Fe)	26	0,4	
Titanium (Ti)	22	0,05	–
Chromium (Cr)	24	–	0,1
Manganese (Mn)	25	0,05	0,3
Copper (Cu)	29	0,05	0,1
Zinc (Zn)	30	0,07	0,1
Other impurities	–	< 0,05	< 0,03
Aluminium (Al)	13	99,5	99,0

which will later be used to securely hold the target-disks in place. The company provides the uncertainty of the cutting procedure as less than 0.05 %. The thickness of the ordered metal sheets was confirmed using an outside micrometer at ten different positions. Total uncertainties on thickness also amount to less than 0.05 %. As a result, both are sufficiently small to have no impact on performance.

Four different aluminium alloys were tested. Unfortunately, many desirable purity (such as “EN AW-1085”) and temper combinations were not available as 0.2 mm (0.1 mm) thick sheets in sizes larger than 200 mm × 200 mm, as is required to produce the target-disks. Therefore, an initial test utilized 0.2 mm thick “EN AW-1050A H11” (99.5 % aluminium, 1/8-hard) and “EN AW-1200 H18” (99.0 % aluminium, full-hard) [77]. The exact alloy compositions including their maximum range of impurities are presented in Table 4.16. The H11 temper of “EN AW-1050A” proved much too malleable, and visible inelastic deformations and bumps occurred during the machining of the target-disk shape (see Figure 4.43a). The much more robust “EN AW-1200 H18” illustrated in Figure 4.43b showed no signs of deformations while still retaining a reasonable amount of elasticity to avoid cracks or breaking when pressure was applied to its surface or edges.

Thus, in the following, 0.2 mm thick “EN AW-1050A” was ordered in a comparable hard temper designation H19 (H19 exceeds H18 by at least 10 MPa). Despite its higher purity, the material performed comparably to “EN AW-1200 H18” and allowed for the machining of target-disks without permanent damage or deformations (see Figure 4.43c).

To incorporate the results of this thesis, the same material was tested with a 0.1 mm thickness. The thinner “EN AW-1050A H19” target-disk is shown in Figure 4.43d. This thickness and temper combination still retains enough of its stability and flexibility to be used as potential target-disk material for COMET. Although the material could be reasonably bent without breaking or permanent deformations, small irregularities around the edges of the disk from the machining were visible and are highlighted as an inset inside the figure. However, despite their distinct visibility, the irregularities amounted to no measurable change in thickness or shape and thus are assumed to be negligible.

In summary, 99.5 % pure aluminium with a reasonably high hardening (≥ 18) is suitable to be used for COMET’s muon-stopping target. Thirty target-disks were manufactured from 0.2 mm thick “EN AW-1050A H19” aluminium sheets to further test the installation procedure within their holding structure (see Section 5.1). Moreover, the aluminium temper was tested with a reduced 0.1 mm thickness as thinner target-disks are expected to improve the muon-stopping target’s performance (compare Section 4.28). Although the thin laser-cut target-disk showed small irregularities at its edges, they amounted to no measurable difference in thickness or shape. Consequently, such 0.1 mm thick target-disks are also deemed suitable for the experiment. This allows for the implementation of previously suggested enhancements, such as a muon-stopping target consisting of 66×0.1 mm target-disks, which improves COMET’s SES by 50 %.

4. The Muon-Stopping Target

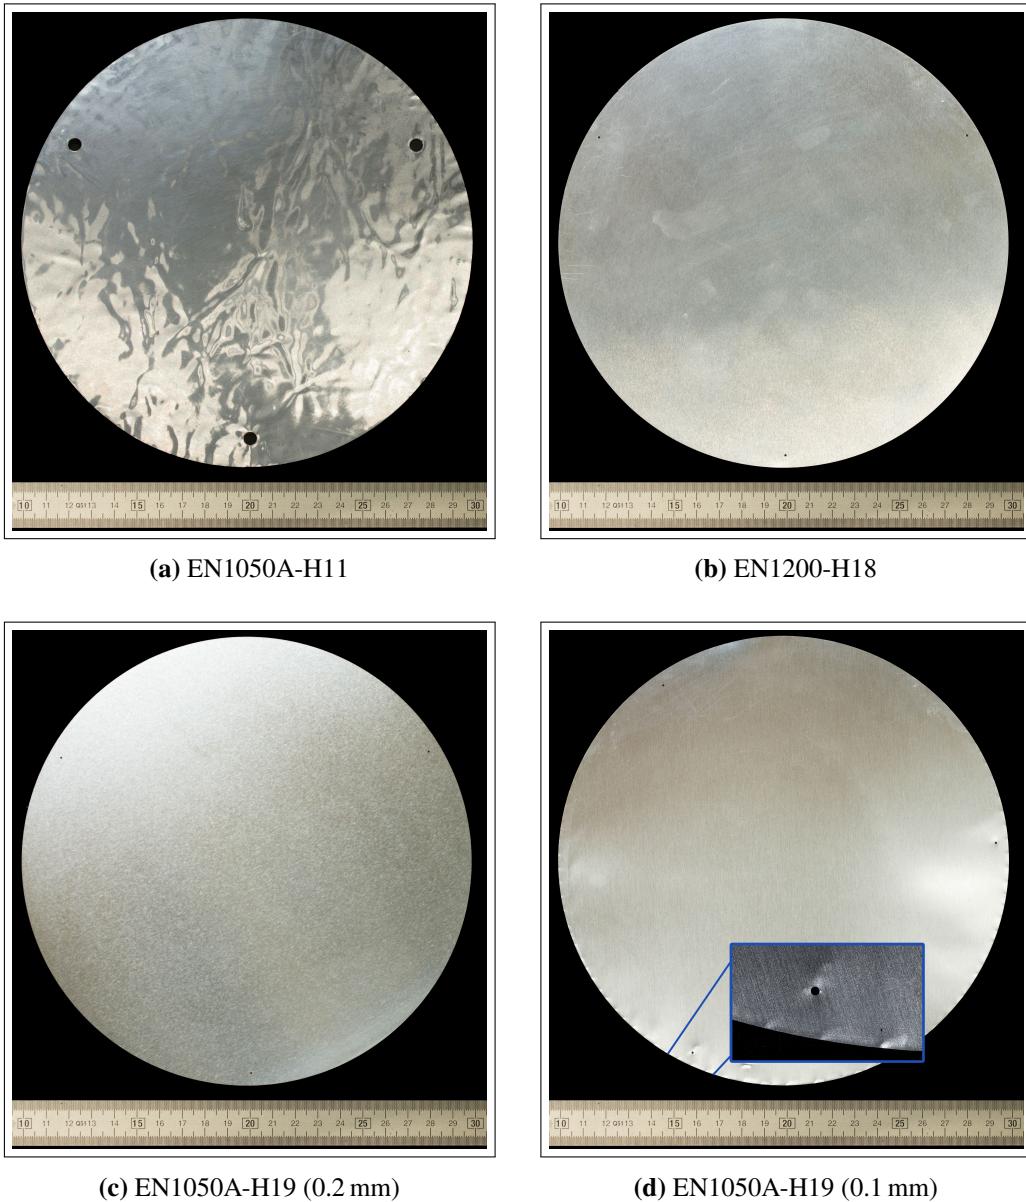


Figure 4.43.: Comparison of four aluminium target-disks made from material with different purity and temper. The disks were laser cut according to the schema shown in Figure 4.42. The (in)purities of the material are presented in Table 4.16. Photographs are retouched to enhance clarity and contrast. The 0.1 mm thick “EN1050A-H19” target-disk (99.5 % purity) in d) showed small irregularities after machining, highlighted inside the inset. However, these amounted to no measurable changes in thickness or shape.

5. The Muon-Stopping Target Suspension

A dedicated structure is necessary to precisely position the target-disks at the center of CyDet. However, the central location of this structure within the detector system imposes numerous requirements. Most importantly, interference of the suspension with the beam and its impact on signal acceptance or general detector background must be minimized. Furthermore, the design must guarantee stable operation throughout COMET’s runtime, and its installation should be easy in the event of exchanging muon-stopping targets between calibration and physics runs.

Experience with such a design was gained during the preceding work presented in [77], and the resulting mock-up target structure was shown in Figure 4.1. However, the focus of the former study was on validating MC results and conducting initial tests of the muon-stopping monitor (see Section 6). The target-disks were positioned in the middle of a square-shaped aluminium frame and held in place using nylon strings – neither of which is suitable for the actual COMET experiment. Based on these preliminary findings, this thesis further refined the design in close collaboration with the local TU Dresden workshop.

In addition to providing a general overview of the suspension and its subsequent production, this section will motivate underlying development decisions and numerically assess their impact on experimental parameters such as background and acceptance.

5.1. Technical Layout of the Suspension Structure

After passing through the bent 90° transport solenoid, charged particle profiles are Gaussian shaped, with their maxima around the center of CyDet (compare profiles shown in Figure 4.13 of Section 4.4.1). Thus, all suspension components must be positioned at as large radii as possible to minimize interaction with beam particles. Additionally, the amount of material between the target-disks and the CDC must be minimized, as conversion electrons hitting the structure will lose too much energy to be accepted by the COMET trigger.

Technical drawings of the resulting design and its implementation inside ICEDUST are shown in Figure 5.1 and Figure 5.2, respectively. Photographs of the fully assembled suspension are presented later in Section 5.5. For this discussion, the suspension is separated into three components: first, the larger ring structure at its beginning, providing support and general stability; second, the three long bars attached to the former, spanning from one CTH

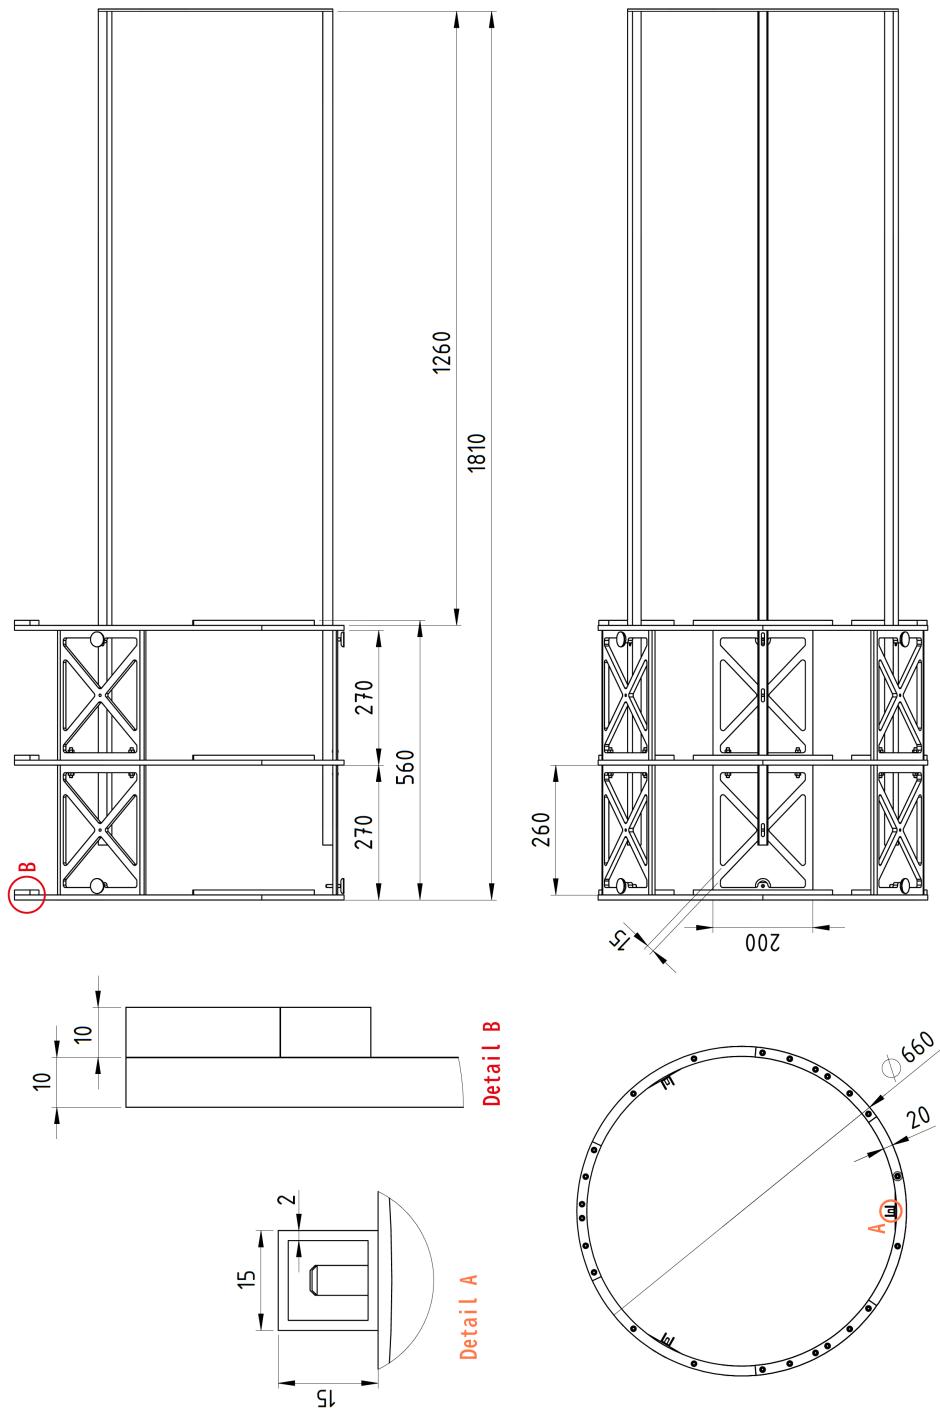


Figure 5.1.: Technical drawing of the suspension structure. Measures are given in mm. For details see text.

side to the other and used to hold the target-disks in place (see Section 5.2); and third, the thin ring attached to the end of those three bars, necessary to stabilize them against each other and to precisely install the structure.

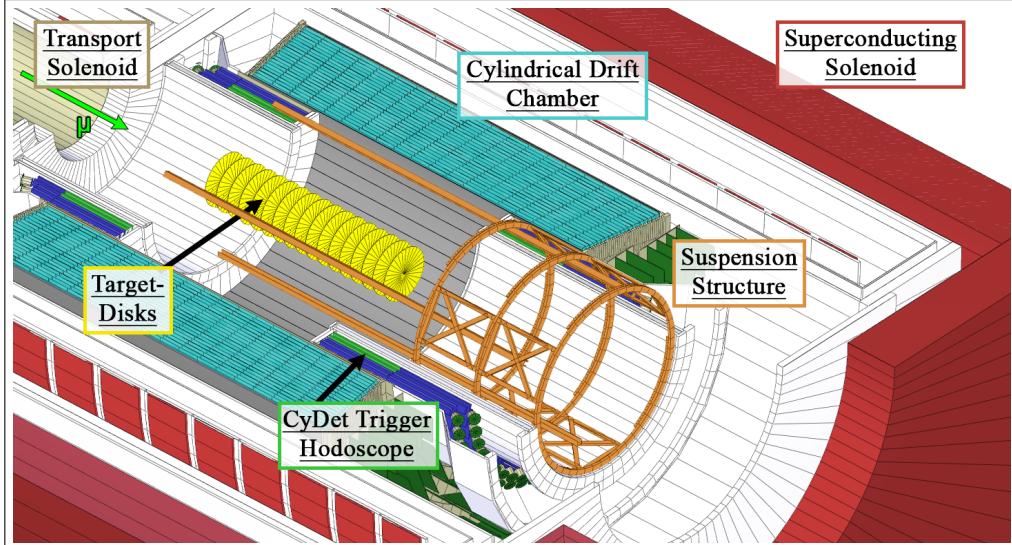


Figure 5.2.: Implementation of the muon-stopping target-disk suspension in ICEDUST. The beam arrives from the US transport solenoid in the top-left corner. The larger ring structure is positioned inside the CTH shielding but behind the scintillation counters (light green squares). The wedges to fine-adjust its positioning at the US CTH side, including the support ring stabilizing the three long bars, are not yet implemented in this figure.

5.1.1. The Outer Ring Structure

A photograph of the large ring structure is shown in Figure 5.3. Its general shape and maximum radius are essentially predefined by the CTH. The outer radius of the rings is adopted as 330 mm, leaving a small 7 mm margin to the inner wall of the CTH for adjustment or changes to the CTH's shielding, which is still under active development. The inner radius of the rings is set to 310 mm, and their total material width is 10 mm. Each ring is segmented into three 120° parts and secured by slightly shorter separate arcs with otherwise identical measures, facilitating quick and straightforward (dis)assembly.

Six plates, each 260 mm long, 200 mm wide, and 10 mm thick, space apart and support the rings. Cut-outs are utilized to minimize material and weight. Each spacing plate provides a height-adjustable mounting foot to anchor the entire suspension structure inside the CTH (see Section 5.3). The CTH inner shielding contains a thick 25 mm lead layer, causing it to

weigh multiple tons. However, the CTH can only be fastened to the outer ends of CyDet because of its position inside the CDC tube (compare Figure 3.3b of Section 3.3). Thus, the sides pointing toward the middle of the CDC of both CTH tubes are anticipated to sag downwards by approximately $1 \sim 2$ mm. In consequence, adjustable feet are preferable over a more direct connection like rails to counteract this sagging and guarantee a level suspension.

All parts of the ring structure are constructed from aluminium “EN AW-7075”, known for having the highest strength values among heat-treatable aluminium alloys. The material is paramagnetic to avoid interference with the magnetic field and offers sufficient mechanical strength with a better strength-to-weight ratio than alloys like stainless steel [86]. As shown in Figure 5.2, the rather large structure is completely positioned behind the CTH scintillators to avoid interference with the detectors.

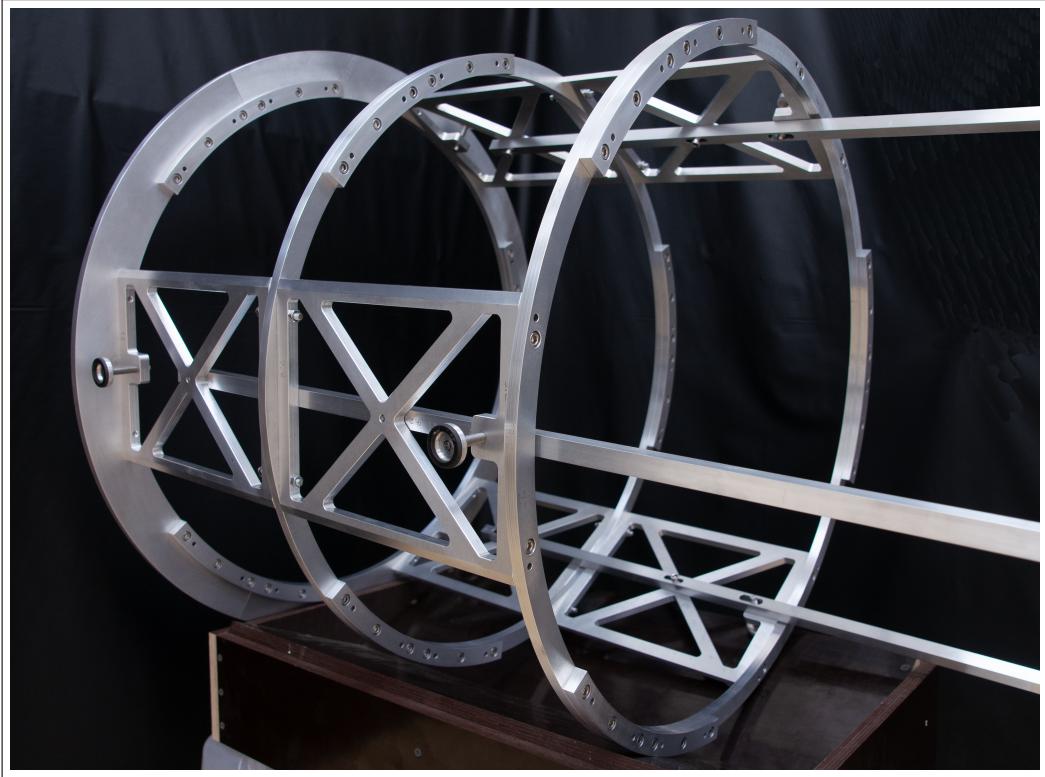


Figure 5.3.: Close-up photograph of the large ring component of the muon-stopping target suspension. In comparison to the technical drawing, the outmost ring segment has a larger radius of 365 mm. This part will remain outside the CTH and the suspension structure will be further stabilized by clamping this larger ring to the CTH endplate (see Section 5.3).

5.1.2. The Three Long Bars

Three 1700 mm long metal bars bridge the gap between US and DS CTH. Each bar is connected to two of the previously discussed spacing plates that connect the rings, as seen in Figure 5.3. A larger cut-out around the connecting screws leaves room for fine-adjustments of approximately ± 20 mm in the longitudinal (beam) direction.

Three different bar profiles were considered: round, u-shaped, and squared. Round profiles, with a diameter equal to the width or height of a square profile, require approximately 20 % less material for the same length and are the most robust against compression. However, compared to correctly aligned square profiles, they have less material in the direction of bending (vertical), making them less resistant to such a bend. Furthermore, precisely attaching round profiles to other components is quite challenging, leading to their disqualification.

If aligned correctly, u-shaped profiles offer the same bending resistance as complete squares with an improved weight-to-length ratio. However, as shown in Figure 5.1, each bar is rotated 120° against another, making it impossible within the intended design to align all bars with their hollow side facing straight up or downwards. In consequence, this caused strange deformations to the left or right side due to their weight, rendering u-profiles unsuitable.

Therefore, the chosen profile is hollow squares. Stainless steel, aluminium, and carbon fibre-reinforced plastic bars with varying wall thickness (2 mm to 6 mm) and width (10 mm to 20 mm) combinations were tested regarding their total bend. However, no significant differences were observed among these options, as the increased weight counteracted a potential increase in bending robustness from additional material. Effects on signal acceptance from these bars are discussed in Section 5.4.2. Despite smaller differences, paramagnetic aluminium bars of the “EN AW-6060” alloy, measuring 15 mm \times 15 mm in width and height, were chosen as a balanced solution considering size, stability, and machinability. Their wall thickness was adopted as 4 mm. Measurements performed with a laser level show a total downward bend at the far end of such bars of 13 ± 1 mm (15 ± 1 mm) without (with) the 17 target-disks installed. This bend is small enough not to complicate the installation of the suspension and can be corrected with the wedges at the US CTH side (see Section 5.3).

Note that each bar is made of a solid piece of metal with fixed drilling positions to hold the target-disks. Introducing a flexible solution to adjust the number or vertical position of target-disks would at least require additional drill holes, thereby resulting in a notably more fragile system. Consequently, for each modified muon-stopping target configuration incorporating offsets (see Section 4.8) and/or more but thinner target-disks (see Section 4.28), a new set of such bars has to be manufactured.

5.1.3. The Support Ring

The far ends of the bars are connected by a 4 mm thin “EN AW-5754” aluminium ring. During development, this ring proved essential to provide stability between the bars, especially when tension was applied during and after target-disks installation. Furthermore, the ring is used in combination with wedges inside the US CTH shielding to provide stability in x - and y -direction of the suspension’s far end side and to counteract the sagging of the bars, as discussed in Section 5.3. The added weight of this ring causes a negligible additional sagging of the bars by approximately 1 mm. More lightweight or thinner materials could be chosen, but due to this reasonably small impact, no further fine-adjustments were performed.

5.2. Holding the Target-Disks in Place: The Spokes

To secure the target-disks at the CyDet center, they are connected to the three long bars by tungsten wires referred to as “spokes”. One spoke at the top counteracts the target-disks’ weight, while one on the left and right side adjusts their x - and y -positioning. To reduce the probability of muons stopping inside the spokes, it is essential to minimize their overall size. A spoke diameter of 0.1 mm was adopted for the test illustrated in pictures shown throughout this section. The spoke material is selected to avoid signal background from DIO. As discussed in Section 4.9.3, high mass number materials are best suited for this purpose, so tungsten was adopted.

The connection of spokes to the target-disk, as shown in the bottom right of Figure 5.4, utilizes the 1 mm diameter laser-cut holes introduced in Figure 4.42 of Section 4.11. The target-disk is secured by leading the tungsten wire through the hole, forming a loop, and then clamping the short and long parts of the wire using thin aluminium metal sleeves. After clamping, the sleeves used in the test had an average thickness of 0.5 mm and at most a profile size of $5 \text{ mm} \times 3 \text{ mm}$ – small enough to not negatively affect beam quality¹. The use of wide tungsten loops prevents the formation of sharp edges in the wire, which could render the material vulnerable to breaks at these points. Welded connections were considered but decided against due to the delicate handling of the fragile wire and the thin target-disk. Additionally, concerns regarding radiation hardness and the weld’s attachment to the target-disk sides introducing a small but hard-to-correct tilt of the target-disk disfavoured these connections.

The bottom left picture in Figure 5.4 shows the connection of the spokes to the metal bars. This connection is engineered to prevent the formation of fracture-prone kinks in the wire while permitting the precise adjustments of tensile strength and target-disk positioning without causing twisting of the wire. First, a 4 mm diameter and 33 mm long hollow screw

¹The connection was also tested with smaller (e.g., half-size) sleeves and shows comparable and sufficient stability. However, at the time, such sleeves were not commercially available but might be specifically machined if deemed necessary.

5. The Muon-Stopping Target Suspension

is inserted through the bars and secured with a nut. Subsequently, the tungsten wire is threaded through the screw and, after adjusting the wire's approximate length, secured at its end by clamping the separated top halves of the hollow screw with a bolt. Lastly, the exact length and tension of the wire can be fine-tuned within a range of ± 5 mm by adjusting the protruding end of the screw with the nut.

