



UNIVERSITÀ DI PISA

Department of Physics E. Fermi

Master's Degree in Physics

Commissioning of the Mu2e tracker DAQ, planning for the Vertical Slice Test and pre-pattern recognition studies

Supervisors:

Dr. Pavel Murat

Prof. Simone Donati

Candidate:

Sara Gamba

Academic Year 2023/2024

*Ai miei nonni.
Per essere i miei porti sicuri e pilastri.
Per essere le persone che tanto ammiro.*

*A Jacopo.
Per la tua pazienza e per sapere ascoltare.
Per avermi svegliato quelle mattine nere e per avermi portato a pattinare.
Per essere, sempre, il punto di riferimento.*

*Ai miei genitori.
A mia mamma, per avermi insegnato il coraggio di buttarmi in ogni esperienza.
A mio babbo, per avermi insegnato la determinazione nel perseguire i miei obiettivi.*

*A mia sorella.
Per quei fogli e lettere sotto la porta ad ogni esame in quarantena.
Per essere il più bel regalo che la vita abbia potuto farmi.
Per avere reso casa un luogo felice.*

Acknowledgements

I would like to extend my deepest gratitude to my supervisor, P. Murat, for his invaluable guidance. His teachings went far beyond textbooks, showing me how to think critically and like a physicist. His patience and support have been essential throughout my time at Fermilab.

My sincere thanks also go to Professor S. Donati for his support and understanding during this year. I am grateful to him for giving me the opportunity to have such a significant experience at Fermilab, and for his impartial guidance in helping me shape my career.

I would like to express my immense gratitude to Namitha and Sridhar, who have always helped me with work and software issues, but most importantly, for their friendship, which has been a truly pleasant discovery.

Vorrei ringraziare la mia famiglia materna e paterna, per l'immenso sostegno, l'affetto e la considerazione che mi hanno sempre dato. Grazie per aver condiviso con gioia i traguardi importanti e per avermi aiutato nei momenti complessi. Grazie per avermi fatto sentire vicino a voi nonostante la grande lontananza da casa.

Grazie a tutti gli amici: a quelli dell'infanzia, delle superiori, del mare, al gruppo della pallavolo, a quelli dell'Università, ai bimbi del Fermilab, a chi ho conosciuto più recentemente, a quelli che sono stati importanti punti di riferimento nella mia vita e che l'hanno profondamente segnata, a chi ho avuto modo di conoscere meno, a quelli con cui ho passato momenti felici, brutti, a coloro con cui ho riso, pianto, litigato, viaggiato. Vi porto sempre con me. Un ringraziamento particolare alle mie amiche, Giulia, Rebe, Giulia, Eli, Silvi, Vane, Marti, soprattutto per questo ultimo anno (e non solo), per le parole di conforto, il sostegno e le risate.

Vorrei esprimere la mia più profonda gratitudine a Gianluca e Giuseppina, che si sono presi cura di me in modo davvero impeccabile e che mi hanno dato gli strumenti per smettere di vedere il mondo solo in nero.

Questo ultimo ringraziamento sarà compreso solo da chi ha avuto un animale nella vita. Un grande grazie a Baghy e Chloè, per avermi ascoltato ripetere dalle elementari fino, con qualche pausa nel frangente, all'Università e per avermi dato conforto nei momenti di solitudine. In particolare vorrei menzionare la profonda protezione di Baghy in qualsiasi momento difficile.

Alle persone che ho menzionato e non, grazie per avermi reso chi sono oggi.

Abstract

The primary objective of the Mu2e experiment at Fermilab is to search for the neutrino-less coherent $\mu \rightarrow e$ conversion in the field of an aluminum nucleus ($\mu^- \text{Al} \rightarrow e^- \text{Al}$). The signature of this process is a monochromatic Conversion Electron (CE) with an energy of approximately 104.97 MeV [1]. Within the Standard Model (SM), the branching ratio for this process, including neutrino masses and oscillation, is expected to be less than $\mathcal{O}(10^{-50})$. This value is far beyond current experimental capabilities. However, models of physics beyond the SM predict much higher relative rates, approaching an observable level. The SINDRUM II experiment set an upper limit on muon conversion at 7×10^{-13} (90% CL) on Au target [2], and the Mu2e collaboration aims to improve this limit by four orders of magnitude. Observing this process would provide a clear evidence of physics beyond the Standard Model. A brief discussion of the theoretical and experimental aspects is provided in Chapter 1.

Mu2e adopts a sophisticated experimental setup to achieve its goals, further described in Chapter 2. The central part of the Mu2e detector is the tracker, that consists of 18 tracking stations. The tracker must provide excellent momentum resolution, approximately 1 MeV/c, to distinguish the monochromatic CE signal from the background. To minimize the energy losses, a straw tube tracker will be used [3]. Chapter 3 provides an overview of the straw tracker design and its working principles.