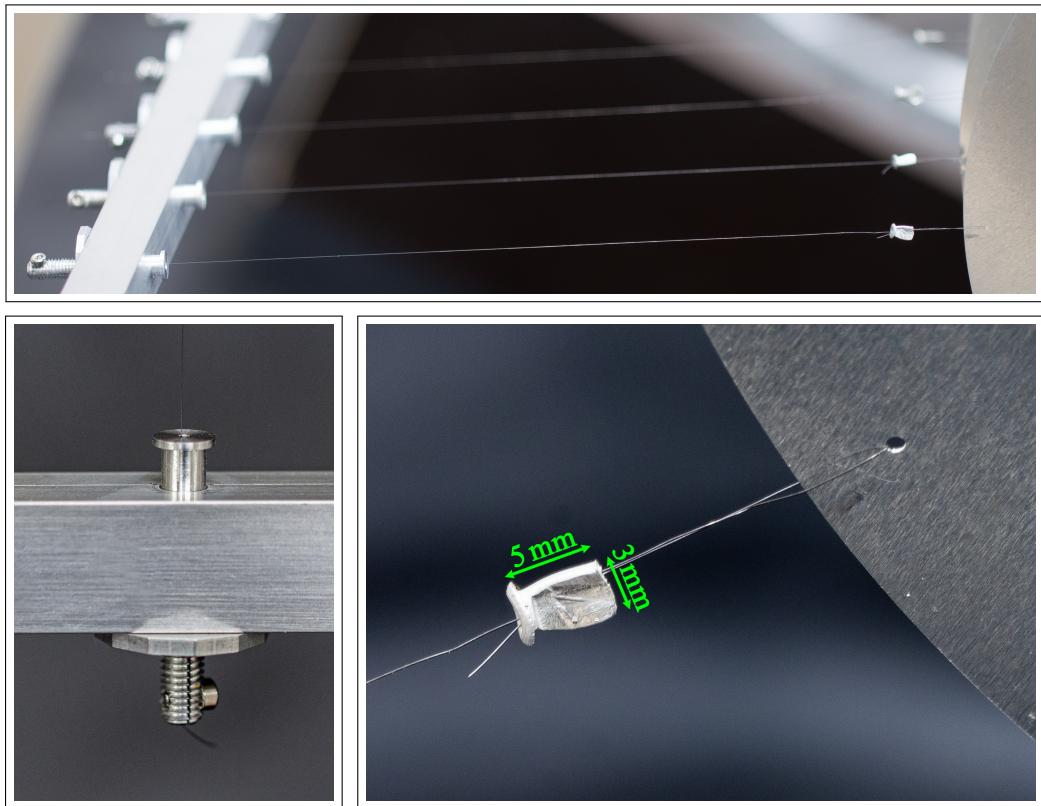


Figure 5.4.: Photographs of the thin tungsten spokes holding the target-disks in place. Top: Overview of the full connection. Bottom-left: Close-up of the connection to the metal bars using hollow screws, with the wire clamped inside the screw at its bottom part. The screw is 33 mm long to permit adjustments to tension and positioning by approximately ± 5 mm. Bottom-right: Close-up of the connection to the target-disk. The wire is fastened using a 0.5 mm thick aluminium sleeve, avoiding fracture-prone kinks or sharp edges in the wire.

Multiple stability tests of the wire and its connections were performed. In these tests, tungsten wires with a purity of 99.95 % and diameters ranging from 0.05 mm to 0.5 mm were suspended on a horizontal frame, and weights were gradually attached to their lower ends until either the connection or the wire failed. Each diameter was tested ten times. For

the thinner wires, the maximum, minimum, and average weight before fracture is given in Table 5.1. The majority of breakdowns occurred at or directly behind the aluminium sleeves that fastened the tungsten loop, likely due to stress applied at these positions during the manual clamping process. In two instances, the wire slid out of the aluminium sleeve, but only around weights that typically led to wire breakdown. For diameters thicker than 0.2 mm the permissible weight exceeded 5 kg. Further tests beyond this threshold were aborted, as such forces are sufficient to securely hold the lightweight target-disks in place.

Table 5.1.: Maximum permissible weight on the spoke connections using tungsten wires with varying diameter. Selected tungsten purity was 99.95 %. Wires were hung on a horizontal frame, increasing the weight on their lower end in steps of 5 g until breakdown. The mean permissible weight represents the average over ten separate measurements.

Diameter [mm]	Applied weight [g]		
	Minimum	Mean	Maximum
0.05	550	632 ± 16	680
0.10 ^a	605	838 ± 29	925
0.10	1630	1738 ± 27	1890
0.20	—	—	5000

^a Values for the ≥ 10 year old tungsten wire provided by the TU Dresden workshop used during installation tests.

Based on these maximum permissible force results, installing all target-disks inside the suspension was tested using 0.1 mm diameter tungsten spokes. The general installation procedure worked as intended. However, due to the lack of a clear indicator of already applied tensile strength, fine-adjustment proved challenging, and multiple wire breakdowns occurred during subsequent stability tests. Notably, these breakdowns were partially due to the use of much older wire during the installation procedure, as revealed by a second series of tensile strength tests conducted on this old wire. While the average permissible weight applied on the old wire was reduced by approximately 50 %, theoretically its strength should have been sufficiently large to withstand forces induced during the installation of the suspension, minor earthquakes, or the solenoid's quench².

²The situation of a small earthquake, not uncommon to the island of Japan, was approximated by reasonably shaking the whole muon-stopping target over an extended period of 30 s. The force was applied to the large outer ring structure, with bars and target-disks free-floating. The solenoid quench is approximated to induce at most up to 0.009 N of force in beam direction on each target-disk, while radial forces are canceled out. The study was kindly performed by a collaboration member, Naoyuki Sumi, using the electromagnetic and electromechanical simulation package OPERA-3D [87].

As a result of all tests, for tungsten, a slightly thicker wire with at least a 0.2 mm diameter is recommended to guarantee a stable operation over the full COMET runtime. Besides the added stability, using more rigid 0.2 mm thick wires significantly simplified the clamping and installation procedure. Alternatively, other materials like paramagnetic stainless steel, primarily composed of iron ($Z = 26$) and chromium ($Z = 24$), could also be used as spokes. A dedicated ICEDUST study shows that even in the case of an exaggerated spoke diameter of 0.5 mm, only approximately 100 muon stops per bunch are expected. Within the simulation's limited statistics, no influence of such spokes on muons stopped inside the target-disk was observed. Furthermore, given the low number of muon stops, any effects on DIO background from the spokes, regardless of material or thickness, are deemed negligible. In summary, 0.2 mm diameter tungsten wires are recommended for their optimal balance between thickness and stability, but other materials or thicker tungsten wires up to 0.5 mm diameter could be used without significant negative consequences.

5.3. The Installation Procedure

At the time of writing, neither the CTH nor the surrounding structure of CyDet was finalized. Consequently, the plans for the installation procedure are still in a preliminary state. The current approach aims to push the suspension into the CDC from the DS side. To avoid excessive movements, a cart to lift up and then smoothly roll the whole structure inside CyDet, starting from a tube with the same radius as CTH, was developed and is shown in Figure 5.5. After reaching the intended position, six height-adjustable mounting feet, one at each spacing plate between the larger rings, are used to fine-adjust the structure's x - and y -position and hold it in place. The largest ring at the suspension's end can be clamped to the DS side of the CyDet or the endplate of CTH to provide further stability. However, most of the structure's weight is around this outer ring part inside the DS CTH position. Thus, securing it via mounting feet should already be sufficient. If not compatible with the final design of CyDet, the clamped connection with the large ring could also be omitted.

It is planned to counteract the sagging of the three long bars holding the target-disks by using wedges inside the US CTH shielding. A schematic drawing of how such a design could be implemented is shown in Figure 5.6. A cut-out at the connection of the bars to the ring structure permits their movement in the longitudinal direction by ± 20 mm. Pushing the bars on top of the wedges lifts the structure's far end but also introduces a minor misalignment of the target-disks in the longitudinal direction. The exact details heavily depend on the final implementation of CTH. After parameters like sagging of the CTH are known in more detail, the shape and positioning of these wedges must be optimized to avoid the need for large adjustments. Currently, an uncertainty of at most 5 mm is envisioned for all directions.

5. The Muon-Stopping Target Suspension

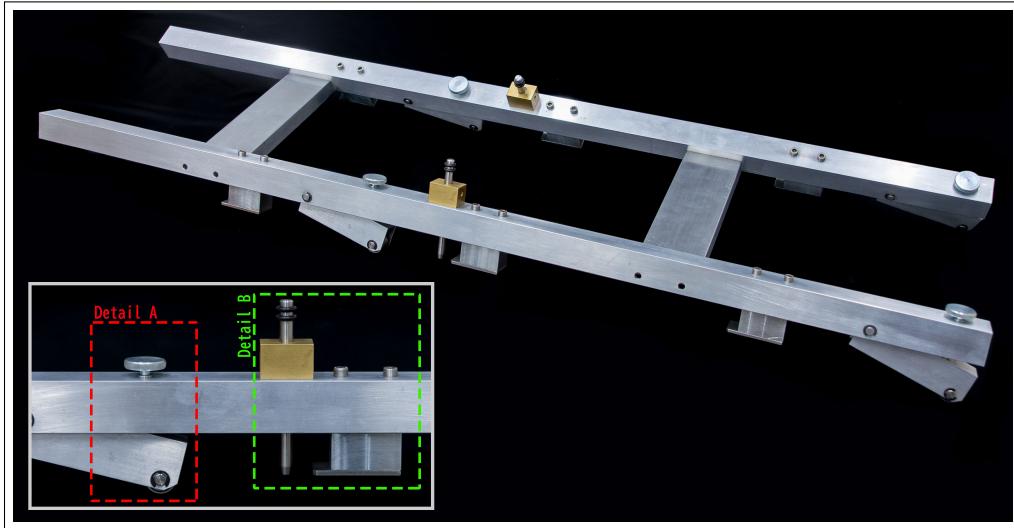


Figure 5.5.: Photograph of the cart used to install the suspension. The suspension's three large rings are placed on the nooks at the bottom of the cart, with a pin used to prevent the structure from sliding out (Detail B). Once the cart is correctly positioned and everything is secured, four wheels (Detail A) are lowered to lift the suspension and move it inside the CTH.

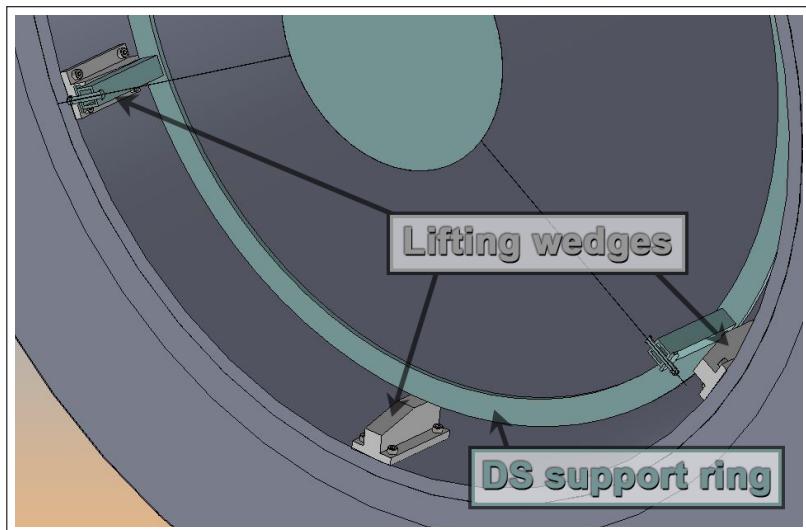


Figure 5.6.: Schematic of wedges inside US CTH, used to fine-adjust the suspension's y-position and provide stability at the far-side of the long bars.

5.4. ICEDUST Study of Key Parameters

5.4.1. Positional Accuracy

Section 5.3 stated the positional uncertainty in all directions to be at most 5 mm. The impact of such target-disk displacements on the number of stopped muons and signal acceptance is illustrated in Figure 5.7. To aid the interpretation of the data, extreme displacements up to 20 mm are included. For simplicity, offsets were applied to the entire muon-stopping target instead of individual target-disks. The results are presented for the default 17×0.2 mm configuration but are identical or less pronounced for configurations with an increased number of target-disks or with offsets at the US and DS side.

The top of Figure 5.7 shows the signal acceptance factor $A_{\mu \rightarrow e}^{\text{GM, ROI}}$, correlating to the probability for conversion electrons started from displaced target-disks to fulfill the COMET trigger condition. The relevance of the longitudinal target-disk position regarding signal acceptance was discussed in detail throughout Section 4.8. In short, signal acceptance around the CyDet center is constant and steeply falls off toward each side (compare Figure 4.34). If the entire muon-stopping target were pushed, for example, by 20 mm too far toward the US side, the outermost US target-disk's signal acceptance would drop by approximately two percentage points. However, signal acceptance of target-disks on the other side would likewise increase. As a net effect, only the two percentage point acceptance loss of the single target-disks getting pushed outside remains, which is negligibly small for the overall signal acceptance of the whole muon-stopping target, even with such an exaggerated displacement.

As long as conversion electrons can still reach the CDC, no effects on signal acceptance are expected for displacements in the horizontal or vertical direction. This assumption is confirmed within the statistical uncertainty of the study. The most probable helix radius for 105 MeV/c (conversion) electrons is approximately 335 mm, as shown in Figure 4.25 of Section 4.5.3. Hence, displacements larger than 20 mm could result in conversion electrons produced at the edges of a target-disk not reaching the required 5th CDC layer ($r \simeq 550$ mm) or cause complications for tracking. However, the expected size of displacements is reasonably small to avoid such undesirable consequences.

The effect of target-disk displacements on the number of muons stopped is shown in the middle of Figure 5.7. In case of a displacement along the beam direction, no significant difference is observed. This result is expected, as the beam profile remains constant over the entire length of the muon-stopping target (compare Figure 4.14 of Section 4.4.2) and further US or DS positioned target-disk are subject to a quasi-identical muon flux. Target-disk displacements in horizontal or vertical direction amount to, at most, a 5 % loss in the number of muons stopped. The shown result is asymmetric due to the small off-center muon beam profile on the target-disks, as previously shown in Figure 4.13 of Section 4.4.1. A displacement in the negative y - or x -direction increases this mismatch while compensating it in the other direction. In both cases, displacements of only a few millimeters cause no significant changes.

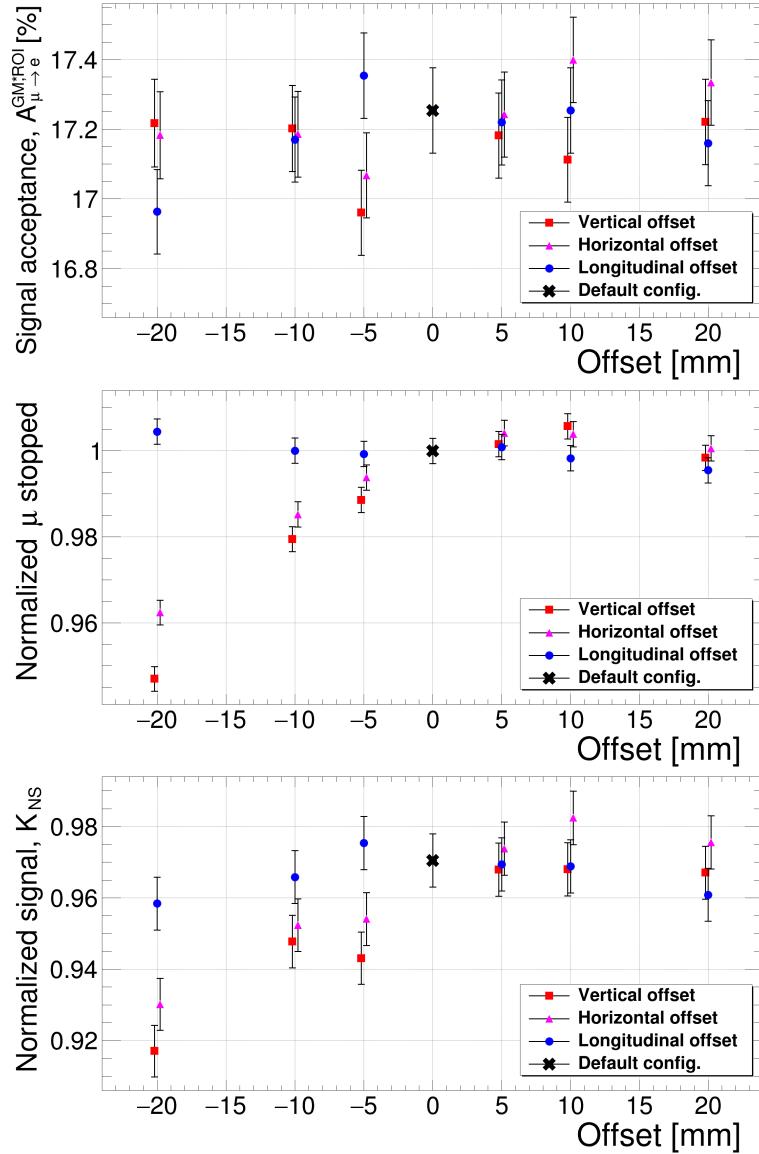


Figure 5.7.: Effect of target-disk displacements on signal acceptance (top), number of muons stopped (middle), and the combined normalized signal (bottom). Results are presented for the default 17×0.2 mm configuration. The data points of different offsets are slightly shifted to increase visibility but refer to their integer value. Shown error bars include only statistical uncertainties. Statistics correspond to 100 000 restarted conversion electrons (signal acceptance) and 14 bunches at 16×10^6 POT (number of muons stopped). A discussion of the results is provided in the text.

The combined value is given as the normalized signal K_{NS} and plotted at the bottom of Figure 5.7. Note that the shown K_{NS} of the default configuration is below one, as the normalization was performed to the default configuration, excluding any holding structure. Exact details of the effects on signal acceptance due to the long bars holding the spokes are discussed in the next section. In general, small positional uncertainties are shown to be negligible. However, vertical displacements appear to be more severe than in the longitudinal direction. Therefore, fine-adjusting the suspension's positioning using the US CTH wedges (see Section 5.3) should be prioritized accordingly. Notably, using the shown results, systematic uncertainties on the discussed COMET observables caused, for example, by slight variances of the beam profile due to differences in the magnetic field during beam transport, can be constrained to remain below a few percent (compare Section 4.2.5). Although these uncertainties are small, the overall result again highlights the importance of a carefully adjusted muon beam and motivates dedicated measurements of beam qualities prior to the physics run.

5.4.2. Size of the Three Long Bars

Three 1700 mm hollow aluminium square bars with a height and width of 15 mm \times 15 mm are used to bridge the gap between DS and US CTH and hold the target-disks (see Section 5.1.2). Conversion electrons interacting with these bars lose too much energy to be accepted by the momentum cut or become impossible to track due to scattering.

This loss in signal acceptance must be minimized. As a first measure, the bars are positioned at the largest possible radius to reduce the solid angle for collisions. A rough estimate for the probability of the bars blocking a conversion electron can be calculated from the blocked fraction of the circumference of a circle:

$$\text{Signal-blocked [%]} = 3 \cdot \frac{100 \cdot w}{2\pi \cdot r}. \quad (5.1)$$

For the used configuration with a radius of $r = 295$ mm and a width of $w = 15$ mm, the three bars block approximately 2.4 % of signal electrons. A more thorough study is performed with ICEDUST and shown in Figure 5.8. Three different materials and a wide range of considered sizes are included. The wall thickness was constant at 4 mm for all shown points. Simulations with thinner or thicker walls were performed but produced identical results and thus are not separately discussed. A higher atomic number generally leads to an increased interaction probability and more energy loss of the electron. Hence, carbon fibre-reinforced plastic bars perform slightly better than aluminium or stainless steel. However, differences around the chosen 15 mm width amount to only 0.1 percentage points and were neglected in favor of the easier machining of aluminium (see Section 5.1.2). Considering the more complex helical movement of electrons, total signal loss due to the three 15 mm wide bars is slightly larger than estimated with Formula 5.1 and amounts to 3.4 %.

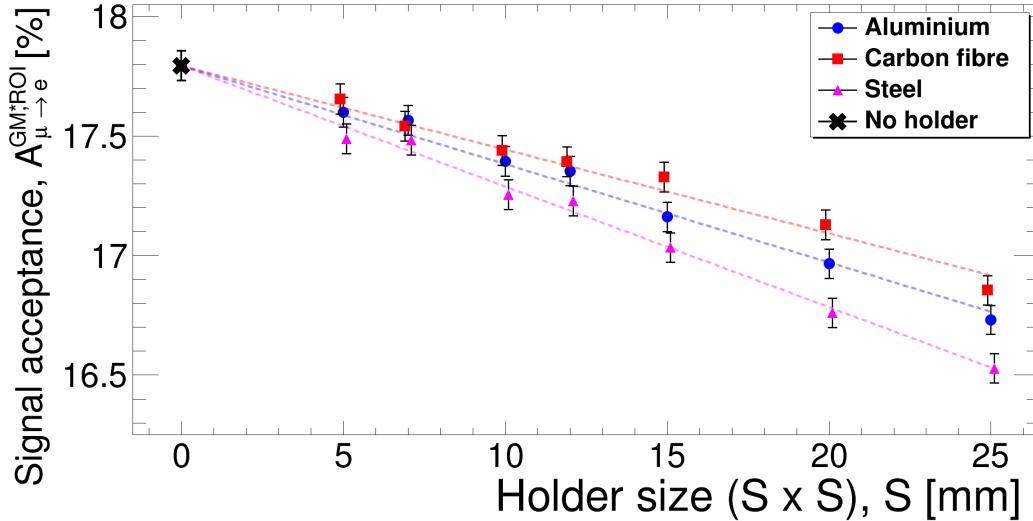


Figure 5.8.: Signal acceptance for varying suspension bar materials and size S of the bars. Results are presented for the default 17×0.2 mm configuration. The data points of different offsets are slightly shifted to increase visibility but refer to their integer value. Shown error bars include only statistical uncertainties, with statistics corresponding to 450 000 restarted conversion electrons. First-order polynomial fits are overlaid as trendlines. The wall thickness of the hollow square bars was fixed at 4 mm. Compared to no holder, the adopted 15 mm \times 15 mm bars cause a 3.4 % loss in signal acceptance.

However, the bars also block background particles originating from the muon-stopping target. Most significantly, protons created after nuclear pion capture, which usually amount to many CDC hits, are absorbed in the aluminium bars. Photons produced in this process do not directly contribute to CDC background. In summary, the total CDC occupancy for the default configuration with 17×0.2 mm aluminium target-disks drops from 32 % to 31 %. The size of this occupancy reduction is consistent for other configurations and equal to the loss in signal acceptance, namely 3.4 %.

5.5. Final Results and Summary

A photograph of the fully assembled suspension is shown in Figure 5.9. The total weight of the aluminium structure, including 17×0.2 mm aluminium target-disks, is approximately 13 kg. With 15 mm \times 15 mm hollow square bars the total signal acceptance but also CDC occupancy is expected to be reduced by 3.4 %. The far end of such bars was measured to bend downwards by 15 ± 1 mm due to self-weight.

The target-disks' x -positioning was initially adjusted by eye, then refined using a laser level and by measuring the length of each spoke. Using the adjustment screws at the connection to the bars, an accuracy of ± 1 mm can be achieved. The y -positions were adjusted accordingly. However, y - and z -direction are correlated due to the wedges inside the US CTH used to counteract the bar's sagging. If the wedges are adequately designed, an uncertainty not larger than 5 mm in both directions can be expected. As the effects of a longitudinal displacement are comparably less pronounced, fine-adjustments should favor the exact vertical positioning. Further details will depend on the final implementation of the CTH. However, no significant impact on the number of muons stopped or the signal acceptance is expected even for the worst-case 5 mm displacement.

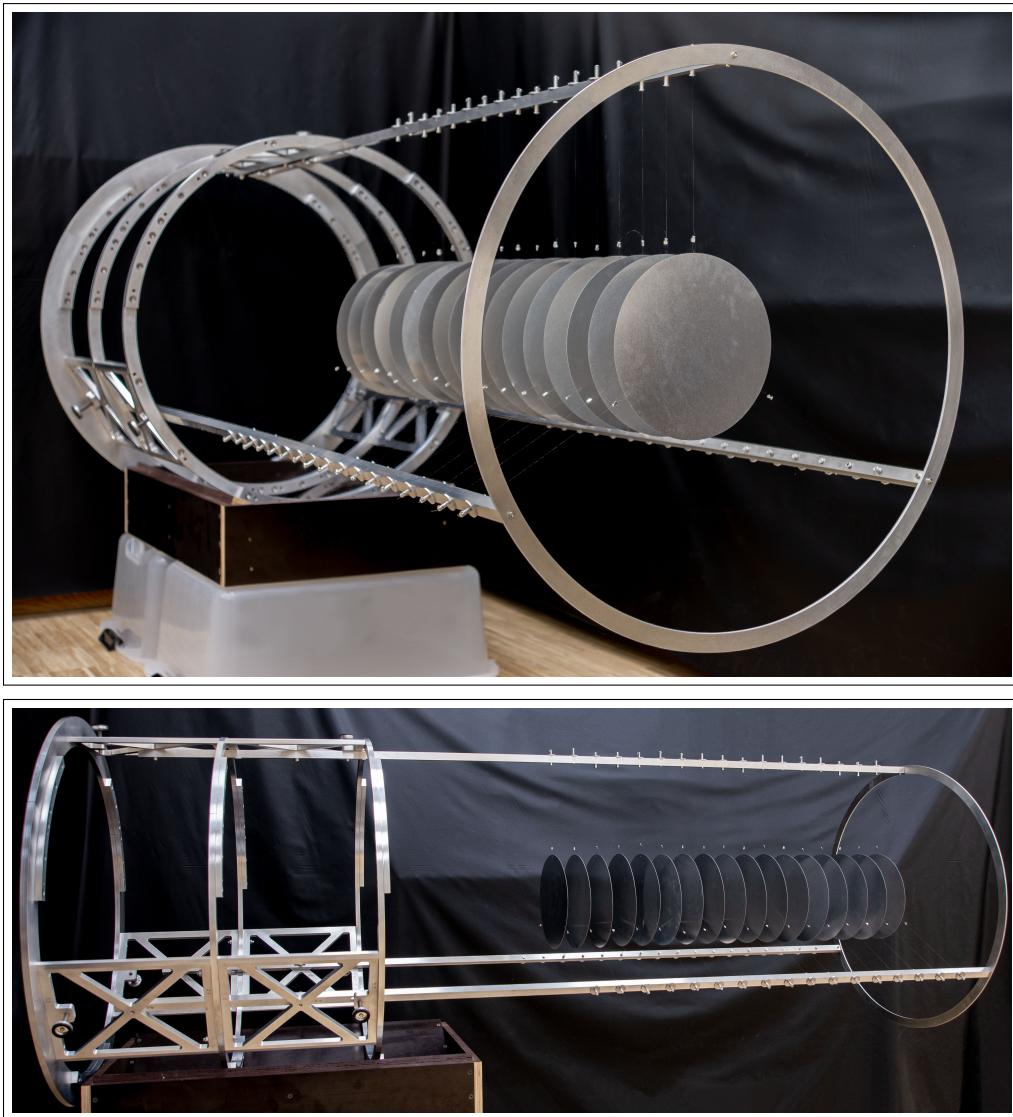


Figure 5.9.: Photographs from the front (US side) and the side of the fully assembled suspension structure. Contrast and sharpness of the images were enhanced to increase clarity. The number of target-disks and their positioning corresponds to the default configuration (see Section 4.1). For details see text.

6. The Muon-Stopping Monitor

The ultimate output of the COMET experiment will be the muon-to-electron conversion rate in aluminium or a limit thereof. To this end, the partial branching ratio $\mathcal{B.R.}$ or the SES introduced in Section 2.2.2 serve as a figure of merit. The numerator of the SES given in Formula 2.4 represents the number of mono-energetic conversion electrons, which are measured with the CyDet detector system. Detection of a conversion electron alone would already confirm the existence of a CLFV process. However, to confirm or refute varying theoretical predictions, additionally the precise measurement of the branching ratio's denominator – represented by the total number of stopped muons – is crucial.

The proposed idea for this measurement utilizes muonic X-rays emitted after muonic shell capture. Using formula:

$$N_{\mu}^{\text{stopped}} = \frac{F_{i \rightarrow f}}{\epsilon_T \cdot P_{i \rightarrow f}}, \quad (6.1)$$

the number of stopped muons N_{μ}^{stopped} can be calculated from the full-energy peak efficiency (or acceptance) ϵ_T , the number of measured events within the full-energy gamma peak $F_{i \rightarrow f}$ and the muonic X-ray line's transition probability (or intensity) $P_{i \rightarrow f}$. As introduced in Section 1.1.2, a stopped muon causes a multitude of X-rays when cascading down to its $1s$ orbit. The energy and intensity of each produced X-ray strongly depend on the element in which the muon is stopped and the transition's initial (i) and final (f) orbit. Fortunately, extensive amounts of precision data on muonic X-rays have been recorded over the past decades due to their relevance for nucleus charge distribution measurements [88]. The energies and transition intensities for the most relevant lines of muonic aluminium are summarized in Table 6.1. With a comparatively high intensity of 79.65 % and an energy of 346.8 keV, the most promising X-ray line for COMET originates from the $2p \rightarrow 1s$ transition. The intensity of the $3d \rightarrow 2p$ line, at 62.50 %, is comparable. However, this alternative transition has a factor five reduced energy of only 66.0 keV. As a result, its measurement is significantly more challenging due to higher detector background and photon attenuation at this energy level.

The X-rays will be measured by an HPGe semiconductor detector. It is planned to use a GMX n-type coaxial HPGe manufactured by ORTEC, which has a relative efficiency of 20 % at 1.33 MeV [89]. The active crystal has a length of 55.4 mm, a diameter 52.5 mm, and a 47.8 mm deep borehole with a diameter of 9.9 mm. The detector is placed at the far

Table 6.1.: Energy E and intensity I of muonic X-rays in aluminium. Measured energies for the muonic Lyman (or K) series are taken from recent publication [90]. Balmer and Paschen series are calculated manually from these levels. Transition intensities are taken from [66]. Note that both publications assume the summed intensity of the Lyman series to match 100 %. Higher series are omitted because their X-ray energies become negligibly small (≤ 30 keV).

n	Lyman: $n \rightarrow 1$		Balmer: $n \rightarrow 2$		Paschen: $n \rightarrow 3$	
	E [keV]	I [%]	E [keV]	I [%]	E [keV]	I [%]
2	346.828	79.65 ± 0.60	—	—	—	—
3	412.87 ± 0.05	7.43 ± 0.29	66.04	62.50 ± 1.80	—	—
4	435.96 ± 0.10	4.72 ± 0.20	89.13	9.70 ± 1.90	23.09	33.50 ± 5.30
5	446.61 ± 0.10	3.89 ± 0.17	99.78	4.23 ± 0.32	33.74	6.07 ± 0.70
6	452.38 ± 0.10	2.35 ± 0.11	105.55	1.44 ± 0.19	39.51	1.98 ± 0.51
(7 to ∞)	—	1.96	—	1.33	—	0.74

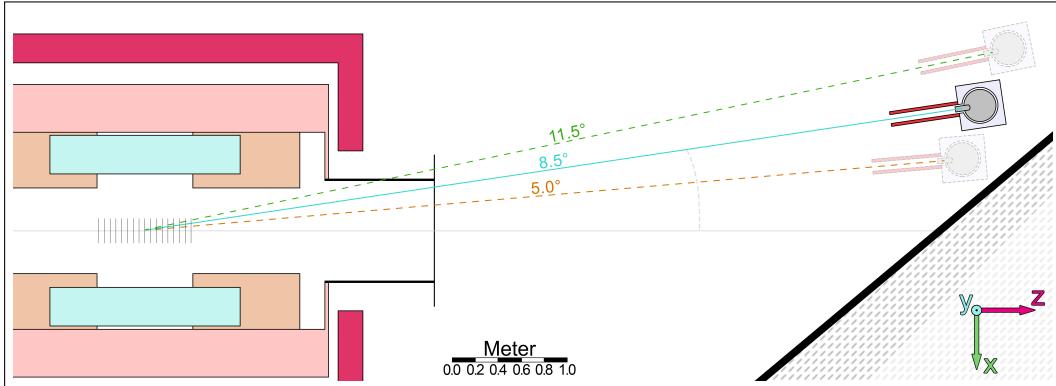


Figure 6.1.: Technical sketch of the HPGe positioning relative to the muon-stopping target. The figure corresponds to a top-down view of the COMET hall. Muons arrive from the US side on the left and hit the target-disks (17 black lines) at the CyDet center. The light-blue and -orange rectangles correspond to the CDC and CTH geometries, respectively. Surrounding magnets and coils are indicated as red squares. The beam duct at the right end of CyDet is shown in black. Besides the adopted position at an angle of 8.5° (cyan), the more central 5.0° (orange dashed) and further outside 11.5° (green) HPGe locations discussed in Sections 6.2 are included as reference.

DS side behind the end of CyDet and its surrounding solenoid¹. A technical sketch of the positioning is shown in Figure 6.1. It is based on the most recent layout of the COMET hall and provides an approximate understanding of the detector's location relative to the muon-stopping target. However, changes in this design must be anticipated as details such as cable-pathing around the CyDet end or the position of the readout DAQ systems were still in an active planning phase at the time of writing. Varying detector positions discussed in Sections 6.2 are already included in the figure.

A HPGe detector is the preferred choice due to its excellent energy resolution, typically around 2 keV@500 keV [89]. This resolution helps to separate the X-ray line from nearby gamma-background² and provides an excellent signal-to-noise ratio for improved peak statistics. However, HPGe detectors are typically used in low(er) count rate environments. Especially when compared to COMET's extremely fast bunch separation of only 1170 ns, the response and signal processing time in the order of $O(10 \mu\text{s})$ will pose challenges. Proposed solutions to mitigate this issue and an assessment of the detector's overall feasibility will be the main topics of this chapter. Since a more detailed summary of the working principle of a HPGe detector exceeds the scope of this work, interested readers are instead referred to the excellent works of Gordon R. Gilmore [93] or a more recent publication by S. Turkat [94].

6.1. Methodology: The Two Main Parameters

The two primary parameters an experimentalist must consider for every detector setup are its expected background (rate) and signal acceptance. Similar to the discussion around the muon-stopping target, both parameters oppose each other. Very stringent shielding and collimation might remove most of the HPGe background but also attenuate too many signal X-rays to achieve a reasonable measurement. A lenient shielding and collimation would allow most X-rays to reach the HPGe, but background events will dominate the signal. As a starting point, this section first highlights key points and the calculations for the detector background and signal acceptance individually. Subsequently, Section 6.2 applies this information to draw conclusions and present results.

6.1.1. Assessing the Germanium Detector Background

Muonic X-rays are directly connected to muon stops and have a characteristic material-dependent energy. Both aspects serve as effective tools to suppress background from other radioactive nuclei produced in the COMET area and from muons stopped in surrounding

¹At this far DS position, the magnetic field strength dropped to approximately 10^{-2} T, low enough for a normal operation of the HPGe.

²For example, the most intense $2p \rightarrow 1s$ aluminium X-ray line at 346.8 keV should be separable from the 351.9 keV or 351.1 keV gamma-lines from ^{214}Pb and ^{211}Bi isotopes produced in the ^{238}U and ^{235}U natural decay chains [91, 92].

components³ or target-disk impurities. However, muonic X-rays are emitted within picoseconds after a muon is stopped. Hence, their measurement competes with the intense background produced by the beam arriving at the CyDet detector region.

Figure 6.2 shows the origin of the expected HPGe background. For now, particle momentum and PID are disregarded, but a thorough discussion can be found later in Section 6.3.1. Without shielding, an average of 182 particles deposit energy within the detector per 16×10^6 POT bunch. Unfortunately, the origins of these background particles are spread out over the COMET hall. Hence, specialized targeted shielding that only focuses on a confined spot or background location is not feasible. However, a notable fraction of particles hit the detector from its rear or side, either after reflecting from the building's back wall or originating from locations not in direct line of sight to the muon-stopping target. Assuming perfect shielding around the back and sides and limiting the HPGe's front view in the direction of the target-disks reduces the number of energy-depositing particles to approximately 13 per bunch. If these 182 (or 13) energy depositions occurred over an extended period, such as within one second, they would not pose any problem. However, from the perspective of the comparatively slow HPGe detector, all particles within the short 1170 ns period between two bunches will arrive almost simultaneously. This results in overlapping energy depositions inside the detector, commonly known as pileup, and the distinct X-ray energy cannot be identified anymore. Two consequences can be drawn from this: Firstly, the commonly used *background rate* is unsuitable for COMET's situation. Instead, the more relevant parameter *probability for a background-free bunch* will be used to assess the HPGe background. Reducing the average number of energy depositions per bunch below one is crucial to raise this probability to a feasible level. Secondly, to achieve such a considerable reduction in background, the HPGe must be positioned further away from the active CyDet area, and a dedicated collimator must be developed to limit the number of particles hitting and depositing energy in the detector from the front.

The detector background is assessed using ICEDUST (see Section 3). Unfortunately, a meaningful background study cannot get around a full simulation accounting for all particles because any particle or interaction around the HPGe could lead to an energy deposition⁴. To make matters worse, a reasonable shielding design limits the number of energy depositions in the detector to less than one per bunch, and as a result, the statistics of the previously used 60-bunch MC5 dataset are insufficient. Hence, the resampling procedure described in Section 3.1 had to be adapted to overcome this problem. In a one-time high statistics run, a 600 bunch-sized dataset of only particles reaching a closely confined space around the HPGe was recorded. Subsequently, all collimator and shielding studies were performed by resampling this dataset.

³Because no other material in line of sight to the detector will consist of aluminium, any characteristic muonic X-ray of aluminium can directly be connected to the muon-stopping target.

⁴This means that previously established procedures to optimize CPUh, specifically those applied in muon-stopping studies where simulations are limited to a few particles of interest such as muons or pions, cannot be used.

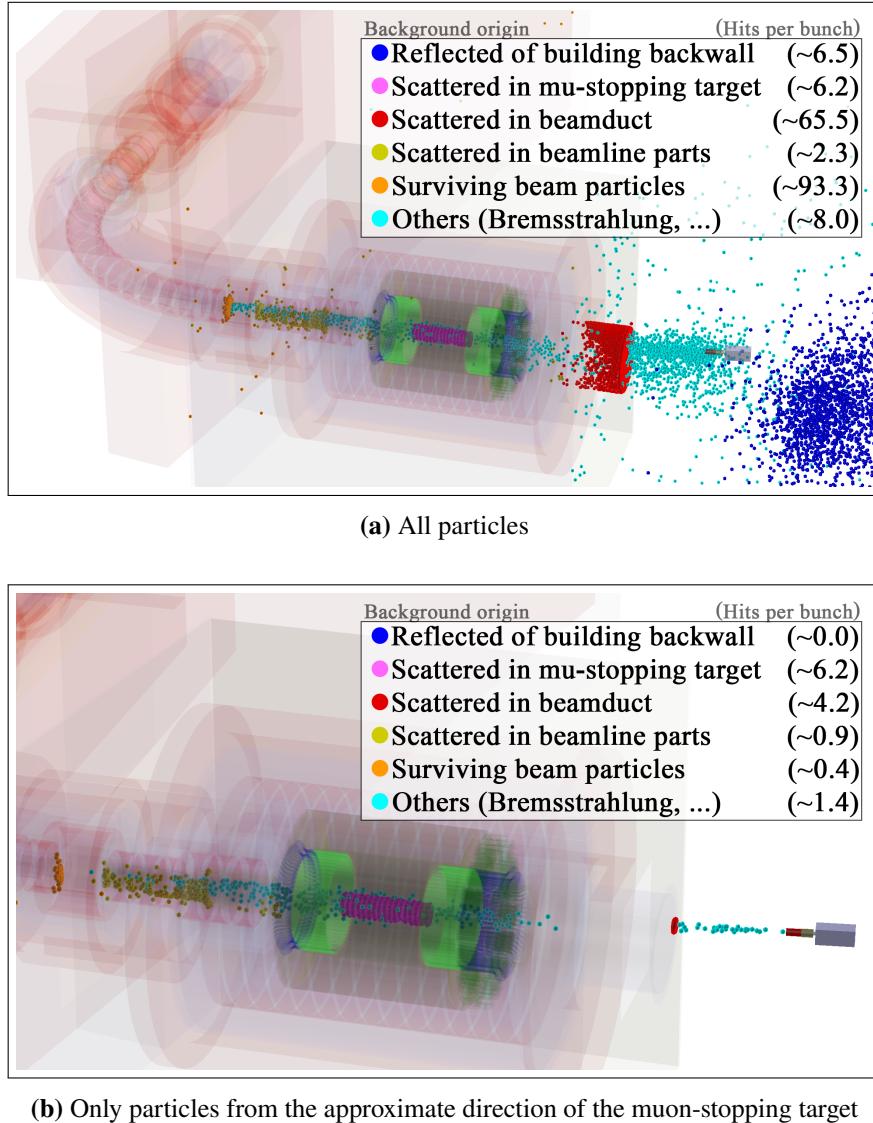


Figure 6.2.: The last point of interaction before particles deposit energy in the HPGe. The data corresponds to 300 bunches with 16×10^6 POT. The HPGe was positioned 4 m DS at an angle of 8.5° relative to the muon-stopping target's center. Interaction points are color-coded corresponding to key geometries. Particles hitting the HPGe without prior interaction within the CyDet area are summarized as 'surviving beam particles' and mostly consist of electrons. Particle momentum distributions and PID are discussed later in Section 6.2. The total number of HPGe hits without (with) a basic direction filter amounts to 182 (13) per 16×10^6 POT bunch.

Problems with the default Geant4 physics During the preparation of this resampling dataset, discrepancies in the expected number and energy of muonic X-rays became apparent. Within default Geant4, the cascade process of muons to their ground state is only approximated in a generic and highly simplified manner. Taken from the source code, muonic energy levels are calculated according to:

$$E(n) = 13.6 \text{ eV} \cdot \frac{Z^2}{n^2} \cdot \frac{m_\mu \cdot M_N(A, Z)}{m_\mu + M_N(A, Z)} \cdot \frac{1}{m_e} \quad (6.2)$$

with atomic number Z , nuclear mass $M_N(A, Z)$ taken from [95], muon mass m_μ , and electron mass m_e . Simplified formulas to decide between Auger electron or X-ray emission and cascade transition probabilities are taken from comparatively old publications [96] and [9], respectively. Besides the mismatch in transition energies, this causes a significant increase of up to 25 % in the number of Auger electrons at the expense of fewer X-rays for the relevant $2p \rightarrow 1s$ and $3d \rightarrow 2p$ transitions. Therefore, the Lyman; Balmer; and Paschen series of muonic aluminium X-ray transitions in Geant4 were manually adjusted to reflect the data from [90] and [66] as presented in Table 6.1. Transitions of quantum numbers $n \geq 4$ were left unchanged because of their negligible impact. The energy difference between these levels is comparably small, resulting in a cascade primarily characterized by the emission of low-energy Auger electrons. Neither these Auger electrons nor similarly low-energy X-rays can reach the HPGe detector due to the long distance and intervening material.

6.1.2. Calculating the Germanium Detector Signal Acceptance

To study the signal acceptance – used synonymously to the full-energy peak efficiency of the HPGe – a custom MC simulation was developed. This MC code is not based on any external framework like Geant4 but was developed entirely from scratch in the C++ programming language, thus necessitating a more detailed introduction. In comparison to ICEDUST, the custom simulation focuses solely on the X-ray of interest emitted only in the direction of the HPGe⁵, thereby significantly easing CPUh requirements for large statistics. Additionally, the MC is built modularly, which will be used to individually investigate effects from different collimator designs, varying detector solid angles, intrinsic detector efficiency, or X-ray interactions along their way to the HPGe.

The MC follows a straightforward start-and-count approach. X-rays are started from normalized muon-stopping distributions recorded and discussed throughout the previous chapter 4. Each starting position is weighted by its corresponding detector solid angle, which is defined as the fraction of the area of a unit sphere covered by the HPGe detector.

⁵The term “direction of the HPGe” used throughout this section refers to the solid angle in which X-rays can be emitted to reach the *active volume* of the HPGe detector.

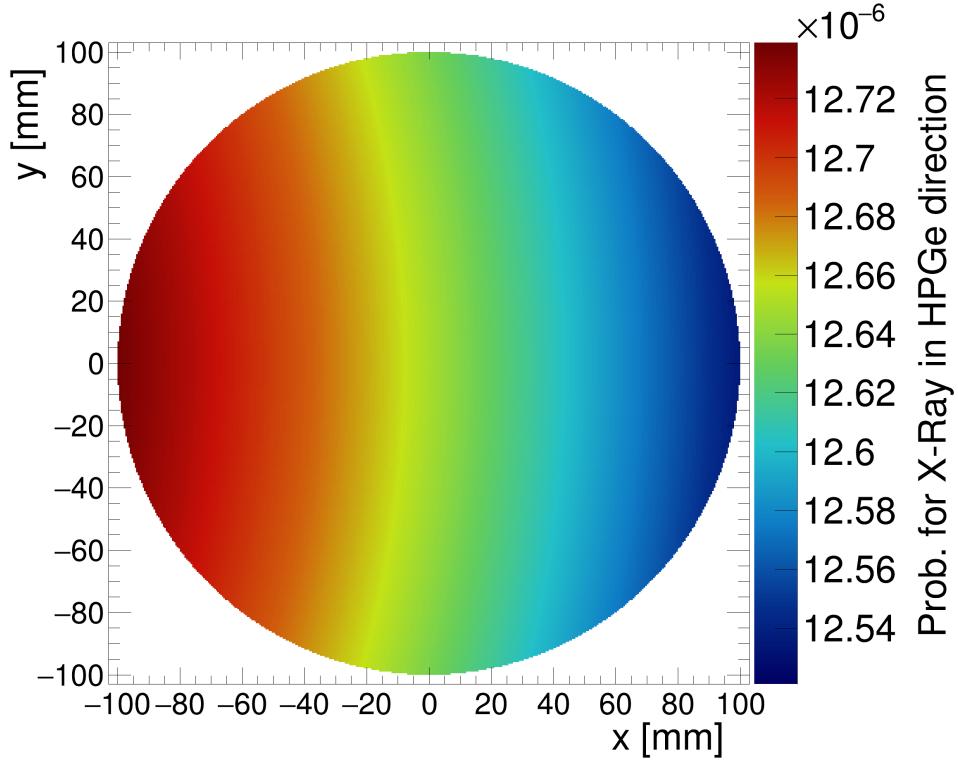


Figure 6.3.: Example of calculated solid angles for the middle target-disk (number 9 of 17). Values are divided by the area of a unit sphere (4π) yielding the probability for X-rays reaching the HPGe detector. The detector was positioned at a distance of 4 m at an angle of 8.5° . Relative to the muon-stopping target, the HPGe is positioned on the left (compare Figure 6.1). Hence, each target-disk's solid angle is dominated by the shorter (left) toward larger (right) distance to the detector and changes by approximately two percent. Differences between neighboring target-disks amount to approximately three percent.

The solid angle for a circular detector surface can be calculated with:

$$\Omega = 2\pi \cdot \left[\frac{r_a^2 + d^2 - d \cdot \sqrt{r_a^2 + d^2}}{r_a^2 + d^2} \right] = 2\pi \cdot \left[1 - \frac{d}{\sqrt{r_a^2 + d^2}} \right], \quad (6.3)$$

using the active detector crystal radius r_a and the distance d between the detector crystal and the X-ray starting position [97]. Formula 6.3 is accurate only for positions at the center of the assumed unit sphere when the distance vector \vec{d} is perpendicular to the detector surface. Following the approach outlined in previous work [77], the solid angle for off-center

positions requires a slightly corrected effective radius r_a^* , which is calculated as:

$$r_a^* = r_a \cdot \cos(\Theta) \quad \text{with: } \Theta = \arccos \left(\frac{\vec{d}_\perp \cdot \vec{d}_{\text{off-center}}}{\|\vec{d}_\perp\| \cdot \|\vec{d}_{\text{off-center}}\|} \right). \quad (6.4)$$

Here, \vec{d}_\perp and $\vec{d}_{\text{off-center}}$ correspond to the perpendicular and off-center distance vectors, respectively, as illustrated in Figure 6.4. Due to the large detector distance compared to the small detector radius, the effect of this correction changes the resulting solid angle by at most 0.1 %. Figure 6.3 illustrates an example of calculated solid angles for the middle target-disk.

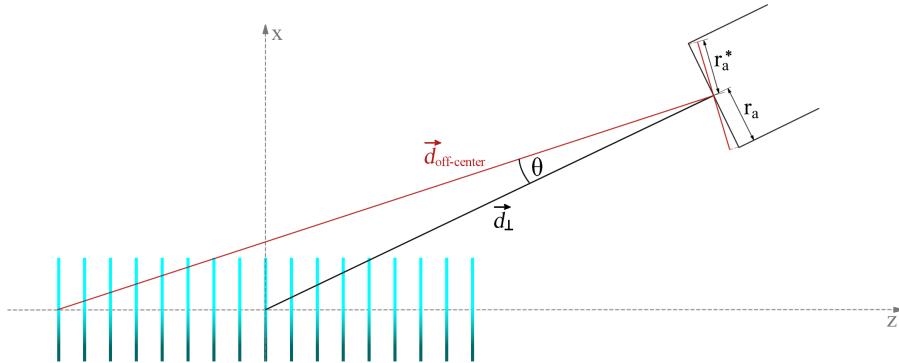


Figure 6.4.: Definition of effective detector crystal radius r_a^* , calculated with Formula 6.4 and accounting for off-center X-ray starting positions.

The full-energy peak efficiency (or signal acceptance) is calculated as the fraction of X-rays depositing their full energy within the HPGe, hereafter referred to as *surviving*, to all X-rays that were emitted. The number of surviving X-rays can be expressed as:

$$N_{\text{surviving}}(x, y, z, i, j) = N_{\text{started}} \times \begin{cases} \text{a) } H(P_{\text{no-interaction}}(x, y, z, i, j) - X_1) \\ \text{b) } H(P_{\text{survive intr.}}(i, j) - X_2) \\ \text{c) } \delta_{\text{Hit-CTH}}(x, y, z, i, j) \\ \text{d) } \delta_{\text{Hit-Collimator}}(x, y, z, i, j) \end{cases} \quad (6.5)$$

The variables (x, y, z) and (i, j) refer to the X-ray's starting position and the hit position on the detector surface, respectively. The function $\delta(\dots)$ equals one if there would be no collision, or zero if there would be a collision with a corresponding geometry part. Variable $H(\dots)$ corresponds to the Heaviside step function and is zero for negative and one for positive arguments. Lastly, $X_{1/2}$ represents random numbers sampled from a uniform distribution on the unit interval $[0, 1]$. Each part in Formula 6.5 contributes multiplicatively and is evaluated separately, as explained below.

Table 6.2.: Parameters used to calculate the interaction probability of the 346.8 keV X-ray associated with the $2p \rightarrow 1s$ transition. Constants are taken from the NIST database [98]. Example path lengths corresponding to the center of the middle target-disk (number 9 of 17), and the HPGe center at an angle of 8.5° and distance of 4 m are provided for reference.

Material k	Mass attenuation coefficient (μ_k/ρ_k) [cm 2 /g]	Density ρ_k [g/cm 3]	Path length d_k [cm]
Helium	0.101	0.18×10^{-3}	266.50
Air ^a	0.101	1.29×10^{-3}	102.00
Steel ^b	0.101	8.00	0.20
Aluminium	0.098	2.70	0.18
Beryllium ^c	0.090	1.85	0.05

^a Assumed 78 % nitrogen, 21 % oxygen, 1 % argon.

^b Average density of type 304 stainless steel. Assumed 72 % iron, 18 % chromium, 10 % nickel.

^c Corresponding to the ultra-thin entry window of the HPGe detector. Effects from the HPGe dead layer, typically thinner than one-tenth of a micrometer, were disregarded.

a) The X-ray can be absorbed through photoeffect or Compton scatter away from the detector. Using the Bouguer-Lambert law, the probability that a photon reaches the HPGe *without* interaction can be calculated with:

$$P_{\text{no-interaction}}(x, y, z, i, j) = \prod^k \exp\left(-\frac{\mu_k}{\rho_k} \cdot \rho_k \cdot d_k\right). \quad (6.6)$$

Variable d_k corresponds to the traversed distance through material k . Assumed values for the density ρ_k and the material- and energy-dependent total mass attenuation coefficient (μ_k/ρ_k) are taken from the NIST database [98] and listed in Table 6.2. An example of calculated probability values is shown in Figure 6.5.

b) An X-ray reaching the detector does not necessarily deposit its full energy within the HPGe crystal. Especially after Compton scattering, the scattered photon can often leave the crystal without further interactions, causing missing energy. The fraction of measured full-energy deposition events to all impinging X-rays defines the intrinsic (full-energy peak) acceptance. For a given photon energy, it primarily depends on detector geometry such as crystal and borehole size. Although the photon's angle of entry in the detector can have an effect, it becomes negligible for a large detector-source distance (> 1 m). The intrinsic efficiency for the employed HPGe was determined with ICEDUST, and a profile highlighting the dependence of the X-ray's detector entry position is shown in Figure 6.6.

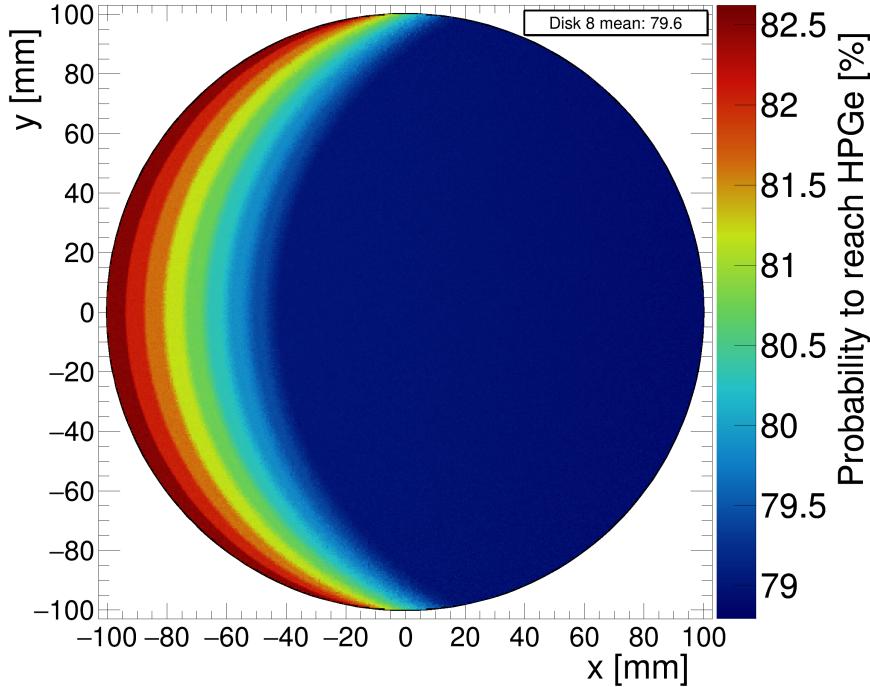


Figure 6.5.: Probability of an X-ray reaching the HPGe detector without being absorbed or scattered in intermediate material. Only photons emitted in the direction of the HPGe are considered. Relative to the muon-stopping target, the HPGe is positioned on the left (compare Figure 6.1). The result is presented for the middle target-disk (number 9 of 17). Each (x, y) point corresponds to an average of 500 started X-rays with randomly chosen HPGe surface points (i, j) . The probability $P_{\text{no-interaction}}(x, y, z, i, j)$ is primarily dominated by the added 0.2 mm aluminium for each additional traversed DS target-disks, resulting in a sickle-shaped pattern. Probability loss from other components, like the steel beam-duct or the traversed air, is approximately 17.5 %.

- c) If the detector is located at large angles ($\gtrsim 9^\circ$) with respect to the muon-stopping target, or for the more US positioned target-disks, the CTH geometry might block the line-of-sight between target-disks and HPGe. As the inner CTH shield consists in large parts of lead, X-rays cannot pass the CTH without being absorbed or scattered. Hence, the X-ray's survival chance is set to zero in these cases.
- d) While in the best case all X-rays that could reach the detector are also accepted, the expected background intensity necessitates a front-facing collimator. The collimator will block selected areas on the detector surface or limit the accepted angle of X-rays. Similar to point c), the survival chance is assumed to be zero if an X-ray's trajectory intersects any part of an investigated collimator design.

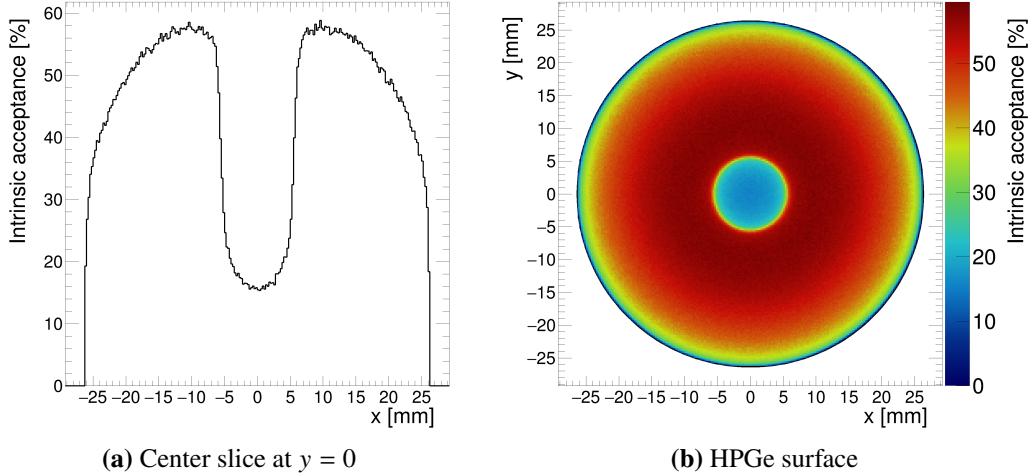


Figure 6.6.: Intrinsic full-energy peak acceptance of the employed HPGe detector, simulated for 346 keV X-rays. These X-rays enter the HPGe crystal approximately perpendicular due to the long distance between the detector and the start position. X-rays Compton scattering at the HPGe’s sides (larger radii) are more likely to exit the crystal prematurely, resulting in missing energy and a smaller intrinsic full-energy peak acceptance. A similar loss is visible around the HPGe’s borehole position, where the crystal thickness is significantly reduced.

Following the described procedure of Formula 6.5 yields the full-energy peak efficiency:

$$\epsilon_T(x, y, z, i, j) = N_{\text{surviving}}(x, y, z, i, j) / N_{\text{started}} \quad (6.7)$$

for a single position on a target-disk. Note that only X-rays in the direction of the HPGe are considered, hence N_{started} includes a weighting factor accounting for the solid angle shown in Figure 6.3. Each (x, y, z) position is sampled and averaged over 200 times to incorporate effects from different (i, j) hit positions on the detector crystal. Finally, the average full-energy peak efficiency $\bar{\epsilon}_T$ of the entire muon-stopping target is determined by computing a weighted sum that accounts for the muon-stopping profile and the number of muons stops of each target-disk.

The procedure was initially developed as part of the preceding work [77] and further refined and adjusted to COMET conditions as part of this thesis. The validity of the custom MC was extensively tested and is demonstrated in Figure 6.7. Compared to results from measurement and ICEDUST, all three are in agreement. However, the latter two only provide the final average full-energy peak efficiency. Hence, the discussion and results in this study are solely based on the custom MC simulation to utilize the information of the modular approach and investigate the interplay of the introduced components a) to d).

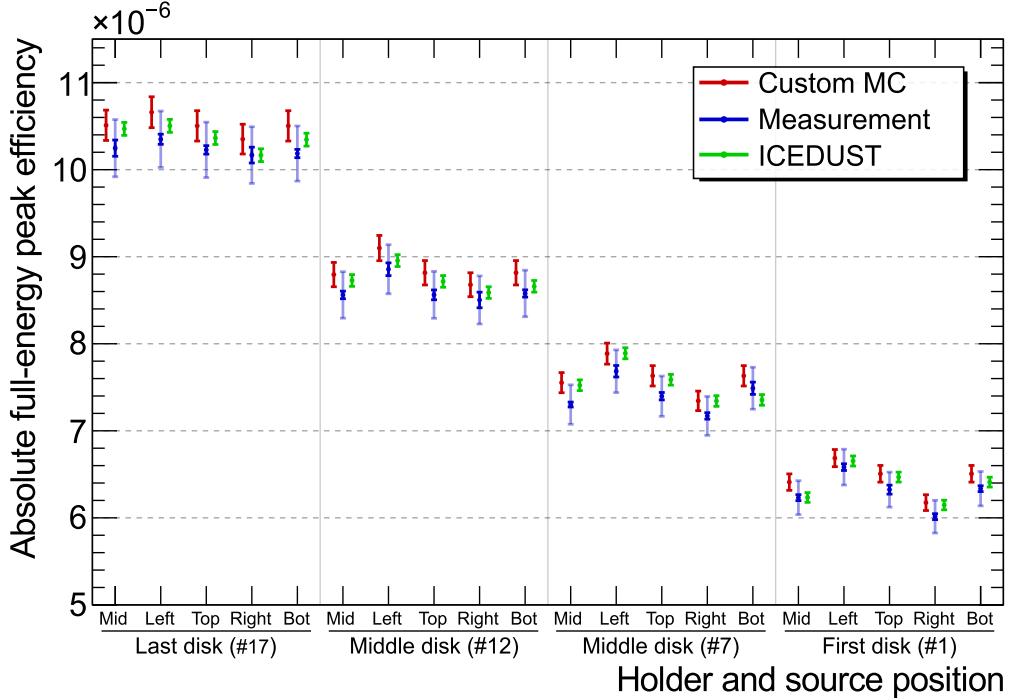


Figure 6.7.: Comparison between full-energy peak efficiencies of measurement, ICEDUST, and the custom MC simulation. The measurement was performed using the 356 keV gamma-ray of a ^{133}Ba source. Note that for this initial study, the HPGe detector was positioned at a larger angle of 15° to the left and a smaller distance of 3.5 m relative to the center of the muon-stopping target. The positional indicators, such as left, middle, or top, denote specific positions of the point source on each target-disk, with left being closest and right being furthest away from the HPGe. The uncertainties shown for both simulations are purely statistical. In the case of the measurement, the dominant uncertainty was systematic and stemmed from the given initial activity of the source. The systematic component is shown as a semi-transparent overlay atop the smaller statistical uncertainty. As can be seen, results from all three methods are in agreement. For further details see [77].

6.2. Results: Positioning and Collimator

In consequence of the discussion regarding Figure 6.2, the HPGe detector will be moved from its initial 4 m position to a location further DS. The maximum distance is limited by the size of the experimental hall (see Figure 6.1). In this study, a distance of 6 m between the center of the muon-stopping target and the center of the HPGe was adopted. This still leaves enough space between the wall and the detector to account for the installation of shielding and cabling. Moving the HPGe further away from the target-disks reduces the solid angle, hence the average full-energy peak efficiency $\bar{\epsilon}_T$ drops from 4.82×10^{-6} at 4 m distance down to 1.91×10^{-6} . Considering common use-case scenarios of HPGe's, like source-activity measurements, such a full-energy peak efficiency might seem too small to measure a significant enough signal. However, for COMET's situation, this distance is acceptable for the science case and unavoidable to reduce the detector background. Slightly adjusting Formula 6.1, the $2p \rightarrow 1s$ X-ray count rate in the HPGe can be estimated with:

$$R_{\text{@HPGe}}(2p \rightarrow 1s) = \underbrace{\left[N_{\mu/\text{bunch}}^{\text{stopped}} \cdot R_{\text{bunch}} \cdot P(2p \rightarrow 1s) \right]}_{R_{\text{@Disk}}(2p \rightarrow 1s)} \cdot P(\text{free-bunch}) \cdot \bar{\epsilon}_T. \quad (6.8)$$

The number of stopped muons per bunch $N_{\mu/\text{bunch}}^{\text{stopped}} = 7834$, the transition probability $P(2p \rightarrow 1s) = 79.65\%$ (see Table 6.1), and the rate of bunches $R_{\text{bunch}} = 153\,318\,\text{s}^{-1}$ yield the rate of produced X-rays from all disks $R_{\text{@Disk}}(2p \rightarrow 1s) \approx 1 \times 10^9\,\text{s}^{-1}$. With a yet arbitrary but very conservative estimate of free bunches $P(\text{free-bunch}) = 5\%$ where X-rays can be measured, and an average full-energy peak efficiency $\bar{\epsilon}_T = 1.91 \times 10^{-6}$, one still gets $R_{\text{@HPGe}}(2p \rightarrow 1s) = 91\,\text{s}$. Hence, enough to obtain a clear peak over the 146 days of COMET Phase-I runtime.

Besides the modified distance, the angle at which the HPGe is placed was investigated. Due to the target-disks being arranged in a row along the beam center, using large detector angles with long collimators will result in X-rays emitted from the first or last target-disks being cut off. On the contrary, the surface from where X-rays are accepted increases for smaller HPGe angles. For a simple 2 m long tube-collimator, total signal acceptance could be improved by approximately 5 % if the detector is placed at only 5°. However, the closer positioning to the center of the passing beam almost doubles the number of particles impinging on the detector's front. The effect is shown in Figure 6.8. Consequently, larger angles are preferable.

The maximum reasonable angle for the HPGe is defined by the beam duct design at the end of CyDet (compare Figure 6.1). The beam duct's side wall is made of 10 mm thick stainless steel, and 45 % of 346 keV X-rays would survive after passing this much material. However, due to the shallow angle of the X-ray's trajectory through the beam duct side, a much larger distance of roughly 60 mm must be traversed, dropping the survival probability down to less than 1 %.

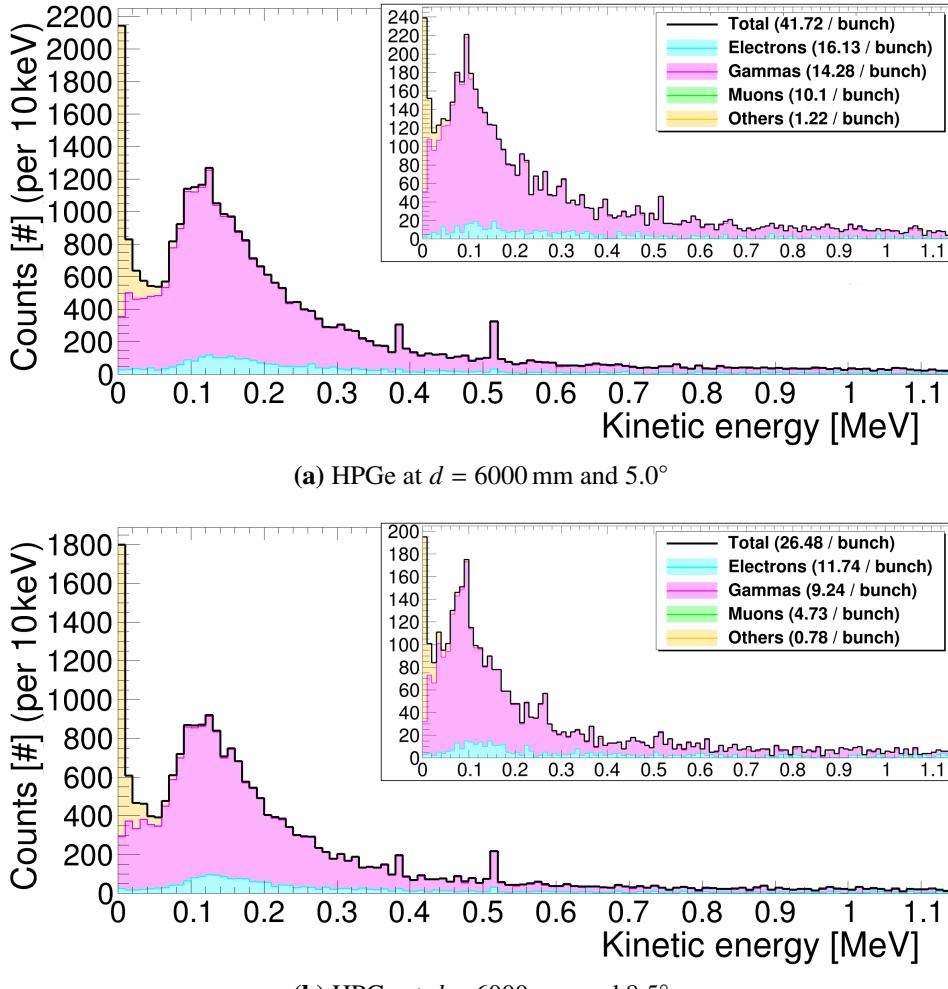


Figure 6.8.: Comparison of energy depositions inside the HPGe detector for positioning under 5.0° (a) or 8.5° (b) with respect to the muon-stopping target. The inlet assumes a perfect shielding at the HPGe sides and only accounts for particles impinging on the detector front facing the target-disks. The study was performed without a collimator. The separation into PID groups was performed at the moment of particle entry into the detector. The category “others” combines energy depositions from mainly protons, alphas, or neutrons. The data is summed over 600 bunches with 16×10^6 POT per bunch. Note that most electrons and muons deposit energies of multiple tens of MeV and thus are not visible in this figure (see Figure 6.13 for further details of both components).

The maximum HPGe angle before X-rays from the first (most US) target-disk would be blocked by the beam duct's side is 8.5° . Hence, this value was adopted for the following studies of collimator design. If current plans to replace the beam duct with a simple, thin wall of steel at the end of CyDet are implemented, the maximum HPGe angle could be increased up to 11.5° – at which point the radius of the hole in the surrounding solenoid will represent the new limiting factor. Note that due to the slanted wall of the COMET hall (see Figure 6.1), a larger angle would also allow for a slightly longer HPGe-to-muon-target distance. Given the manageable impact on full-energy peak efficiency and expected X-ray count rate, implementing such changes in the future is strongly recommended to further relieve the challenging background situation.

The HPGe background without a collimator at a distance of 6 m and an angle of 8.5° is shown in Figure 6.8b. Repeating the conclusion of Figure 6.2, the number of particles depositing energy in the HPGe's front must be reduced below one to make an X-ray measurement possible. The complexity of the collimator design arises from the extensive area of all target-disks from which X-rays can be emitted. If a collimator accounts for all target-disks positions, its opening angle is too large to reduce the background sufficiently. Therefore, a compromise between background reduction and signal acceptance is necessary.

To determine the optimal collimator, its length was increased in 100 mm steps from 100 mm to 3100 mm. Furthermore, various collimator configurations were investigated, as presented in Figure 6.9. Each configuration blocks X-rays and background particles from different areas of the target-disks, as illustrated by the accompanying acceptance profiles in the same figure. This effect is particularly noticeable with longer collimators, where accepted emission positions and angles become more restricted. In each successive (more DS positioned) target-disk, the number of stopped muons decreases by approximately 1 %. However, this difference of around 15 % between the most US and DS target-disk is partly compensated by the improved solid angle resulting from a shorter distance to the HPGe detector. Consequently, no preference was given to specific target-disks, and each collimator's central axis aimed toward the middle of the muon-stopping target, maximizing the total accepted area. An example of how acceptance profiles change between the most US, middle, and most DS target-disks is shown in Figure 6.11.

The results of all investigated collimators are summarized in three plots, demonstrated with Figure 6.10: The first illustrates potential X-ray count rate as calculated with Formula 6.8 but assuming $P(\text{free-bunch}) = 100\%$. The second shows the percentage of background-free bunches, i.e. bunches without background hits during which an X-ray could be measured. Lastly, both parameters are multiplied to yield the performance of the collimator configuration, which can be understood as the measurable number of full-energy depositions of the X-ray of interest. This procedure is demonstrated for tube-shape collimators with varying diameters in Figure 6.10, and analogous figures for other collimator configurations are provided in Appendix A.

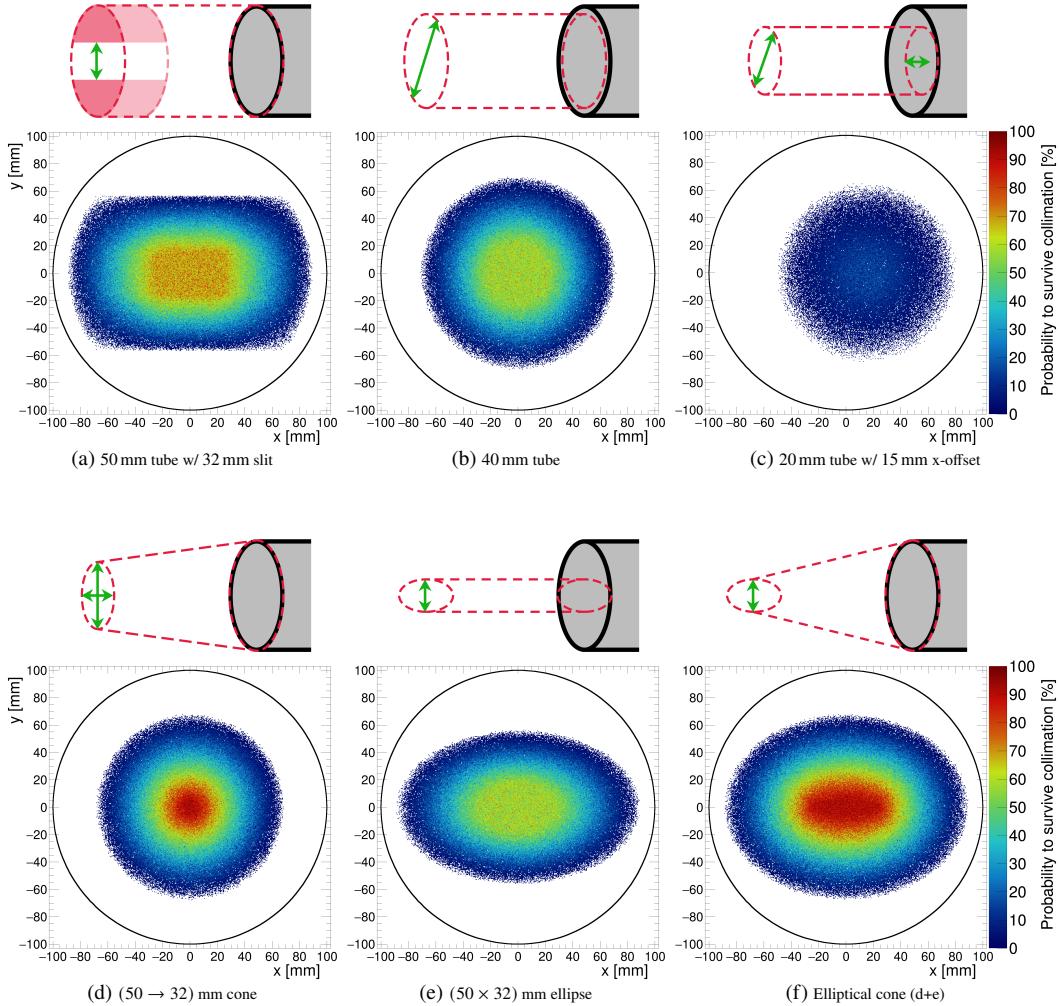


Figure 6.9.: Schematics of the investigated collimator configurations (red dashed lines). The size of the HPGe crystal is indicated as a grey tube on the right. Adjustable variables are indicated with green arrows. The color-coded acceptance profiles illustrate the probability for an X-ray to survive the collimation if emitted in the direction of the HPGe from a respective target-disks position. All profiles are given for the middle target-disk (number 9 of 17). Profiles appear symmetrical because the middle target-disk is on-axis of the HPGe. Comparable profiles for off-axis target-disks are shown in Figure 6.11. A black ring indicates the size of the muon-stopping target (disks). White areas are completely blocked by the collimator and correspond to probability zero. The collimator length of the narrow tube (c) was 1500 mm, all others correspond to $L = 2500$ mm.

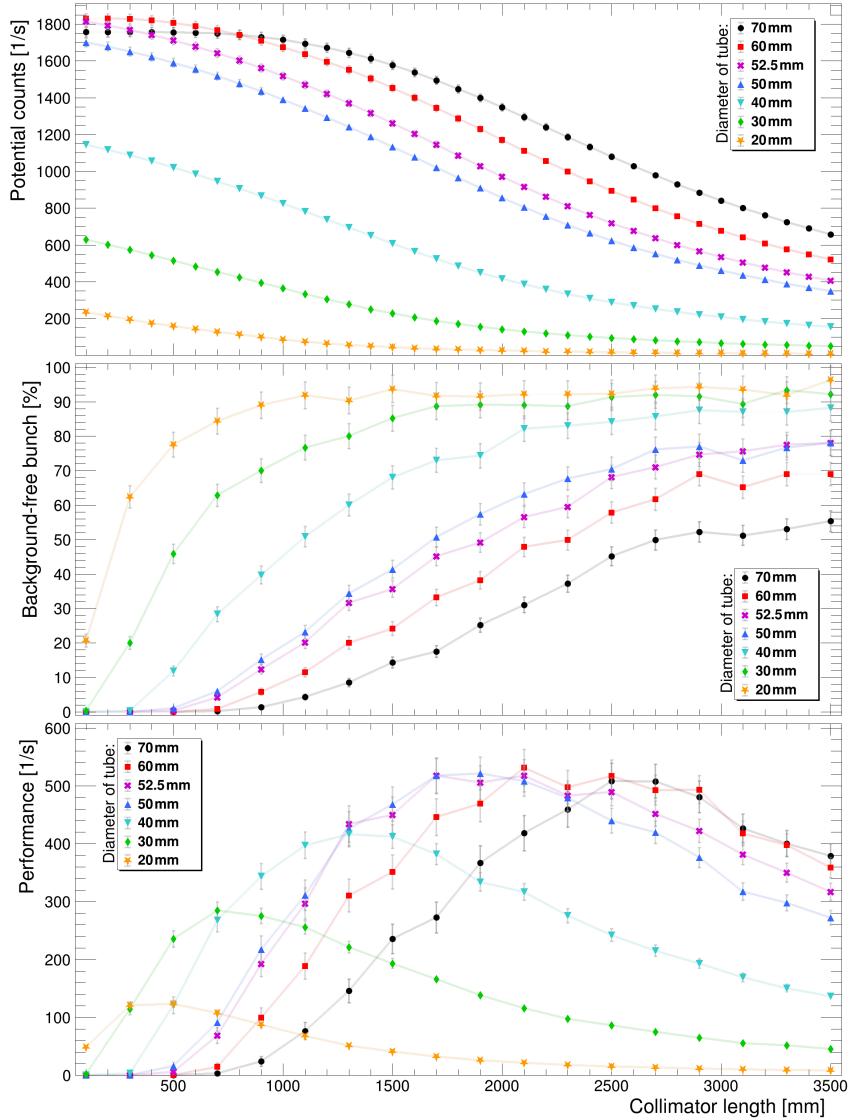


Figure 6.10.: Summary for the tube collimator, depicting the potential X-ray count rate (top), the percentage of background-free bunches (middle), and the product of both given as a performance value (bottom). Colors correspond to different tube diameters. Values were obtained according to Section 6.1. Straight lines connecting the data points are included to increase clarity. Error bars include statistical uncertainties and an assumed 1 % systematic uncertainty on the number of muons stopped per bunch and the full-energy peak efficiency. See text for discussion.

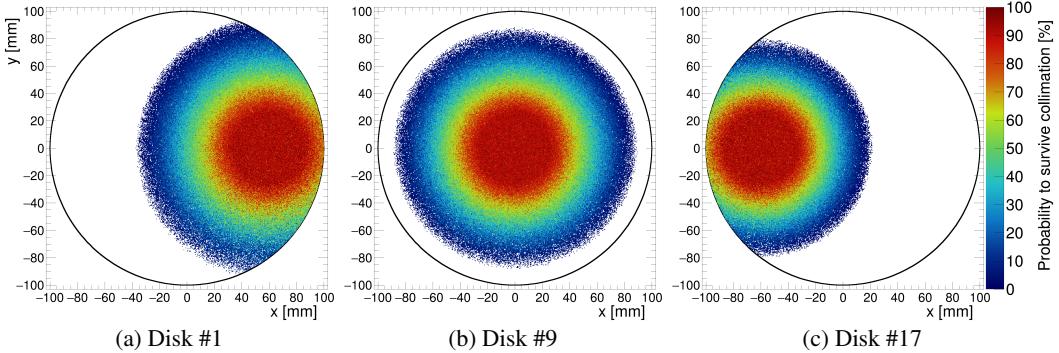


Figure 6.11.: Collimator acceptance profiles for the first (a), middle (b), and last (c) target-disk. The chosen collimator corresponds to a tube with a diameter of 50 mm and a length of 2500 mm. Because the collimator's central axis is aimed toward the muon-stopping target's center, the profile is perfectly symmetrical and centered on the target-disks at this position. However, with the HPGe positioned left relative to the muon-stopping target, the acceptance profile shifts right or left for further US (a) or DS (b) target-disks, respectively.

The key finding from Figure 6.10 is the confirmation of the HPGe detector's suitability for the X-ray measurement in COMET. Specifically, even a simple tube-shaped collimator can result in over 60 % background-free bunches. Secondly, the figure shows that a large collimator diameter up to the crystal's size yields the maximum performance. Despite smaller collimators achieving better background suppression at much shorter lengths, their maximum possible performance is reduced due to limitations in accepted emission angles. Small diameters also result in a lower average intrinsic efficiency by blocking parts of the high-efficiency zones at the HPGe's front. Hence, even if an X-ray survives collimation, it has a reduced chance of a full-energy deposition. In the case of a small 20 mm diameter, offsetting the collimator at the HPGe's front can partially recover this loss in intrinsic efficiency but still results in comparatively poorer performance. Increasing the tube's diameter beyond the crystal size yields a higher possible count rate but reduces the probability of no overlapping background energy depositions. Both effects cancel each other out, thus maximum performance cannot be further improved. In summary, the optimum collimator tube diameter is found to be 50 mm. This size allows for large X-ray emission angles, preserves the sensitive zone around the HPGe center, and only cuts negligible low intrinsic efficiency edges of the HPGe crystal. It shows the highest possible performance (compared to smaller tubes) at the shortest necessary collimator length (compared to larger tubes).

As shown in Figure 6.9, more complex collimator shapes like cones, ellipses, elliptical cones, and tubes with blockers at the top and bottom were investigated. Summary pictures are provided in Appendix A. As a consequence of the tube results, the collimator diameter

close to the HPGe was maintained at 50 mm if possible. Unsurprisingly, a smaller opening at the collimator's front, facing the muon-stopping target, always reduces the potential X-ray count rate. These differences from varying opening areas are most noticeable for short collimators. However, similar to the results of the tube, no distinct features emerge when increasing the collimator length. Although with varying gradients, the potential count rate simply decreases and eventually converges toward an offset value determined solely by the front-facing opening size. This opening size is also the dominant factor for how sharply the percentage of background-free bunches grows with increasing collimator length. While exact numbers vary among configurations, the conclusion remains unchanged: short collimators favor a more restricted opening area. However, if the collimator length must not be constrained, the maximum performance is always archived with a larger opening.

A summary and comparison of select investigated collimator designs is presented in Figure 6.12. Rather than focusing on collimator length, the figure shows a collimator's performance in relation to its level of background-free bunches. While optimal performance is favored, the importance of achieving a high probability for background-free bunches should not be underestimated. A minor loss in performance may be warranted if it results in a noticeable increase in background-free bunches, as this would reduce the overall rate of energy depositions in the HPGe, which in turn lowers crystal damage and data rate. Moreover, the previous performance definition accounted for background-free bunches linearly. However, in scenarios with prolonged detector dead time (see Section 6.3.2), the need for two consecutive background-free bunches may arise, which would translate to a quadratic contribution of background-free bunches to the performance.

Surprisingly, the plain tube-shaped collimator consistently outperforms all other designs, yielding the best combinations of performance and percentage of background-free bunches. Its only drawback is its comparatively long length, around 1.9 m corresponding to the configuration with the best performance of the 50 mm tube. In comparison, the best performance of the (50 → 40) mm cone collimator or the tube collimator with a 40 mm slit is achieved at lengths of around 1.5 m and 1.7 m, respectively. However, considering the increased machining difficulty of these designs and the available space between the HPGe and the end of CyDet (as shown in Figure 6.1), the conclusion of this study is evident: a tube with a 50 mm diameter close to the HPGe crystal's size, should be utilized. While the maximum signal count rate is obtained at an approximate length of 1.9 m, the maximum of 75 % background-free bunches is achieved at approximately 2.7 m. If stronger background suppression is required, the front-facing opening size of the collimator must be further reduced. However, differences between investigated designs become negligible at very high background suppression. Therefore, the addition of height blockers inside a 50 mm tube forming the slit-type collimator represents the most straightforward solution. As an example, by using a 2.7 m long tube and restricting the opening to a 20 mm slit, it becomes possible to achieve close to 90 % background-free bunches.

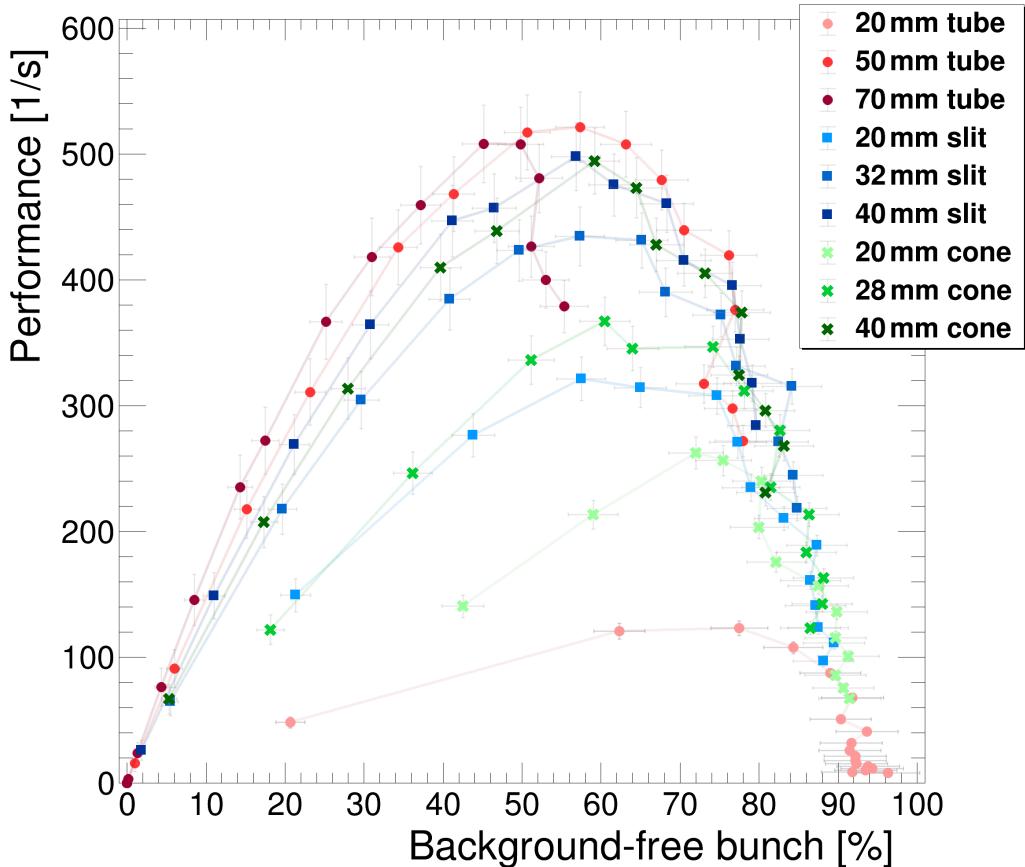


Figure 6.12.: Comparison between select collimator designs, with definitions for cone, slit, and tube shown in Figure 6.9. Values were obtained according to Section 6.1. Each point corresponds to a set collimator length, generally becoming longer from left to right with an increasing percentage of free bunches. Exact lengths can be obtained from comparison with Figure 6.10 or figures in Appendix A. Straight lines connecting the data points are included to increase clarity. Error bars include statistical uncertainties and an assumed 1 % systematic uncertainty on the number of muons stopped per bunch and the full-energy peak efficiency. See text for discussion.

6.3. Discussion and Outlook

6.3.1. Moving Toward a Realistic Shielding

The values discussed in the previous section are based on the assumption of perfect shielding and collimation. However, in realistic scenarios, particles from the HPGe sides or back may penetrate such shielding and deposit energy within the detector. Details heavily depend on the exact particle type and energy or the chosen material and thickness of the shielding. As a starting point for this discussion, Figure 6.13 shows the expected particle fluence impinging on the collimator as well as the HPGe side, back, and front.

In all panels of Figure 6.13, the most significant contribution up to approximately 10 MeV/c stems from photons. These photons mainly originate from Bremsstrahlung or deexcitation after ionization of surrounding material but are easily shielded by materials with a high atomic number. A common practice, not limited to HPGe detectors, is to enclose the active detector components with lead. Using the Bouguer-Lambert law (see Formula 6.6), one can calculate that just 10 cm of lead is adequate to shield 99 % of all impinging photons⁶. Typically, a graded shielding approach is advised to suppress characteristic X-rays from lead excitations. Lead X-rays fall within the energy range of 75 to 80 keV, and a thin layer of copper is sufficient for their suppression. Therefore, an initial design to mitigate the intense photonic background could consist of a 1 cm pure copper layer surrounded by a 10 cm thick lead layer.

Above 10 MeV/c, electrons represent the most prominent background in Figure 6.13. Unlike heavier particles, however, electrons at such energies not only lose energy by ionization after collisions but also by emitting Bremsstrahlung after scattering in the electric field of a nucleus [99]. Which of both interactions dominates strongly depends on the mass number Z of the shielding material. It is common to define a critical energy E_c for which both loss rates are equal, which can be approximated with:

$$E_c = \frac{800 \text{ MeV}}{Z + 1.2}. \quad (6.9)$$

For lead one quickly finds $E_c \simeq 9.6 \text{ MeV/c}$. Thus, even though the thickness of the above-proposed lead shielding is enough to suppress the electron background⁷, such shielding would produce significant amounts of high-energy bremsstrahlung. Generally, achieving a high percentage of (electron) shielding with high- Z materials requires a disproportionately large thickness to attenuate this bremsstrahlung component. Therefore, a layer of lighter,

⁶The mass attenuation coefficient of lead is at a minimum for 5 MeV photons and was taken from [98] as $(\mu_k/\rho_k) = 4.27 \times 10^{-2} \text{ cm}^2 \text{ g}^{-1}$. The lead density was taken as $\rho_k = 11.34 \text{ g cm}^{-3}$.

⁷The traveled distance of electrons in a material can be estimated using the Continuous Slowing Down Approximation (CSDA). Values for various materials are conveniently provided at [100]. For example, a 50 MeV/c electron in lead has a CSDA range of 14.6 g cm^{-2} , hence travels approximately 1.3 cm. Note that electrons undergo multiple scattering and their resulting vertical projected range is often significantly shorter than their CSDA distance.

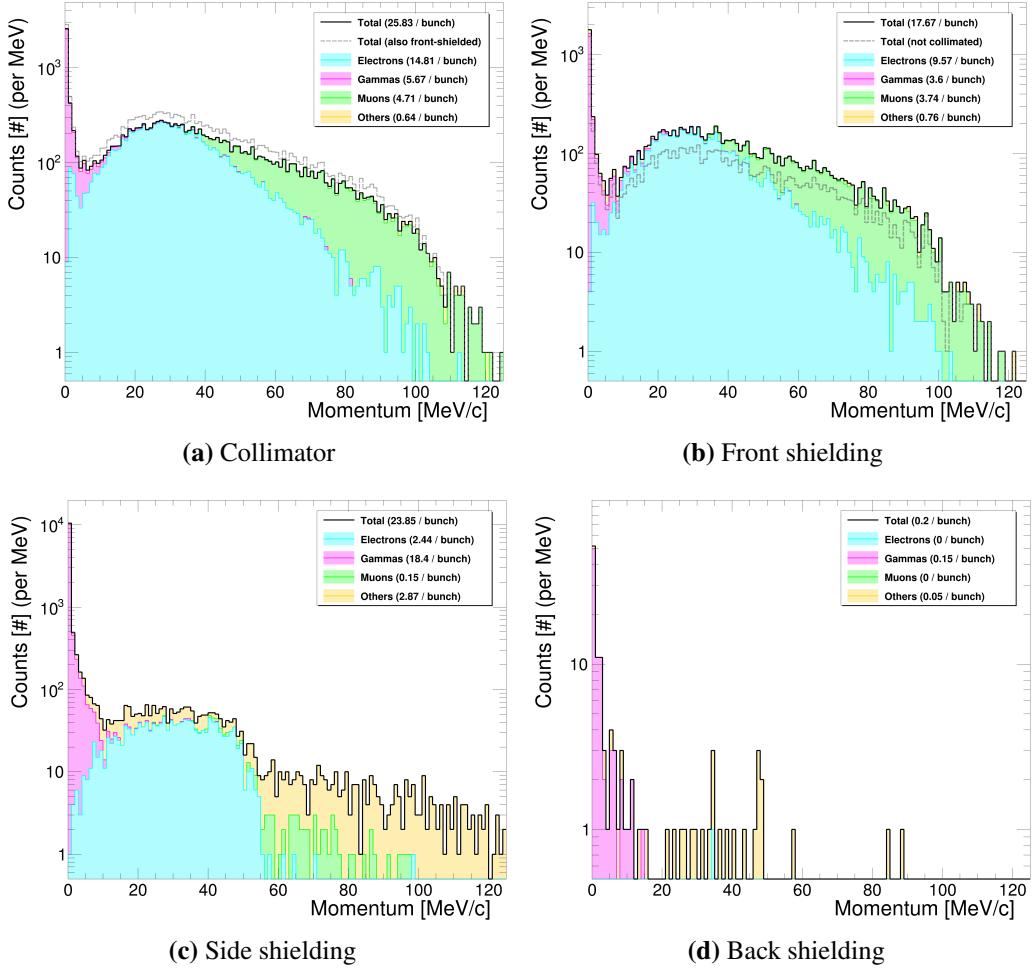


Figure 6.13.: Momentum distribution of particles impinging on the shielding or the collimator, only accounting for particles which would deposit energy inside the HPGe. The data is summed over 600 bunches with 16×10^6 POT. The HPGe was positioned 6 m DS at an angle of 8.5° relative to the muon-stopping target's center. The collimator was a 2.5 m long and 50 mm diameter tube. Particles intersecting both the collimator (a) and the front shielding (b) are counted toward the latter. The gray dashed line in plot (a) indicates the fluence if these particles were counted toward the collimator. On the contrary, the grey dashed line in plot (b) corresponds to particles that do not traverse the collimator but only the front shielding.

cheaper, and easier-to-handle low- Z material like 10 to 20 cm polystyrene⁸ should be used as an outer layer.

The collimator and HPGe front will be hit by many high-momentum muons between 35 MeV/c and 120 MeV/c, which were not stopped within the target-disks. The mean rate of energy loss of these muons is well described by the so-called Bethe-Bloch formula:

$$\left\langle -\frac{dE}{dx} \right\rangle = K \cdot \frac{Z}{A} \rho \cdot \frac{z^2}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 (\beta\gamma)^2 W_{\max}}{I^2} - \beta^2 \right] = \frac{Z}{A} \rho \cdot f(\beta, I). \quad (6.10)$$

Besides a variety of constants combined into K , it depends on the incident particle's charge z , its velocity expressed in terms of $\beta = v/c$ and $\gamma = 1/\sqrt{1 - \beta^2}$, and material constants like density ρ , mass number A , atomic number Z , and mean excitation potential I . The formula only accounts for collision-ionization. However, the cross-section for Bremsstrahlung is approximately 40 000 times smaller for muons than for electrons and becomes only relevant around a few hundred GeV as it depends on the inverse square of a particle's mass. The energy loss rate reaches a broad minimum around $\beta\gamma = p/m_\mu c = 3.0 - 3.5$, at which shielding becomes less effective. However, this value corresponds to a muon momentum of around 300 MeV/c and is above the expected muon background. Toward lower momenta, a steep rise in stopping power can be observed. Integration of Formula 6.10 yields the CSDA range. For a 100 MeV/c muon in lead one finds $R/M \approx 180 \text{ g cm}^{-2} \text{ GeV}^{-1}$ [101], which corresponds to a range of only 19 g cm^{-2} or 1.7 cm^{\dagger} . Consequently, the already proposed lead shielding for photons is also sufficient to stop the impinging muons.

The remaining background in Figure 6.13 is summarized as “others” and mainly consists of neutrons, protons, or alpha particles. Energy loss of the latter two can be described in a similar fashion to muons with Formula 6.10. Hence, both particles are sufficiently shielded with the proposed lead layer. However, neutrons require special treatment due to their neutral nature. Fast neutrons mainly lose energy through elastic and inelastic scattering until they are thermalized and captured by a nucleus. The exact cross-section strongly depends on the shielding material and its exact nuclear structure⁹. However, neutrons are generally very penetrating as their capture cross-section is inversely proportional to the neutron's velocity.

A significant number of neutrons is produced at the initial proton-proton collision at the pion-production target but also further DS through processes like muon nuclear capture. Preliminary estimates show that the unshielded HPGe would be exposed to a 1-MeV-neutron equivalent fluence of up to $10^{12} \text{ n}_{\text{eq}}/\text{cm}^2$ over a runtime of 150 days. However, significant damage can be expected after exposure to more than $10^9 \text{ n}_{\text{eq}}/\text{cm}^2$ [104]. Typical neutron

⁸Polystyrene is a synthetic polymer made from hydrocarbons. Because of its low mass number, energy loss of a 50 MeV/c electron through ionization is twice as likely as through radiation. Its CSDA range, taken from [100], is approximately 21 g cm^{-2} .

[†]The estimation assumed $m_\mu = 105.7 \text{ MeV/c}^2$ as muon mass and $\rho_{\text{lead}} = 11.34 \text{ g cm}^{-3}$ as lead density.

⁹Data libraries with tabulated neutron cross-sections exist and can be conveniently accessed via web-based interfaces at [102] or [103].

6. The Muon-Stopping Monitor

shielding consists of hydrogenous materials for efficient neutron slowdown combined with an element exhibiting a high thermal neutron capture cross-section, like borated polyethylene [105]. The material should be placed outside the lead layer in order to absorb photons stemming from the neutron capture reactions. At the time of writing, a collaboration-wide effort was underway to reduce the overall neutron fluence in the CyDet hall. Following this investigation, the thickness of a HPGe specific shielding must be adjusted to lower neutron fluence below the $10^9 \text{ n}_{\text{eq}}/\text{cm}^2$ critical value.

The schematic drawing in Figure 6.14 summarizes the discussion. However, a more detailed ICEDUST study should still be performed in the future to determine the required material thicknesses more accurately. Note that the polystyrene layer might not be necessary, as the lightweight borated polyethylene used for neutron shielding may already sufficiently attenuate the electron background without significantly increasing Bremsstrahlung. However, polystyrene is inexpensive, easy to handle, and can help to further thermalize the neutron background, hence it is still included in the figure. The above discussion only yet accounts for the shielding at the HPGe sides (and front/back). The required materials for the collimator are much harder to assess. The proposed layered design of the shielding is too bulky to employ as a roughly 2 m long tube. Additionally, particles will not traverse the collimator surface perpendicular and could easily miss layers. However, the shallow angle of particles impinging on the collimator also translates to much longer traversed material within significantly thinner shielding. Hence, a preliminary design could be a layer of lightweight polystyrene wrapped around a steel(pipe), which is easy to machine and has a medium mass number. If necessary, thin disks within the collimator to absorb low-energetic gammas or a dipole magnet to sweep away electrons could be additionally employed to reduce background. Details should be optimized in future MC campaigns.

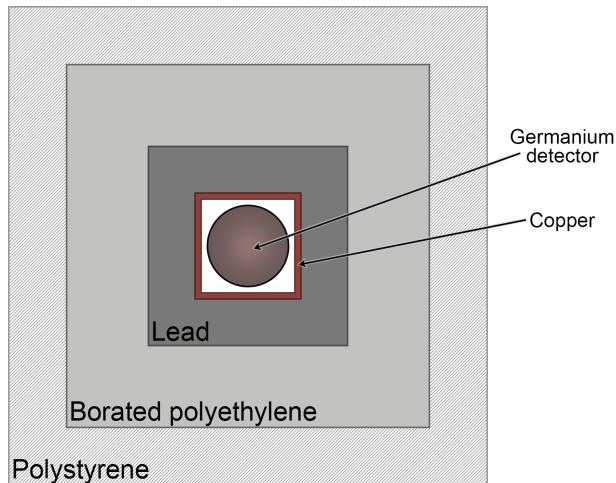


Figure 6.14.: Conceptual shielding design for COMET's HPGe detector.

6.3.2. HPGe Data Acquisition: Read-out and Dead Time

The output of the HPGe is essentially an electrical charge proportional to the amount of energy deposited within the detector. Therefore, a conventional DAQ system consists of the following three parts, which are quickly be introduced in the following:

- A preamplifier to collect charge carriers within the HPGe
- An amplifier to process the height information from the preamplifier pulse by amplification and pulse shaping
- An Analog to Digital Converter (ADC) to measure the height of the shaped pulse

The preamplifier The preamplifier serves as the direct interface to the HPGe crystal. Despite its name, it does not amplify the pulse but rather collects the charge generated within the germanium crystal. In modern HPGe systems, the preamplifier is closely attached to the HPGe crystal (capsule) and an integrated part of the purchased detector setup. This close attachment allows for efficient (and passive) cooling of this critical component to reduce thermal noise within its electronics.

Two common types of preamplifier systems exist: Resistive Feedback Preamplifiers (RFPs) and Transistor Reset Preamplifiers (TRPs). Schematic examples of pulse shapes for both types are shown in Figure 6.15. A shared characteristic is the steep rising edge, with rise times typically in the range of 100 to 700 ns depending on the amount of collected charge or the size of the HPGe crystal. In RFP systems, this rising edge is followed by a long pulse decay (roughly 50 μ s) to restore the initial baseline height. In contrast, the voltage level in TRP systems will stepwise increase for each incoming pulse and ‘reset’ back to baseline in a single step after a predetermined threshold is passed.

RFP systems are not recommended for high count rate environments¹⁰. If the time difference between pulses becomes too short for the RFP circuit to restore its baseline, the detector essentially starts to ‘lock up’ (see example in Figure 6.15). Such a lock-up cannot occur with TRP systems, and therefore it is the natural choice for COMET. However, the resets of a TRP can significantly contribute to the systems’s dead time – an effect introduced and discussed in detail further below.

The amplifier and ADC Traditionally, the sharp peak shape from the preamplifier was not suitable to accurately determine the pulse height. Instead, a set of high-pass (differentiator) and low-pass (integrator) filters combined in an amplifier were used to convert the preamplifier pulse to a semi-Gaussian shape that is much easier to handle for the ADC [106]. However, besides issues like pole-zero correction or proper pile-up rejection, these systems are relatively slow. Depending on the desired energy resolution, the processing of a single pulse (or energy deposition) can take from 3 μ s up to 60 μ s – much too long compared to the 1.17 μ s bunch separation of COMET.

¹⁰High count rate refers to everything above $100\,000\,\text{s}^{-1}$. While the limit of an RFP system depends on the size of the feedback resistor, for 1 MeV photons (or energy depositions), Mirion (former Canberra) states a maximum rate of around $2 \times 10^6\,\text{s}^{-1}$ with their fastest RFP.

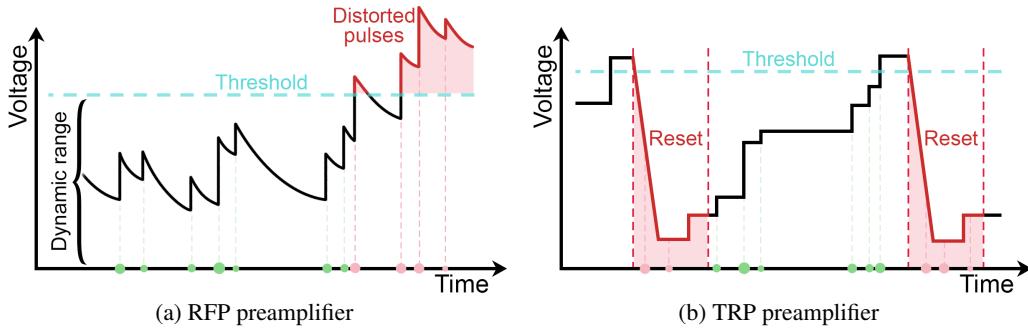


Figure 6.15.: Example pulse shapes of a RFP and TRP preamplifier systems.

Energy depositions in the crystal are shown as dots on the time axis. The dot color indicates if the system is ready to process the signal (green) or if the energy deposition would be lost (red). The dot size corresponds to the amount of deposited energy. In both systems, it is proportional to the step height of the pulse. If count rates are too high, the RFP system becomes too slow to restore voltage levels to its baseline. Above its dynamic range threshold, linearity between step height and energy deposition is impaired, yielding to distorted pulses. The transistor within the RFP preamplifier circuit saturates for even higher count rates, essentially ‘locking up’ the detector. Such problems do not exist for a TRP preamplifier system. However, its periodic reset causes a unavoidable dead time of the system.

Thankfully, modern flash ADCs can easily achieve sampling rates of $100 \text{ MHz}^{\ddagger}$, translating to one pulse height measurement every $0.01 \mu\text{s}$. This enables the direct digitization of the preamplifier’s rising edge, eliminating the need for analog signal shaping using the before-described amplifier system. Instead, pulse shaping is performed digitally and often integrated into the same device as the ADC. Unbound from the restrictions of analog electronics, these digitizer systems can employ an improved trapezoidal pulse filter. If one accepts a compromise on energy resolution, trapezoidal shaping times of around $2 \mu\text{s}$ seem possible¹¹ with commercially available devices from CAEN or ORTEC [107]. Note that such systems often come equipped with some buffer, which decouples incoming pulses from shaping and processing. While these buffers may allow for faster pulse processing than the $2 \mu\text{s}$ shaping time (over short periods), the specifics heavily depend on hardware, and buffers are therefore disregarded for simplicity throughout this discussion.

[‡]At the time of writing, standard CAEN digitizer systems are available with up to 500 MHz.

¹¹Much to the regret of the author, no functioning HPGe system with TRP preamplifier was available at the time of writing. Hence, the values provided in this section still require proper testing and verification. Especially the exact quantification of loss in the HPGe’s energy resolution for such fast shaping times and in an ultra-high rate environment should be extensively studied in the future to find a set of optimum parameters for COMET’s condition.

Detector Dead Time Given a $2\text{ }\mu\text{s}$ DAQ chain and assuming ideal conditions with no background interference, the HPGe detector could process a signal every second bunch. However, signals of in-between bunches cannot be detected and processed. This period, during which the detector cannot register any energy depositions, is defined as its dead time. For COMET’s HPGe detector, the dead time will be primarily characterized by the electronic recovery time. Naturally, it also directly depends on the expected rate of events (background and X-rays) generating signals within the detector. Additionally, the resetting mechanisms of the TRP preamplifier will significantly contribute¹².

Dead time values can be directly obtained from the software accompanying a HPGe’s digitizer systems. However, at best, such values are mere estimates that depend on underlying assumptions. In particular, events are considered independent and randomly distributed in time [108]. This assumption is not applicable for COMET’s conditions, as the X-rays follow the external bunch timing instead of being Poisson distributed. Moreover, a general dead time spanning the entire measurement time holds little significance. Instead, the parameter of interest is the dead time around the time X-rays are emitted, characterized in the previous section as the probability for a background-free bunch.

More precise estimates for dead time losses accounting for pulsed environments can be found in [109] or the more recent publication [110]. However, even these sophisticated calculations depend on various assumptions about exact pulse and dead time durations. Additionally, apart from the dead time caused by the DAQ chain, COMET’s high background environment will lead to significant signal counting rate losses due to pileup with background events. Given the dead time’s linear impact on the estimated number of muons stopped, an external measurement of these effects is necessary to accurately assess (and later correct for) the fraction of true versus measured X-ray events.

A detailed discussion of several reliable and well-understood methods for live signal-loss measurements are given in [106]. However, methods relying on reference count rates of one or two external sources are not applicable to COMET’s pulsed-beam situation. Hence, the author recommends using the *pulser method*. The method is well discussed in literature [111, 112]. In essence, external pulses from a pulse generator are periodically injected into the HPGe preamplifier. Identical to real detector signals, this pulser signal is processed by the full DAQ chain and may be affected by pileup due to parallel occurring energy depositions or be lost because of dead time if the DAQ is already processing another event. The signal counting rate loss is estimated offline by comparing the number of counts within the measured spectrum to the known¹³ number of injected pulses.

The method was employed successfully in other experiments and proven to yield the loss on signal count rate with a relative systematic error below 1 % [113]. For optimal results,

¹²The rate of TRP resets depends on factors such as the number and size of energy depositions, along with the capacitance of the built-in capacitor. Typical reset rates, in the absence of sources near the HPGe, fall within the range of 1 to 10 Hz [93].

¹³For a direct measurement of the injected number of pulses, a T-splitter can be used to feed the same pulser signal into a separate unused channel of the digitizer.

the pulser's signal shape should be configured to closely resemble an energy deposition within the HPGe. The signal height (equal to the energy) should be chosen reasonably large to clearly separate pulser and real measurement events. For COMET, the pulse generator should be connected to an external clock to emit pulses in synchronization with possible X-ray emission, i.e., when muons reach the muon-stopping target. The frequency of pulses must be chosen reasonably low to avoid a significant impact on dead time by the pulser itself. However, given COMET's long measurement time, rates as low as 10 ms (every $\sim 10\,000$ bunches) will still produce a reasonable pulser statistic.

6.3.3. Estimation of Accuracy

The main aim of the HPGe measurement is to determine the number of stopped muons in the muon-stopping target. The accuracy of this measurement is crucial, as the number of stopped muons contributes linearly to the SES. In this section, an estimate of the uncertainty is determined by revisiting Formula 6.1, where three factors contribute:

$$N_{\mu}^{\text{stopped}} = \frac{F_{i \rightarrow f}}{\bar{\epsilon}_T \cdot P_{i \rightarrow f}}, \quad (6.11)$$

First, as taken from the literature, the muonic X-ray line transition probabilities $P_{i \rightarrow f}$. The relative uncertainty of $P_{2 \rightarrow 1}$ is given as less than 1 % (compare Table 6.1). Second, the uncertainty on the number of measured events within the full-energy X-ray peak $F_{i \rightarrow f}$. For one, this uncertainty depends on the signal-to-background ratio and the total number of counts in the final spectrum. However, as seen in Figure 6.10, the theoretical maximum count rate equals 500 s^{-1} . Hence, even if 90 % of the signal were lost due to pileup or dead time, a clear high-statistic peak should be visible above the otherwise flat background at the 347 keV region. For another, the dead time and pileup correction, which contributes linearly to $F_{i \rightarrow f}$. The pulser method discussed in the section above was shown to yield an accuracy within the order of $\mathcal{O}(1\%)$ [112, 113], even for pileup as high as 30 %. However, tests performed at TU Dresden show increasing uncertainties toward higher count rates [114]. As COMET could have a combined dead time and pileup correction of 60 % or more, a future in-depth study of this method is advised. Lastly, the average full-energy peak efficiency $\bar{\epsilon}_T$. Its uncertainty is a combination of inaccuracies in target-disk positioning, HPGe positioning and angle, collimator accuracy, and variations in the muon beam, including profile and momentum distribution. Values were estimated by individually adjusting parameters within the custom MC described in Section 6.1.

Simplified estimates of uncertainties on the HPGe measurement are provided in Table 6.3. For the stated assumptions, the $^{+0.6}_{-24.4}\text{ %}$ relative uncertainty on the average full-energy peak efficiency dominates the total error budget. The uncertainty is strongly asymmetric because an error-free alignment of the collimator resembled its optimum. Hence, variations in the collimator's tilt angle or x - y -placement in front of the HPGe always cause a loss in $\bar{\epsilon}_T$.

Table 6.3.: Contributions to the overall uncertainty on the number of muons stopped. Values shown correspond to the $2p \rightarrow 1s$ transition at an X-ray energy of $E = 346.8$ keV. The HPGe was positioned 6 m DS at an angle of 8.5° relative to the muon-stopping target's center. The collimator was a 2 m long and 50 mm diameter tube.

Parameter K	$\Delta K/K [\%]$	Comments
Transition probability: $P_{2 \rightarrow 1}$	0.8	Literature value from [66]
Counts in the X-ray peak: $F_{2 \rightarrow 1}$		
Count rate	$\lesssim 1.0$	Estimate based on expected count rate
Pileup and dead time	~ 1.0	Estimate based on [112] and [113]
Total	1.4	Dominated by dead time correction
Avg. full-energy peak efficiency: $\bar{\epsilon}_T$		
Target-disk positioning	± 0.8	Target-disks displaced by up to ± 5 mm
HPGe distance	∓ 0.4	Varried by ± 25 mm
HPGe position angle	∓ 0.7	Varried by $\pm 0.3^\circ$
Collimator radius	± 5.5	Varried by ± 0.5 mm
Collimator length	∓ 0.6	Varried by ± 10 mm
Collimator tilt angle	$\{ \mp 5.4$ $-9.5 \}$	Varried by $\pm 0.3^\circ$
Collimator x-position	-1.6	Varried by ± 2 mm
Collimator y-position	$\{ \mp 0.9$ $-2.3 \}$	Varried by ± 2 mm
Muon beam profile distribution ^a	± 13.9	
Muon beam momentum distribution ^b	± 0.3	
Total^c	$\{ \pm 0.6$ $-24.4 \}$	

^a The x - y -stopping profiles used as input for X-ray starting positions were changed to a 2nd order 2D-polynomial, a Gaussian, and a homogenous distribution. Stated uncertainty corresponds to the maximum deviation. ϵ_T increases (decreases) for more centrally confined (broader spread out) distributions due to a focusing-like effect of the collimator.

^b The z -stopping profile used as input for X-ray starting positions was changed by adjusting the number of muons stopped inside individual target-disks. While the z -stopping profile normally peaks around the second and drops by approximately 1 % per subsequent target-disks, this peak was moved further DS (more high-energetic muons) or made more pronounced (more low-energetic muons).

^c To account for correlations between contributing parameters, they were uniformly varied within their stated uncertainties over a sample size of 750. The uncertainty on ϵ_T is given with respect to the calculated true value of 8.94×10^{-7} and corresponds to the 95 % confidence interval around the distribution's median at 8.62×10^{-7} .

Note that variations in each component of the efficiency were chosen quite conservatively to assess their individual severity. The most significant factors are the collimator radius and its tilt angle, as well as the profile of muons stopped within target-disks. The first, collimator radius, can be expected to be significantly more accurate than 0.5 mm. Assuming a more realistic uncertainty on the tube-radius of ± 0.1 mm lowers its contributing relative uncertainty to mere 1.1 %. The latter two highlight that the efficiency calculation is most sensitive to the area of muon stops (equal to X-ray start positions) on which the collimator focuses its view. Fine adjustments of the collimator's tilt angle might prove challenging because the muon-stopping target will already be secured inside CyDet at a 6 m distance at the time of the HPGe installation. The assumed 0.3° variation amounts to an offset of only 3 cm at the target-disks. A dedicated alignment scheme using theodolites should be employed in order to achieve an accuracy of less than $\pm 0.1^\circ$ (± 1 cm), which could lower this relative uncertainty contribution to ${}^{+0.2}_{-1.4}\%$. Regarding the influence of the muon beam (or stopping) profile, the assumed variations from the ICEDUST profile to a Gaussian or 2nd order polynomial were chosen rather extreme. The profile is expected to be known to much higher precision during data taking. This precision will be obtained through a separate beam measurement using the COMET's StrECAL system (see Section 2.1), which is further motivated by this result.

In summary, the exact placement of the HPGe detector, its shielding, and especially its collimator is paramount to control systematic uncertainties that otherwise will dominate the result on the number of muons stopped. However, with proper alignment of all parts, achieving a relative uncertainty around an envisioned design value of 5 % seems feasible. A significant factor will be the accuracy of the muon beam measurement using COMET's StrECAL system. Note that the discussion assumed an uncertainty from the proposed dead time and pileup correction of only 1 %. While this result was shown achievable in other experiments, a proper investigation of the performance under COMET's high background situation is advised.

6.3.4. COMETs Timing Structure and Supplementary Measurements

A large fraction of the HPGe background surviving shielding and collimator is Bremsstrahlung caused by beam electrons. These electrons are almost massless compared to muons, so their transport through the 90° beamline is slightly faster. Studies on MC truth data reveal that electron-caused background in the HPGe is clearly separated from muons arriving at the target-disks by approximately 100 ns. It certainly seems tempting to use this information for timing-based cuts. However, the significance of this proposal should not be overestimated as the slow electronic response of the HPGe of up to 700 ns (see Section 6.3.2) will pose a significant challenge in utilizing this timing information in the actual measurement scenario. In fact, any timing cut based on events within COMET's very fast-paced 1.17 μ s bunch-spacing will require careful study of actual measurement data to

reveal if such cuts may be applicable¹⁴.

With no realistic 1.17 μ s bunch-spaced measurement data available, further analysis of such timing-cuts was postponed. However, an additional external timing structure exists in COMET. Proton extraction from the J-PARC proton main ring occurs in 2.5 s long cycles, from which only 500 μ s are used to transfer protons to the experiment. This so-called 1/5 duty factor leads to periods of approximately 2 s in which no beam-induced background particles hit the HPGe. The absence of protons naturally means that no muons are produced and stopped within target-disks, thus no muonic X-rays can be measured. However, the measurement of decays from nuclides produced during previous muon captures is possible.

Approximately 60.9 % of stopped muons are captured by the aluminium nucleus. From these capture reactions, 9(6) % occur without neutron emission¹⁵, producing ^{27}Mg in reactions such as $^{27}\text{Al}(\mu^-, \nu)^{27}\text{Mg}$ or $^{27}\text{Al}(\mu^-, \nu\gamma)^{27}\text{Mg}$ [90]. As shown in Figure 6.16, the isotope has a half-life of around 9.5 min and decays by beta emission back into excited states of ^{27}Al . The aluminium nucleus then de-excites into its ground state by emission of 844 keV (71.8 % absolute intensity) and 1014 keV (28.2 % absolute intensity) gamma-rays [70]. The delayed nature of this gamma-ray emission allows its measurement within the proton-free and background-reduced period of an extraction cycle. However, great care must be taken to separate the yield of ^{27}Mg between muon capture and competing processes like $^{27}\text{Al}(n, p)^{27}\text{Mg}$ or $^{27}\text{Al}(\pi, \gamma)^{27}\text{Mg}$. Additionally, the uncertainty on the production rate of ^{27}Mg by muon capture will strongly dominate the final measurement result. Hence, a dedicated measurement campaign to establish the yield of both characteristic ^{27}Al gamma-rays in relation to the number of stopped muons is currently pursued by the AlCap experiment [72].

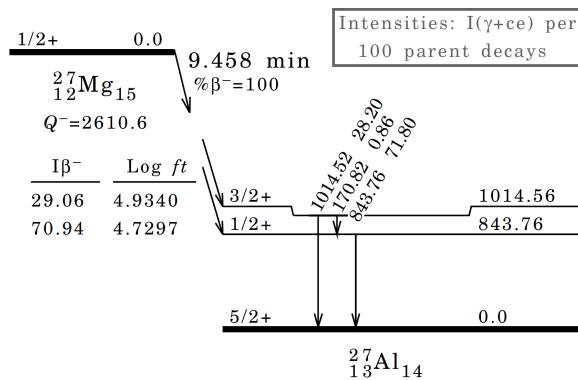


Figure 6.16.: Decay chain of ^{27}Mg , adapted from [115]

¹⁴Cuts based on timing information might help improve the signal-to-background ratio for the X-ray peaks. However, such cuts will be applied offline and cannot be used to ‘veto’ HPGe events within a bunch to relieve the challenging dead time situation.

¹⁵Emission of charged particles like proton, deuteron, or alpha is possible but has a negligible low rate compared to neutron emission, hence will be disregarded in this discussion.

6. The Muon-Stopping Monitor

The vast majority of muon captures, 75(10) %, occur accompanied by single neutron emission. Hence, another supplementary measurement could be the detection of deexcitation gamma-rays from ^{26}Mg produced in reactions like $^{27}\text{Al}(\mu^-, \nu n\gamma)^{26}\text{Mg}$. Especially the transition from the $J^\pi = 2^+$ level to the ground state of ^{26}Mg , producing a 1809 keV gamma-ray, has a sizeable yield of 51(5) % per muon capture [90]. However, the lifetime of muon capture on aluminium is in the order of $\mathcal{O}(1\text{ }\mu\text{s})$. Hence, the process competes with beam background and X-ray hits in the HPGe. Additionally, similar to ^{27}Mg , corrections accounting for competing processes producing ^{26}Mg without the involvement of muons are necessary but challenging.

In summary, analyzing measured magnesium gamma-rays within the HPGe spectrum can yield valuable additional information on the number of muons stopped inside the target-disks. Unfortunately, both discussed measurements are expected to have large uncertainties propagated from their production cross-sections or competing production processes. However, these measurements are not mutually exclusive and should be used complementary to restrain the uncertainty on the HPGe's final result: the number of muons stopped.

7. Conclusion

This thesis focused on the number of muons stopping within the detector area of the COMET experiment. Phase-I of this experiment is designed to improve the current best limit on the μ -e conversion process by two orders of magnitude, aiming at a SES of 3.1×10^{-15} . Detecting such a signal would provide evidence of CLFV and serve as a direct confirmation of new physics beyond the Standard Model. The number of available muons is crucial in achieving this sensitivity within a reasonable timeframe. Utilizing a high-intensity proton beam from the J-PARC complex in Japan, approximately 1.6×10^{17} muons are expected to reach the detector area within a runtime of 146 days. However, for conversion to occur, these muons must be brought to a stop within a muon-stopping target, placed at the center of the detector system and the muon beam.

The initial design of the muon-stopping target consisted of 17×0.2 mm thin aluminium target-disks with a radius of 100 mm, positioned along the limited 800 mm space inside the CDC detector. However, with this default configuration, only about 10 % of the arriving muons can be stopped, as its shape and amount of material are insufficient to slow down and stop the large fraction of muons with more than 50 MeV/c momentum.

State-of-the-art high statistics simulations were employed to investigate and optimize the muon-stopping target, particularly to capture these high-momentum muons. Larger target-disk radii up to 150 mm can improve the conversion signal by 40 % but also cause a steep rise in background from Bremsstrahlung, disfavoring such a design. Instead, the approach to enhance the amount of stopping material by employing more target-disks was pursued. The optimal number of 0.2 mm thick target-disks was found to be 67. Placing this much material inside the muon beam triples the number of stopped muons and improves the conversion signal by 140 %. However, it also increases CDC background to an occupancy of 70 % – far above its operational limit around 50 %. To remain below this threshold, the total aluminium thickness had to be limited to approximately 6.6 mm corresponding to 33×0.2 mm target-disks.

Further improvements were achieved by implementing thinner target-disks and positioning them closer together to enable offsets around the sides toward the CTH detector. Within the 50 % CDC occupancy limitation, the optimal configuration was found to be 66×0.1 mm target-disks with 100 mm offsets on both the US and DS sides. This configuration improves the SES by almost 50 %, reaching 1.53×10^{-15} . The result provides COMET with a substantial buffer for its physics goal or allows for a significant reduction in runtime while maintaining its original sensitivity. Alternatively, reducing the proton beam intensity to ease challenging detector backgrounds becomes feasible, as the improved muon-stopping power compensates for the loss in muons reaching the muon-stopping target.

7. Conclusion

New theoretical predictions of the DIO electron spectrum were used to update background estimates. For a momentum window of $P_e \in [103.6, 106.0] \text{ MeV}/c$ and an assumed partial branching ratio of $\mathcal{B.R.}(\mu \rightarrow e) = 3.0 \times 10^{-15}$, the fraction of DIO electrons in the signal region within the Phase-I runtime is shown to be reduced from 1×10^{-2} to 3.2×10^{-3} events. For the selected optimal configuration, the value is further reduced to 1.4×10^{-3} events. However, this study did not perform CDC track reconstruction but instead applied naive Gaussian smearing to the momentum of DIO tracks. Therefore, the results are expected to be slightly increased under actual CDC conditions. Additionally, first-time estimates of DIO background from non-aluminium elements verified that contributions from target-disk impurities or stopped muons in the surrounding helium atmosphere are completely negligible compared to the aluminium DIO rate.

Various aluminium alloys were tested for suitability as target-disk material. Aluminium “EN1050A” (99.5 % purity) with temper H19 (full-hard) was found suitable in both 0.2 mm and 0.1 mm thicknesses. A suspension structure to hold these target-disks was designed and prototyped. Three 15 mm \times 15 mm hollow square bars made from “EN AW-6060” aluminium alloy bridge the gap between the US and DS CTH. By placing all components at the outermost possible radii, around 300 mm, interaction with beam particles can be avoided, and signal acceptance loss from blocking conversion electrons is limited to 3.4 %. Target-disks are anchored to the bars using thin tungsten wire, for which a diameter of at least 0.2 mm was proven necessary to ensure stability. Installation and stability tests demonstrated a positional accuracy of $\pm 1 \text{ mm}$ in x and $\pm 5 \text{ mm}$ in y and z direction – sufficient to avoid negatively affecting the number of stopped muons.

Lastly, the feasibility of assessing the number of stopped muons by measuring characteristic X-rays with a high-purity germanium detector was successfully demonstrated. The detector should be positioned at a large distance of around 6 m and under an angle of 8.5° relative to the center of the muon-stopping target. Collimation using a plain 50 mm diameter and 1.9 m long collimator tube showed the best performance, yielding around 55 % of background-free bunches and a potential X-ray count rate of 500 s^{-1} . In case that a new beam duct design at the end of CyDet is adopted by the collaboration, the detector distance should be maximized up to around 7.5 meter for an angle of 11.5° to gain additional background suppression. The performed study was limited to MC truth data and assumed perfect shielding and collimation. Further work is required to assess the detector’s performance for realistic shielding materials. Additionally, careful consideration must be given to optimizing the neutron shielding. This should align with collaborative-wide efforts addressing this background to ensure neutron exposure remains below $10^9 \text{ n}_{\text{eq}}/\text{cm}^2$ within the 150 days Phase-I runtime.

In conclusion, this work significantly contributes to the development of the COMET experiment and its pursuit of measuring μ -e conversion. The presented information translates to a direct improvement of the experiment’s SES and lays the foundation for transitioning from a simulation-based muon-stopping target to its real-world implementation. Although many aspects remain based on simulation, these results will help to realize the COMET experiment and to ensure the upcoming Phase-I is a resounding success.

List of Acronyms

ADC	Analog to Digital Converter
BSM	Physics Beyond the Standard Model
CDC	Cylindrical Drift Chamber
CLFV	Charged Lepton Flavour Violation
CMT	Configuration Management Tool
COMET	COherent Muon to Electron Transition
CPUh	ComPUting hours
CSDA	Continuous Slowing Down Approximation
CTH	CyDet Trigger Hodoscope
CyDet	Cylindrical Detector system
DAQ	Data AcQuisition
DIO	Decay In Orbit
DS	DownStream
ENSDF	Evaluated Nuclear Structure Data File
HPGe	High Purity Germanium
ICEDUST	Integrated COMET Experiment Data User Software Toolkit
J-PARC	Japan Proton Accelerator Research Complex
LFV	Lepton Flavour Violation
MC	Monte Carlo
MC5	Monte Carlo dataset 5
MC6	Monte Carlo dataset 6
NMC	Nuclear Muon Capture
PID	Particle IDentification
POT	Proton On Target
RFP	Resistive Feedback Preamplifier
RMC	Radiative Muon Capture
ROI	Region Of Interest
SES	Single Event Sensitivity
SM	Standard Model of particle physics
StrECAL	Straw-tube tracker and Electromagnetic CALorimeter
TDR	Technical Design Report
TRP	Transistor Reset Preamplifier
US	UpStream

References

- [1] S. H. Neddermeyer and C. D. Anderson. “Note on the Nature of Cosmic-Ray Particles”. In: *Phys. Rev.* 51 (1937), pp. 884–886. doi: [10.1103/PhysRev.51.884](https://doi.org/10.1103/PhysRev.51.884).
- [2] J. C. Street and E. C. Stevenson. “New Evidence for the Existence of a Particle of Mass Intermediate Between the Proton and Electron”. In: *Phys. Rev.* 52 (1937), pp. 1003–1004. doi: [10.1103/PhysRev.52.1003](https://doi.org/10.1103/PhysRev.52.1003).
- [3] Y. Nishina, M. Takeuchi, and T. Ichimiya. “On the Nature of Cosmic-Ray Particles”. In: *Phys. Rev.* 52 (1937), pp. 1198–1199. doi: [10.1103/PhysRev.52.1198](https://doi.org/10.1103/PhysRev.52.1198).
- [4] H. Yukawa. “On the Interaction of Elementary Particles. I”. In: *Prog. Theor. Exp. Phys. Supplement* 1 (1955), pp. 1–10. doi: [10.1143/PTPS.1.1](https://doi.org/10.1143/PTPS.1.1).
- [5] G. Bonomi, P. Checchia, et al. “Applications of cosmic-ray muons”. In: *Prog. Part. Nucl. Phys.* 112 (2020), p. 103768. doi: [10.1016/j.ppnp.2020.103768](https://doi.org/10.1016/j.ppnp.2020.103768).
- [6] L. W. Nordheim. “On the Nature of the Meson Decay”. In: *Phys. Rev.* 59 (1941), pp. 554–555. doi: [10.1103/PhysRev.59.554](https://doi.org/10.1103/PhysRev.59.554).
- [7] Y. Kuno and Y. Okada. “Muon decay and physics beyond the standard model”. In: *Rev. Mod. Phys.* 73 (2001), pp. 151–202. doi: [10.1103/RevModPhys.73.151](https://doi.org/10.1103/RevModPhys.73.151).
- [8] G. Danby, J.-M. Gaillard, et al. “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”. In: *Phys. Rev. Lett.* 9 (1962), pp. 36–44. doi: [10.1103/PhysRevLett.9.36](https://doi.org/10.1103/PhysRevLett.9.36).
- [9] C. S. Wu and L. Wilets. “Muonic Atoms and Nuclear Structure”. In: *Annu. Rev. Nucl. Sci.* 19.1 (1969), pp. 527–606. doi: [10.1146/annurev.ns.19.120169.002523](https://doi.org/10.1146/annurev.ns.19.120169.002523).
- [10] J. Hüfner, F. Scheck, and C. S. Wu. “Muon Physics”. In: *ISBN: 0123606039. Published by Academic Press, New York* (1977).
- [11] E. Fermi and E. Teller. “The Capture of Negative Mesotrons in Matter”. In: *Phys. Rev.* 72.5 (1947), pp. 399–408. doi: [10.1103/physrev.72.399](https://doi.org/10.1103/physrev.72.399).
- [12] T. von Egidy, D. H. Jakubassa-Amundsen, and F. J. Hartmann. “Calculation of muonic Coulomb-capture probabilities from electron binding energies”. In: *Phys. Rev. A* 29 (1984), pp. 455–461. doi: [10.1103/PhysRevA.29.455](https://doi.org/10.1103/PhysRevA.29.455).
- [13] F. J. Hartmann, T. von Egidy, et al. “Measurement of the Muonic X-Ray Cascade in Metallic Iron”. In: *Phys. Rev. Lett.* 37 (1976), pp. 331–334. doi: [10.1103/PhysRevLett.37.331](https://doi.org/10.1103/PhysRevLett.37.331).

[14] D. Measday. “The nuclear physics of muon capture”. In: *Phys. Rep.* 354.4 (2001), pp. 243–409. doi: [10.1016/S0370-1573\(01\)00012-6](https://doi.org/10.1016/S0370-1573(01)00012-6).

[15] A. Czarnecki, M. Dowling, et al. “Michel decay spectrum for a muon bound to a nucleus”. In: *Phys. Rev. D* 90 (2014), p. 093002. doi: [10.1103/PhysRevD.90.093002](https://doi.org/10.1103/PhysRevD.90.093002).

[16] J. Heeck, R. Szafron, and Y. Uesaka. “Isotope dependence of muon decay in orbit”. In: *Physical Review D* 105, 053006 (2022). doi: [10.1103/PhysRevD.105.053006](https://doi.org/10.1103/PhysRevD.105.053006).

[17] A. Czarnecki, X. G. i Tormo, and W. J. Marciano. “Muon decay in orbit: Spectrum of high-energy electrons”. In: *Phys. Rev. D* 84.1 (2011), p. 013006. doi: [10.1103/physrevd.84.013006](https://doi.org/10.1103/physrevd.84.013006).

[18] R. W. Huff. “Decay rate of bound muons”. In: *Ann. Phys.* 16.2 (1961), pp. 288–317. doi: [10.1016/0003-4916\(61\)90039-2](https://doi.org/10.1016/0003-4916(61)90039-2).

[19] K. Langanke and G. Martínez-Pinedo. “The role of electron capture in core-collapse supernovae”. In: *Nucl. Phys. A* 928 (2014), pp. 305–312. doi: [10.1016/j.nuclphysa.2014.04.015](https://doi.org/10.1016/j.nuclphysa.2014.04.015).

[20] D. Bajpai, L. Baudis, et al. *The MONUMENT Experiment: Ordinary Muon Capture studies for $0\nu\beta\beta$ decay*. 2024. doi: [10.48550/arXiv.2404.12686](https://doi.org/10.48550/arXiv.2404.12686).

[21] F. Wilczek. *Beyond The Standard Model: An Answer and Twenty Questions*. 1998. doi: [10.48550/arXiv.hep-ph/9802400](https://doi.org/10.48550/arXiv.hep-ph/9802400).

[22] P. A. M. Dirac. “The quantum theory of the Electron. Part II”. In: *Proc. R. Soc. Lond. A* 118 (1928), pp. 351–361. doi: [10.1098/rspa.1928.0056](https://doi.org/10.1098/rspa.1928.0056).

[23] F. Jegerlehner. *The Anomalous Magnetic Moment of the Muon*. Vol. 274. Springer, 2017. doi: [10.1007/978-3-319-63577-4](https://doi.org/10.1007/978-3-319-63577-4).

[24] G. W. Bennett, B. Bousquet, et al. “Final report of the E821 muon anomalous magnetic moment measurement at BNL”. In: *Phys. Rev. D* 73 (2006), p. 072003. doi: [10.1103/PhysRevD.73.072003](https://doi.org/10.1103/PhysRevD.73.072003).

[25] T. Aoyama, N. Asmussen, et al. “The anomalous magnetic moment of the muon in the Standard Model”. In: *Phys. Rep.* 887 (2020), pp. 1–166. doi: [10.1016/j.physrep.2020.07.006](https://doi.org/10.1016/j.physrep.2020.07.006).

[26] D. P. Aguillard, T. Albahri, et al. “Measurement of the Positive Muon Anomalous Magnetic Moment to 0.20 ppm”. In: *Phys. Rev. Lett.* 131 (2023), p. 161802. doi: [10.1103/PhysRevLett.131.161802](https://doi.org/10.1103/PhysRevLett.131.161802).

[27] R. Davis, D. S. Harmer, and K. C. Hoffman. “Search for Neutrinos from the Sun”. In: *Phys. Rev. Lett.* 20 (1968), pp. 1205–1209. doi: [10.1103/PhysRevLett.20.1205](https://doi.org/10.1103/PhysRevLett.20.1205).

[28] Y. Fukuda, T. Hayakawa, et al. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. doi: [10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562).

[29] Q. R. Ahmad, R. C. Allen, et al. “Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by 8B Solar Neutrinos at the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 87 (2001), p. 071301. doi: [10.1103/PhysRevLett.87.071301](https://doi.org/10.1103/PhysRevLett.87.071301).

[30] S. Weinberg. “A Model of Leptons”. In: *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266. doi: [10.1103/PhysRevLett.19.1264](https://doi.org/10.1103/PhysRevLett.19.1264).

[31] G. Bellini, L. Ludhova, G. Ranucci, and F. L. Villante. “Neutrino Oscillations”. In: *Adv. High Energy Phys.* 2014 (2014). doi: [10.1155/2014/191960](https://doi.org/10.1155/2014/191960).

[32] M. Kobayashi and T. Maskawa. “CP-Violation in the Renormalizable Theory of Weak Interaction”. In: *Prog. Theor. Phys.* 49.2 (1973), pp. 652–657. doi: [10.1143/PTP.49.652](https://doi.org/10.1143/PTP.49.652).

[33] M. Aker, M. Balzer, et al. “KATRIN: status and prospects for the neutrino mass and beyond”. In: *Phys. G: Nucl. Part. Phys.* 49.10 (2022), p. 100501. doi: [10.1088/1361-6471/ac834e](https://doi.org/10.1088/1361-6471/ac834e).

[34] S. T. Petcov. “The Processes $\mu \rightarrow e + \gamma$, $\mu \rightarrow e + \bar{e}$, $\nu' \rightarrow \nu + \gamma$ in the Weinberg-Salam Model with Neutrino Mixing”. In: *Sov. J. Nucl. Phys.* 25 (1977), p. 340.

[35] R. Bernstein and P. S. Cooper. “Charged lepton flavor violation: An experimenter’s guide”. In: *Phys. Rep.* 532.2 (2013), pp. 27–64. doi: [10.1016/j.physrep.2013.07.002](https://doi.org/10.1016/j.physrep.2013.07.002).

[36] W. J. Marciano, T. Mori, and J. M. Roney. “Charged Lepton Flavor Violation Experiments”. In: *Ann. Rev. Nucl. Part. Sci.* 58 (2008), pp. 315–341. doi: [10.1146/annurev.nucl.58.110707.171126](https://doi.org/10.1146/annurev.nucl.58.110707.171126).

[37] S. Dimopoulos and H. Georgi. “Softly broken supersymmetry and SU(5)”. In: *Nucl. Phys. B* 193.1 (1981), pp. 150–162. doi: [10.1016/0550-3213\(81\)90522-8](https://doi.org/10.1016/0550-3213(81)90522-8).

[38] J. Bernabéu, E. Nardi, and D. Tommasini. “ μ – e in nuclei and Z' physics”. In: *Nucl. Phys. B* 409.1 (1993), pp. 69–86. doi: [10.1016/0550-3213\(93\)90446-V](https://doi.org/10.1016/0550-3213(93)90446-V).

[39] M. Bauer and M. Neubert. “Minimal Leptoquark Explanation for the $R_{D^{(*)}}$, R_K , and $(g-2)_\mu$ Anomalies”. In: *Phys. Rev. Lett.* 116 (2016), p. 141802. doi: [10.1103/PhysRevLett.116.141802](https://doi.org/10.1103/PhysRevLett.116.141802).

[40] K. N. Abazajian, M. A. Acero, et al. *Light Sterile Neutrinos: A White Paper*. 2012. doi: [10.48550/arXiv.1204.5379](https://doi.org/10.48550/arXiv.1204.5379).

[41] W. Altmannshofer and J. Zupan. *Snowmass White Paper: Flavor Model Building*. 2023. doi: [10.48550/arXiv.2203.07726](https://doi.org/10.48550/arXiv.2203.07726).

[42] N. Teshima. *Status of the DeeMe Experiment, an Experimental Search for μ – e Conversion at J-PARC MLF*. 2019. doi: [10.48550/arXiv.1911.07143](https://doi.org/10.48550/arXiv.1911.07143).

[43] The COMET Collaboration et al. “COMET Phase-I technical design report”. In: *Prog. Theor. Exp. Phys.* 2020.3 (2020). doi: [10.1093/ptep/ptz125](https://doi.org/10.1093/ptep/ptz125).

[44] L. Bartoszek, E. Barnes, et al. *Mu2e Technical Design Report*. 2015. doi: [10.48550/arXiv.1501.05241](https://doi.org/10.48550/arXiv.1501.05241).

[45] S. Davidson, B. Echenard, et al. *Charged Lepton Flavor Violation*. 2022. doi: [10.48550/arXiv.2209.00142](https://doi.org/10.48550/arXiv.2209.00142).

[46] S. Davidson and B. Echenard. “Reach and complementarity of $\mu \rightarrow e$ searches”. In: *Eur. Phys. J. C* 82 (2022). doi: [10.1140/epjc/s10052-022-10773-4](https://doi.org/10.1140/epjc/s10052-022-10773-4).

[47] J. Aystö, A. Baldini, et al. *Physics with Low-Energy Muons at a Neutrino Factory Complex*. 2001. doi: [10.48550/arXiv.hep-ph/0109217](https://doi.org/10.48550/arXiv.hep-ph/0109217).

[48] W. Bertl, R. Engfer, et al. “A search for μ -e conversion in muonic gold”. In: *Eur. Phys. J. C* 47.2 (2006), pp. 337–346. doi: [10.1140/epjc/s2006-02582-x](https://doi.org/10.1140/epjc/s2006-02582-x).

[49] K. Oishi. “Study of Sensitivity to Search for a Charged Lepton Flavor Violating Process”. In: *Ph.D. thesis* (2021). doi: <https://hdl.handle.net/2324/4474936>.

[50] M. Dubouchet. “Sensitivity and background estimates towards Phase-I of the COMET muon-to-electron conversion search”. In: *Ph.D. thesis* (2022). doi: [10.25560/105498](https://doi.org/10.25560/105498).

[51] H. Chiang, E. Oset, et al. “Coherent and incoherent (μ^- , e^-) conversion in nuclei”. In: *Nucl. Phys. A* 559.4 (1993), pp. 526–542. doi: [10.1016/0375-9474\(93\)90259-Z](https://doi.org/10.1016/0375-9474(93)90259-Z).

[52] E. Tiesinga, P. J. Mohr, D. B. Newell, and B. N. Taylor. “CODATA recommended values of the fundamental physical constants: 2018”. In: *Rev. Mod. Phys.* 93 (2021), p. 025010. doi: [10.1103/RevModPhys.93.025010](https://doi.org/10.1103/RevModPhys.93.025010).

[53] K. Arndt, H. Augustin, et al. “Technical design of the phase I Mu3e experiment”. In: *Nucl. Instrum. Methods Phys. Res. A* 1014 (2021), p. 165679. doi: [10.1016/j.nima.2021.165679](https://doi.org/10.1016/j.nima.2021.165679).

[54] Y. Fujii. “Studies of optical fibre and SiPM readout system for the Cylindrical Trigger Hodoscope in COMET Phase-I”. In: *The 30th International Symposium on Lepton Photon Interactions at High Energies* (2022). doi: [10.5281/zenodo.6781367](https://doi.org/10.5281/zenodo.6781367).

[55] D. F. Measday, T. J. Stocki, B. A. Moftah, and H. Tam. “ γ rays from muon capture in ^{27}Al and natural Si”. In: *Phys. Rev. C* 76 (2007), p. 035504. doi: [10.1103/PhysRevC.76.035504](https://doi.org/10.1103/PhysRevC.76.035504).

[56] T. S. Kosmas and I. E. Lagaris. “On the muon–nucleus integrals entering the neutrinoless $\mu^- \rightarrow e^-$ conversion rates”. In: *J. Phys. G: Nucl. Part. Phys.* 28.12 (2002), pp. 2907–2920. doi: [10.1088/0954-3899/28/12/302](https://doi.org/10.1088/0954-3899/28/12/302).

[57] T. Siiskonen, J. Suhonen, and T. S. Kosmas. “New limits for lepton-flavor violation from the $\mu^- \rightarrow e^-$ conversion in ^{27}Al ”. In: *Phys. Rev. C* 60.6 (1999), p. 062501. doi: [10.1103/physrevc.60.062501](https://doi.org/10.1103/physrevc.60.062501).

[58] T. Siiskonen, J. Suhonen, and T. S. Kosmas. “Realistic nuclear matrix elements for the lepton-flavor violating $\mu^- \rightarrow e^-$ conversion in ^{27}Al and ^{48}Ti ”. In: *Phys. Rev. C* 62 (2000), p. 035502. doi: [10.1103/PhysRevC.62.035502](https://doi.org/10.1103/PhysRevC.62.035502).

[59] B. E. Krikler. “Sensitivity and background estimates for phase-II of the COMET experiment”. In: *Ph.D. thesis* (2016). doi: [10.25560/45365](https://doi.org/10.25560/45365).

[60] J. Allison, K. Amako, et al. “Geant4 developments and applications”. In: *IEEE Transactions on Nuclear Science* 53.1 (2006), pp. 270–278. doi: [10.1109/tns.2006.869826](https://doi.org/10.1109/tns.2006.869826).

[61] E. Gillies. “COMET Phase-I Track Reconstruction using Machine Learning and Computer Vision”. In: *Ph.D. thesis* (2018). doi: [10.25560/68463](https://doi.org/10.25560/68463).

[62] R. Kitano, M. Koike, and Y. Okada. “Detailed calculation of lepton flavor violating muon-electron conversion rate for various nuclei”. In: *Phys. Rev. D* 66 (2002), p. 096002. doi: [10.1103/PhysRevD.66.096002](https://doi.org/10.1103/PhysRevD.66.096002).

[63] S. Davidson, Y. Kuno, and A. Saporta. “‘Spin-dependent’ $\mu \rightarrow e$ conversion on light nuclei”. In: *Eur. Phys. J. C* 78 (2018). doi: [10.1140/epjc/s10052-018-5584-8](https://doi.org/10.1140/epjc/s10052-018-5584-8).

[64] S. Davidson, Y. Kuno, and M. Yamanaka. “Selecting $\mu \rightarrow e$ conversion targets to distinguish lepton flavour-changing operators”. In: *Phys. Lett. B* 790 (2019), pp. 380–388. doi: [10.1016/j.physletb.2019.01.042](https://doi.org/10.1016/j.physletb.2019.01.042).

[65] V. Cirigliano, R. Kitano, Y. Okada, and P. Tuzon. “Model discriminating power of $\mu \rightarrow e$ conversion in nuclei”. In: *Phys. Rev. D* 80 (2009), p. 013002. doi: [10.1103/PhysRevD.80.013002](https://doi.org/10.1103/PhysRevD.80.013002).

[66] F. J. Hartmann, R. Bergmann, et al. “Measurement of the muonic x-ray cascade in Mg, Al, In, Ho, and Au”. In: *Z. Phys. A* 305 (1982), pp. 189–204. doi: doi.org/10.1007/BF01417434.

[67] D. F. Measday, T. J. Stocki, B. A. Moftah, and H. Tam. “ γ rays from muon capture in ^{27}Al and natural Si”. In: *Phys. Rev. C* 76.3 (2007), p. 035504. doi: [10.1103/physrevc.76.035504](https://doi.org/10.1103/physrevc.76.035504).

[68] Geant4 Collaboration. *Guide For Physics Lists, Release 11.1*. URL: <https://geant4-userdoc.web.cern.ch/UsersGuides/PhysicsListGuide/fo/PhysicsListGuide.pdf>.

[69] A. J. M. Plompen, O. Cabellos, et al. “The joint evaluated fission and fusion nuclear data library, JEFF-3.3”. In: *Eur. Phys. J. A* 56.7 (2020). doi: [10.1140/epja/s10050-020-00141-9](https://doi.org/10.1140/epja/s10050-020-00141-9).

[70] J. Tuli. “Evaluated nuclear structure data file”. In: *Nucl. Instrum. Methods Phys. Res. A* 369.2-3 (1996), pp. 506–510. doi: [10.1016/s0168-9002\(96\)80040-4](https://doi.org/10.1016/s0168-9002(96)80040-4).

[71] Geant4 Collaboration. *Physics Reference Manual, Release 11.1*. URL: <https://geant4-userdoc.web.cern.ch/UsersGuides/PhysicsReferenceManual/fo/PhysicsReferenceManual.pdf>.

[72] A. Edmonds, J. Quirk, et al. “Measurement of proton, deuteron, triton, and α particle emission after nuclear muon capture on Al, Si, and Ti with the AlCap experiment”. In: *Phys. Rev. C* 105 (2022), p. 035501. doi: [10.1103/PhysRevC.105.035501](https://doi.org/10.1103/PhysRevC.105.035501).

[73] C. Arnault. *CMT: A software configuration management tool*. In Proceedings of CHEP, Padova, Italy, 2000.

[74] Kitware, Inc. *CMake Documentation*. URL: <https://cmake.org/cmake/help/latest/>.

[75] A. S. Wightman. “Moderation of Negative Mesons in Hydrogen I: Moderation from High Energies to Capture by an H_2 Molecule”. In: *Phys. Rev.* 77 (1950), pp. 521–528. doi: [10.1103/PhysRev.77.521](https://doi.org/10.1103/PhysRev.77.521).

[76] J. S. Cohen. “Slowing down and capture of negative muons by hydrogen: Classical-trajectory Monte Carlo calculation”. In: *Phys. Rev. A* 27.1 (1983), pp. 167–179. doi: [10.1103/physreva.27.167](https://doi.org/10.1103/physreva.27.167).

[77] A. Jansen. “Experimentelle Untersuchung des Targets des COMET Myon-Elektron Konversionsexperimentes”. Deutsch. In: *Master thesis* (2017).

[78] M. Moritsu, Y. KUNO, et al. “Construction and performance tests of the COMET CDC”. In: 2019, p. 538. doi: [10.22323/1.340.0538](https://doi.org/10.22323/1.340.0538).

[79] C. Wu, T. Wong, et al. “Test of a small prototype of the COMET cylindrical drift chamber”. In: *Nucl. Instrum. Methods Phys. Res. A* 1015 (2021), p. 165756. doi: [10.1016/j.nima.2021.165756](https://doi.org/10.1016/j.nima.2021.165756).

[80] M. G. Catanesi, E. Radicioni, et al. “Large-angle production of charged pions with 3–12.9 GeV/c incident protons on nuclear targets”. In: *Phys. Rev. C* 77 (2008), p. 055207. doi: [10.1103/PhysRevC.77.055207](https://doi.org/10.1103/PhysRevC.77.055207).

[81] K. Oishi. “COMET Phase- α Experiment to Investigate COMET’s New Muon Beamline at J-PARC”. In: Zenodo, 2023. doi: [10.5281/zenodo.8370320](https://doi.org/10.5281/zenodo.8370320).

[82] H. W. Koch and J. W. Motz. “Bremsstrahlung Cross-Section Formulas and Related Data”. In: *Rev. Mod. Phys.* 31.4 (1959), pp. 920–955. doi: [10.1103/revmodphys.31.920](https://doi.org/10.1103/revmodphys.31.920).

[83] W. R. Leo. *Techniques for nuclear and particle physics experiments: a how-to approach*. Springer Science & Business Media, 2012.

[84] S. M. Seltzer and M. J. Berger. “Energy Loss Straggling of Protons and Mesons: Tabulation of the Vavilov Distribution”. In: *Studies in Penetration of Charged Particles in Matter* National Academy of Sciences Publication 1133, Nuclear Science Series Report No. 39 (1964).

[85] T. Suzuki, D. F. Measday, and J. P. Roalsvig. “Total nuclear capture rates for negative muons”. In: *Phys. Rev. C* 35 (1987), pp. 2212–2224. doi: [10.1103/PhysRevC.35.2212](https://doi.org/10.1103/PhysRevC.35.2212).

[86] Thyssenkrupp. *Material datasheet for the aluminum-zinc alloy EN AW-7075*. URL: https://datenblaetter.thyssenkrupp.ch/en_aw_7075_0717.pdf.

[87] SIMULIA. *OPERA 3D – The electromagnetic and electromechanical simulation pacakge*. URL: <https://www.3ds.com/products-services/simulia/products/opera/>.

[88] G. Fricke, C. Bernhardt, et al. “Nuclear Ground State Charge Radii from Electromagnetic Interactions”. In: *At. Data Nucl. Data Tables* 60.2 (1995), pp. 177–285. doi: [10.1006/adnd.1995.1007](https://doi.org/10.1006/adnd.1995.1007).

[89] ORTEC. *GMX Series Coaxial HPGe Detector Product Configuration Guide*. URL: <https://www.ortec-online.com/-/media/ametekortec/brochures/gamma-x.pdf?la=en&revision=c6249801-7468-4c42-88b1-37cf43bf291c>.

[90] D. F. Measday, T. J. Stocki, B. A. Moftah, and H. Tam. “ γ rays from muon capture in ^{27}Al and natural Si”. In: *Phys. Rev. C* 76 (2007), p. 035504. doi: [10.1103/PhysRevC.76.035504](https://doi.org/10.1103/PhysRevC.76.035504).

[91] P. Bossew. “A very long-term HPGe-background gamma spectrum”. In: *Appl. Radiat. Isot.* 62.4 (2005), pp. 635–644. doi: [10.1016/j.apradiso.2004.09.006](https://doi.org/10.1016/j.apradiso.2004.09.006).

[92] The AlCap Collaboration et al. *A Measurement of Proton, Deuteron, Triton and Alpha Particle Emission after Nuclear Muon Capture on Al, Si and Ti with the AlCap Experiment*. 2022. doi: [10.48550/arXiv.2110.10228](https://doi.org/10.48550/arXiv.2110.10228).

[93] G. Gilmore. “Practical Gamma-Ray Spectrometry: Second Edition”. In: *ISBN: 978-0-470-86196-7. Published by Wiley-VCH Verlag, Weinheim, Germany* (2008). doi: [10.1002/9780470861981](https://doi.org/10.1002/9780470861981).

[94] S. Turkat. “Primordial nuclides and low-level counting at Felsenkeller”. In: *Ph.D. thesis* (2023).

[95] G. Audi, A. Wapstra, and C. Thibault. “The AME2003 atomic mass evaluation: (II). Tables, graphs and references”. In: *Nucl. Phys. A* 729.1 (2003), pp. 337–676. doi: [10.1016/j.nuclphysa.2003.11.003](https://doi.org/10.1016/j.nuclphysa.2003.11.003).

[96] N. C. Mukhopadhyay. “Nuclear Muon capture”. In: *Phy. Rep.* 30.1 (1977), pp. 1–144. doi: [10.1016/0370-1573\(77\)90073-4](https://doi.org/10.1016/0370-1573(77)90073-4).

[97] N. J. Zaluzec. “Analytical Formulae for Calculation of X-Ray Detector Solid Angles in the Scanning and Scanning/Transmission Analytical Electron Microscope”. In: *Microsc. Microanal.* 20.4 (2014), pp. 1318–1326. doi: [10.1017/S1431927614000956](https://doi.org/10.1017/S1431927614000956).

[98] M. J. Berger, J. H. Hubbell, et al. *XCOM: Photon Cross Sections Database*. National Institute of Standards and Technology. doi: [10.18434/T48G6X](https://doi.org/10.18434/T48G6X).

[99] Y.-S. Tsai. “Pair production and bremsstrahlung of charged leptons”. In: *Rev. Mod. Phys.* 46 (1974), pp. 815–851. doi: [10.1103/RevModPhys.46.815](https://doi.org/10.1103/RevModPhys.46.815).

[100] M. J. Berger, J. S. Coursey, M. A. Zucker, and J. Chang. *ESTAR, PSTAR, and ASTAR: Computer Programs for Calculating Stopping-Power and Range Tables for Electrons, Protons, and Helium Ions*. National Institute of Standards and Technology. doi: [10.18434/T4NC7P](https://doi.org/10.18434/T4NC7P).

[101] P.A. Zyla et al. (Particle Data Group). “Review of Particle Physics”. In: *Prog. Theor. Exp. Phys.* 083C01 (2020). doi: [10.1093/ptep/ptaa104](https://doi.org/10.1093/ptep/ptaa104).

[102] N. Soppera, M. Bossant, and E. Dupont. “JANIS 4: An Improved Version of the NEA Java-based Nuclear Data Information System”. In: *Nucl. Data Sheets* 120 (2014), pp. 294–296. doi: [10.1016/j.nds.2014.07.071](https://doi.org/10.1016/j.nds.2014.07.071).

[103] V. Zerkin and B. Pritychenko. “The experimental nuclear reaction data (EXFOR): Extended computer database and Web retrieval system”. In: *Nucl. Instrum. Methods Phys. Res. A* 888 (2018), pp. 31–43. doi: [10.1016/j.nima.2018.01.045](https://doi.org/10.1016/j.nima.2018.01.045).

[104] H. W. Kraner, R. H. Pehl, and E. E. Haller. “Fast Neutron Radiation Damage of High-Purity Germanium Detectors”. In: *IEEE Trans. Nucl. Sci.* 22.1 (1975), pp. 149–159. doi: [10.1109/TNS.1975.4327633](https://doi.org/10.1109/TNS.1975.4327633).

[105] M. Stone, L. Crow, V. Fanelli, and J. Niedziela. “Characterization of shielding materials used in neutron scattering instrumentation”. In: *Nucl. Instrum. Methods Phys. Res. A* 946 (2019), p. 162708. doi: [10.1016/j.nima.2019.162708](https://doi.org/10.1016/j.nima.2019.162708).

[106] G. F. Knoll. “Radiation detection and measurement”. In: *ISBN: 978-0-470-13148-0. Published by John Wiley & Sons* (2010).

[107] R. M. Keyser and T. R. Twomey. “Optimization of Pulse Processing Parameters for HPGe Gamma-ray Spectroscopy Systems Used in Extreme Count Rate Conditions and Wide Count Rate Ranges”. In: *J. Radioanal. Nucl. Chem.* 296 (2013), pp. 503–508. doi: [10.1007/s10967-012-2113-3](https://doi.org/10.1007/s10967-012-2113-3).

[108] ORTEC. *UM1934 - CAENComm User & Reference Manual*. URL: https://www.npl.washington.edu/TRIMS/sites/sand.npl.washington.edu.TRIMS/files/manuals-documentation/CAENDigitizer_SW_User_Manual_rel.5.pdf.

[109] C. H. Westcott and M. L. E. Oliphant. “A study of expected loss rates in the counting of particles from pulsed sources”. In: *Proc. R. Soc. (London), Ser. A* 194.1039 (1948), pp. 508–526. doi: [10.1098/rspa.1948.0094](https://doi.org/10.1098/rspa.1948.0094).

- [110] U. von Möllendorff. “Note on dead-time losses for pulsed radiation”. In: *Nucl. Instrum. Methods Phys. Res. A* 356.2 (1995), pp. 468–469. doi: [10.1016/0168-9002\(94\)01255-5](https://doi.org/10.1016/0168-9002(94)01255-5).
- [111] P. C. Johns and M. J. Yaffe. “Correction of pulse-height spectra for peak pileup effects using periodic and random pulse generators”. In: *Nucl. Instrum. Methods Phys. Res. A* 255.3 (1987), pp. 559–581. doi: [10.1016/0168-9002\(87\)91227-7](https://doi.org/10.1016/0168-9002(87)91227-7).
- [112] R. M. Lindstrom and R. F. Fleming. “Dead time, pileup, and accurate gamma-ray spectrometry”. In: *Radioactivity and Radiochemistry* 6.2 (1995).
- [113] V. Mossa, K. Stöckel, et al. “Setup commissioning for an improved measurement of the D(p,γ)³He cross section at Big Bang Nucleosynthesis energies”. In: *Phys. J. A* 56.144 (2020). doi: [10.1140/epja/s10050-020-00149-1](https://doi.org/10.1140/epja/s10050-020-00149-1).
- [114] M. Osswald. “Test und Kalibrierung der Reinstgermaniumdetektoren für die ³He(α,γ)⁷Be-Messung am Felsenkeller”. In: *Bachelor thesis* (2020).
- [115] M. Shamsuzzoha Basunia. “Nuclear Data Sheets for A = 27”. In: *Nucl. Data Sheets* 112.8 (2011), pp. 1875–1948. doi: [10.1016/j.nds.2011.08.001](https://doi.org/10.1016/j.nds.2011.08.001).

A. Collimator Performances

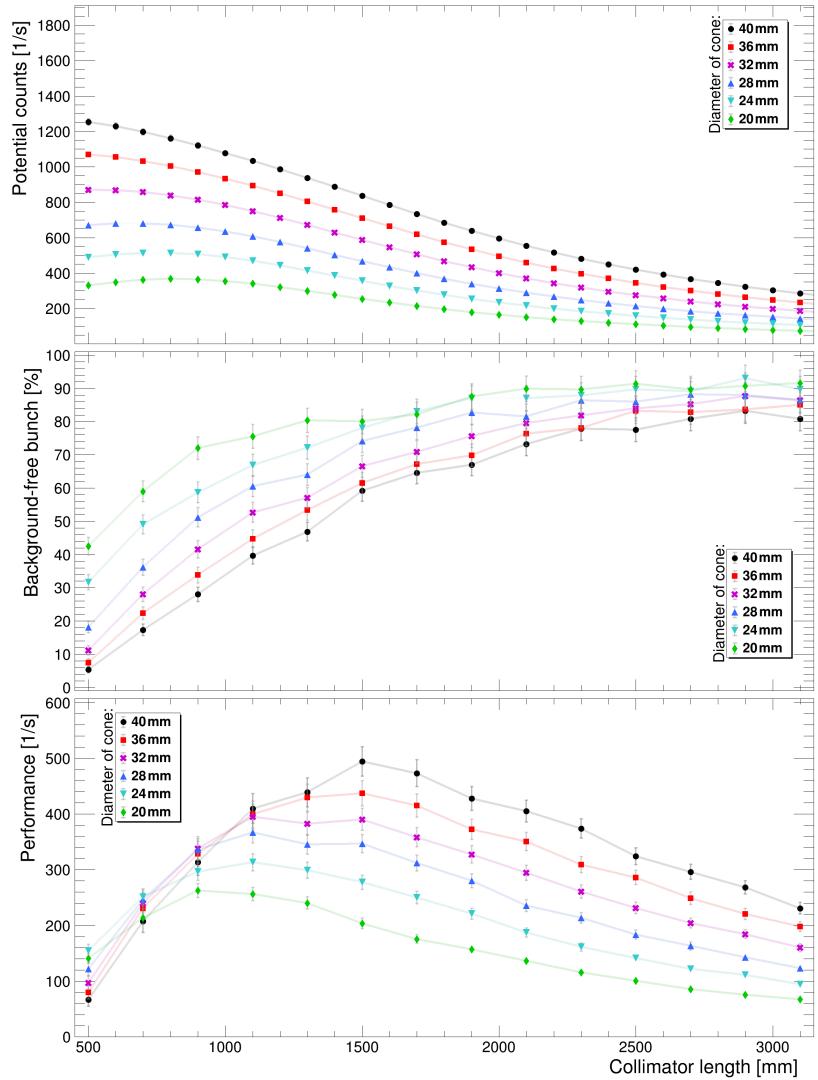


Figure A.1.: Summary for the cone collimator, depicting the potential X-ray count rate (top), the percentage of background-free bunches (middle), and the product of both given as a performance value (bottom). Colors correspond to the cone's diameter at the HPGe-far side. The diameter at the HPGe-close side is set to 50 mm. Values were obtained according to Section 6.1. Straight lines connecting the data points are included to increase clarity. Error bars include statistical uncertainties and an assumed 1 % systematic uncertainty on the number of muons stopped per bunch and the full-energy peak efficiency. See Section 6.2 for discussion.

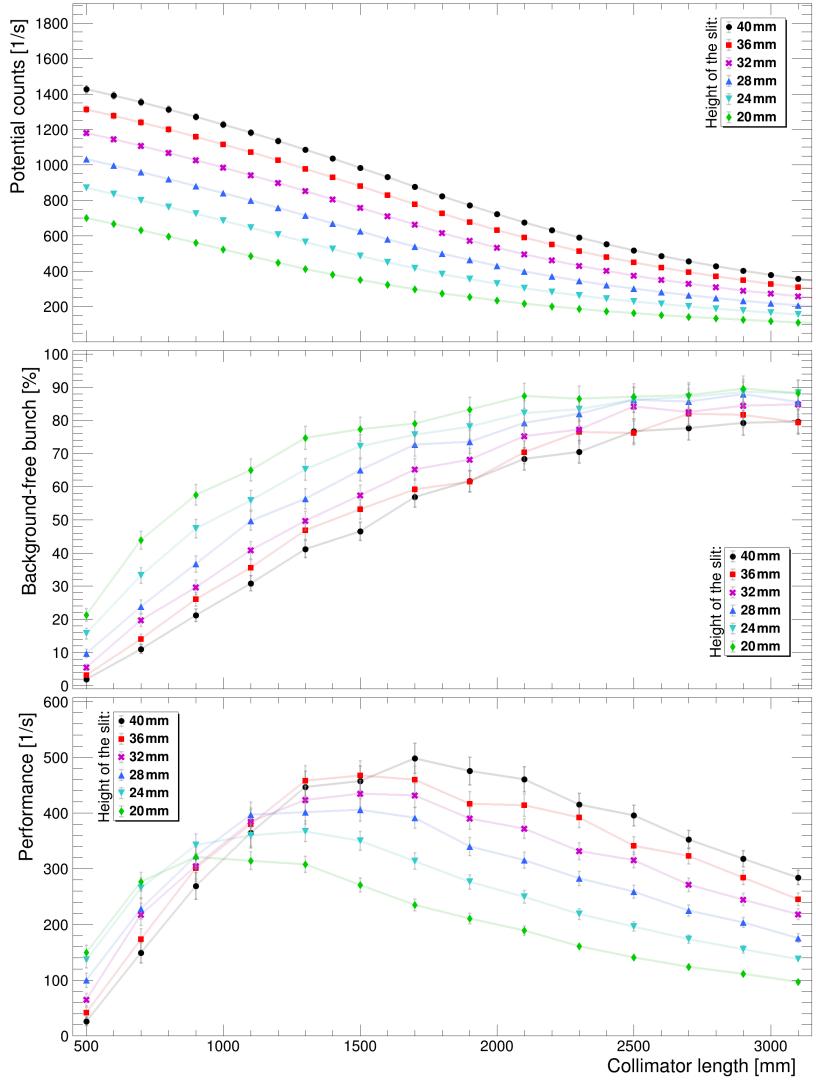


Figure A.2.: Summary for the slit-type tube collimator, depicting the potential X-ray count rate (top), the percentage of background-free bunches (middle), and the product of both given as a performance value (bottom). The tube's diameter is set to 50 mm. Colors correspond to the height of a slit, determined by height blockers positioned within the top and bottom half of the tube. Values were obtained according to Section 6.1. Straight lines connecting the data points are included to increase clarity. Error bars include statistical uncertainties and an assumed 1 % systematic uncertainty on the number of muons stopped per bunch and the full-energy peak efficiency. See Section 6.2 for discussion.

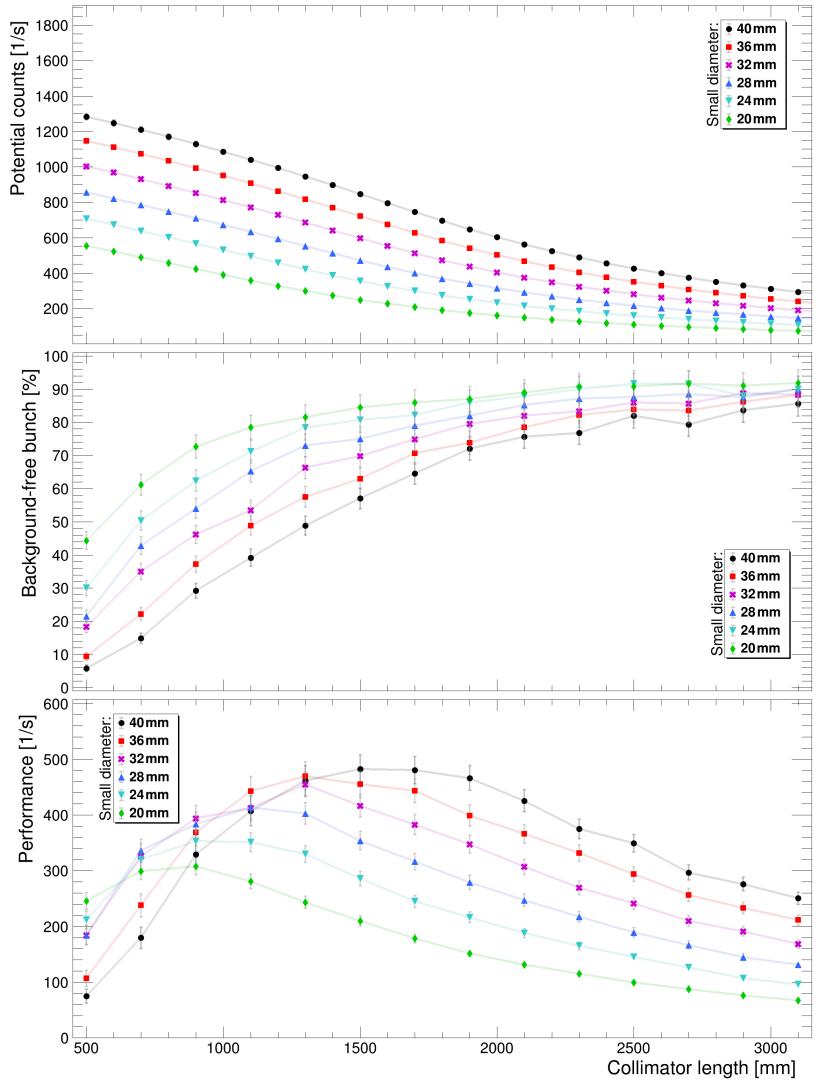


Figure A.3.: Summary for the ellipse collimator, depicting the potential X-ray count rate (top), the percentage of background-free bunches (middle), and the product of both given as a performance value (bottom). Colors correspond to the small diameter of the ellipse. The large diameter is set to 50 mm. Values were obtained according to Section 6.1. Straight lines connecting the data points are included to increase clarity. Error bars include statistical uncertainties and an assumed 1 % systematic uncertainty on the number of muons stopped per bunch and the full-energy peak efficiency. See Section 6.2 for discussion.

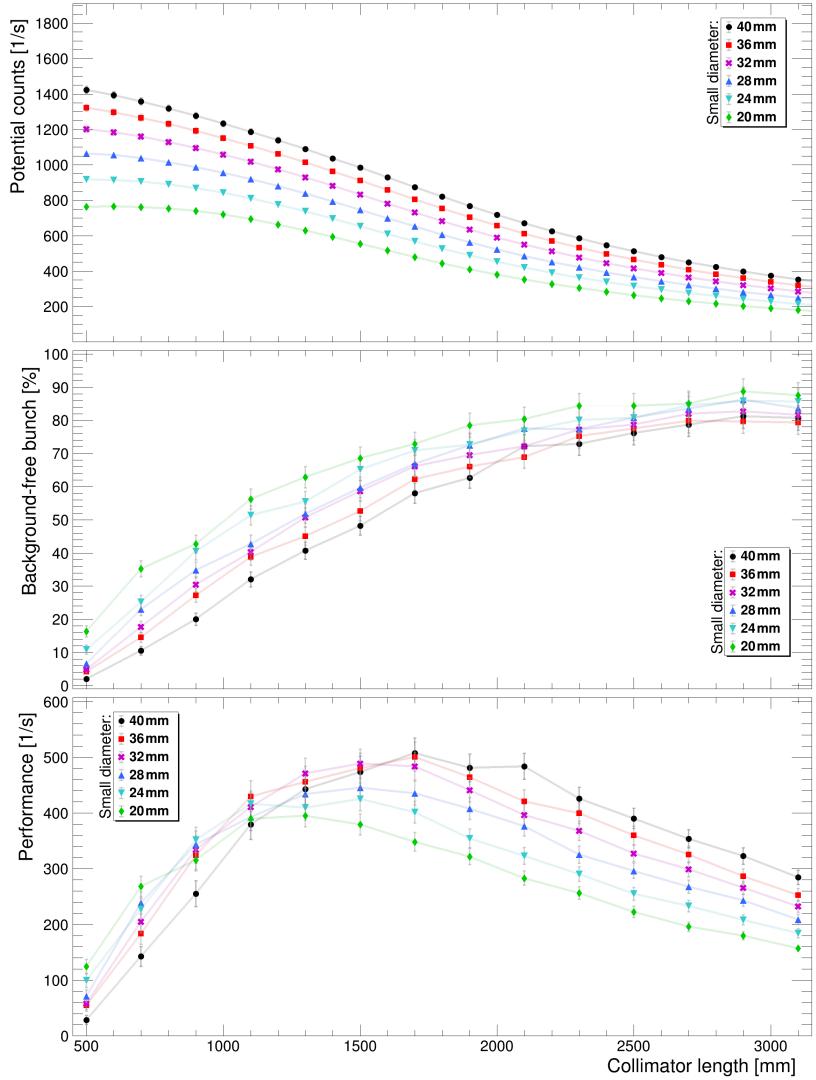


Figure A.4.: Summary for the elliptical cone collimator, depicting the potential X-ray count rate (top), the percentage of background-free bunches (middle), and the product of both given as a performance value (bottom). Colors correspond to the small diameter of the ellipse at the HPGe-far side. Its large diameter is set to 50 mm. The collimator morphs to a tube with a 50 mm diameter towards the HPGe-close side. Values were obtained according to Section 6.1. Straight lines connecting the data points are included to increase clarity. Error bars include statistical uncertainties and an assumed 1 % systematic uncertainty on the number of muons stopped per bunch and the full-energy peak efficiency. See Section 6.2 for discussion.

Acknowledgments

“I can’t believe it... it is finally done!” – me, the day before submission, and probably everybody who followed my painful odyssey of completing this PhD. After (too many) years of research and (what feels like even more years of) writing, this work not only stands for my hardest-ever attempt at *something*, but it resembles a journey – a long journey through countless doors and countless faces, each earning my eternal gratitude, as this personal milestone would not have been possible without any of you.

First of all, I want to express my gratitude toward *Prof. Kai Zuber*. Kai, not only did you provide me with the opportunity to embark on this journey to begin with, but over the years, you became more than just my doctoral advisor. Thank you for offering guidance and support for every idea I had, and for gracefully managing the balance between friend and supervisor. Without question, I also want to extend my thanks to *Prof. Arno Straessner*. I will always be grateful that you inspired me to stay in physics with your incredible particle-physics lectures, numerous impromptu discussions about basically anything really in front of my room, and your consistently encouraging and uplifting attitude. Much of the latter sentiment can also be extended to *Prof. Satoshi Mihara*. I deeply admire your composed, supportive, and always open attitude toward us young research students in the collaboration. Neither of you hesitated for even a split-second when I asked you to undertake the daunting task of serving as my referees. Thank you so much!

I also want to extend a heartfelt thank you to all of my collaborators on the COMET experiment. You have endured me for so many years through countless discussions and naive questions at collaboration meetings. Specifically, I want to express my gratitude to *Prof. Yoshitaka Kuno* – not only for your truly tremendous work as COMET’s spokesperson but also for your interest in my work, insightful commentary, thoughtful questioning, reinforcing feedback, and for providing me with the opportunity for my half-year stay at Osaka. During this stay in Japan, I had the incredible fortune to meet two remarkable people: *Sam* and *Dorian*. You two are the first people I think of when returning to Japan, and you always provide me with wonderful and unforgettable memories, for which I am incredibly thankful.

Also in Germany, I am grateful for the endless support I received in everything surrounding my workplace at IKTP. Thank you, *Nicole* and *Birgit*, for always trying your best to shield me from the monstrosity that is German bureaucracy, for your open ears and hearts, and for working hard to make the IKTP a genuinely better place. Thank you, *Thomas* and *Theresa*, for the countless rants I needed to relieve from my heart and for the time I stole from you with my physics questions. Thank you, *Andreas* – not for stealing my name, but for providing me with all my everyday electronics and general working-tool needs. And,

last but not least, thanks to all my current and former group members of the IKTP neutrino group – just to name a few: *Alex; Birgit; Christoph; Hans; Heinrich; Jonas; Juliane; Marie; Stefan; Steffen; Xianke; and Yingjie*. I am indebted to every one of you for your constant feedback in meetings and relieving discussions during lunch. In particular, I want to thank *Juliane* for enduring my constant nagging in our office when you want to concentrate and get things done and for deliberately taking time out of your busy days to correct some of this work. *Marie*, for providing delicious cakes, an uplifting spirit, and emotional support in the final days of this work. And – the best guys come last – you two: *Steffen* and *Heinrich*. I honestly struggle to fathom what an incredible ride it was, and I can say without shame that I would not write this today without you. While with my submission, I feel like one incredible era is drawing to a close, I truly hope that this just marks the beginning of a new story and countless more holidays and memories to come.

Arne and *Tom*, I am proud to say that you two have been my best friends for so many years. Our weekly sessions of DnD, board games, cooking, or just talking and enjoying our time together have helped me tremendously through many stressful days and weeks. I am never bored when we are together, and I hope that this friendship, which has already lasted for so many years, will continue to last tenfold as many years into the future.

Mama und *Papa*, ihr seid immer an meiner Seite und deckt mir den Rücken. Obgleich der manchen Frage nach “dem Abschluss deines Doktors” weiß Ich doch stets, dass ihr immer nur das Beste für mich wollt(et) und mich auch viele weitere Jahre noch unterstützt hättest. *Bruderherz*, auch wenn wir so einige Reiberein in unserer Kindheit hatten, bist du mir einfach nur ans Herz gewachsen. Danke für die vielen besuchten und natürlich noch auf viele weitere Konzerte zusammen. Ich danke euch wirklich allen aus tiefsten Herzen!

Lastly, but by no means least, I want to thank you, *Anna-Caroline*. You endure me through my best, but most importantly, also through my worst days and nights. You are my sunshine, the light in the lighthouse I call home, and the support I never realized I needed. Thank you for your support, your caring, your endurance, and your patience through all these last years. I know you went above and beyond, handling more than just the occasional housework when I was silently overwhelmed by my ‘PhD stress’. I want you to know that I am dedicated to becoming a better version of myself. I love you.