This Thesis presents a comprehensive study of the Mu2e tracker, covering complementary aspects from initial commissioning to optimization and first steps of the calibration processes. My work at Fermilab has been focused on the complete Data Acquisition (DAQ) testing from both hardware and software perspectives. I was involved in the commissioning of the Mu2e DAQ system and the Vertical Slice Test (VST) of the tracker. The VST encompasses the entire testing chain, from the straws to the readout, and to processed data on disk. I was also focused on the offline analysis, especially on pre-pattern recognition studies, to explore the best methods for identifying δ -electrons during the data taking.

Chapter 4 details the commissioning of the tracker DAQ system, emphasizing the importance of understanding of the readout process before the data acquisition. This includes validating the readout logic and firmware through Monte Carlo simulations to confirm functionality and buffering, monitoring the quality of the data from the tracker preamplifiers and front-end electronics, and assessing overall DAQ performance to ensure reliability during future calibration and data-taking.

Chapter 5 discusses the initial steps towards the tracker calibration. The ultimate goal is to perform a time calibration of the first assembled station of the tracker using cosmic muons, aiming for a longitudinal hit position resolution better than 4 cm. This involves determining the signal propagation times and channel-to-channel delays. I performed a Monte Carlo study to determine the impact of

the station orientation on the quality of the calibration, in particular on the cosmic track reconstruction, focusing on potential biases that could arise. These studies provide essential insights into the operation, optimization, and calibration of the Mu2e tracker system.

Given the high data volume expected during Mu2e operations, estimated at approximately 7 PBytes per year, optimizing memory usage and minimizing CPU consumption are critical. A significant challenge lies in effectively flagging δ -electron hits, which are the primary source of hits in the tracker, without compromising the efficiency of CE hit detection and track reconstruction. A detailed study of pre-pattern recognition and a thorough comparison of two δ -electron flagging algorithms is provided in Chapter 6.

In Chapter 7, the findings are concisely summarized, offering a comprehensive synthesis of the research and emphasizing the key insights derived from this study.

Contents

1	Charged Lepton Flavour Violation	11
1.1	Theoretical introduction	11
1.1.1	The Standard Model	11
1.1.1.1	Bosons	11
1.1.1.2	Fermions	12
1.1.2	History of flavour	13
1.1.3	Lepton sector in the Standard Model	14
1.1.4	Overview of Charged Lepton Flavour Violation	15
1.1.5	CLFV in the SM with massive neutrinos	16
1.1.6	Beyond the Standard Model	18
1.1.6.1	CLFV in Supersymmetry	18
1.1.6.2	Two Higgs Doublet model	19
1.1.6.3	Leptoquark models	19
1.1.6.4	Additional Neutral Gauge Boson	19
1.2	Experiments searching for CLFV	20
1.2.1	μ Channels	20
1.2.1.1	$\mu^+ \rightarrow e^+ \gamma$	22
1.2.1.1.1	The MEG Experiment	23
1.2.1.1.2	The MEG II experiment	24
1.2.1.2	$\mu^+ \rightarrow e^+ e^- e^+$	25
1.2.1.2.1	SINDRUM I	26
1.2.1.2.2	The Mu3e Experiment	26
1.2.1.3	$\mu^- N \rightarrow e^- N$	27
1.2.1.3.1	SINDRUM II	29
1.2.1.3.2	Mu2e	30
1.2.1.3.3	DeeMe	30
1.2.1.3.4	COMET	30
1.2.2	τ Channels	31
2	The Mu2e Experiment	33
2.1	Experiment concept	33
2.2	Signal and backgrounds	34
2.2.1	Conversion Electron signal	34

2.2.2	Backgrounds	34
2.2.2.1	Cosmic rays	35
2.2.2.2	Intrinsic backgrounds	35
2.2.2.2.1	Muon Decay-in-Orbit	35
2.2.2.2.2	Radiative Muon Capture	36
2.2.2.3	Prompt processes	37
2.2.2.3.1	Radiative Pion Capture	37
2.2.2.3.2	π -DIF and μ -DIF	38
2.2.2.3.3	Beam electrons	38
2.2.2.4	Delayed processes from antiprotons	38
2.2.3	Background estimates and signal sensitivity	39
2.3	Experimental setup	40
2.4	Accelerator system and the proton beam	40
2.4.1	Pulsed proton beam	40
2.4.2	Proton beam extinction and Extinction Monitor	43
2.4.3	Production Target	44
2.5	Solenoids	44
2.5.1	The Production Solenoid	46
2.5.2	The Transport Solenoid	46
2.5.3	The Detector Solenoid	47
2.6	Stopping Target	48
2.7	The Mu2e detectors	50
2.7.1	The straw tracker	51
2.7.1.1	Momentum scale calibration	52
2.7.2	The electromagnetic calorimeter	52
2.7.2.1	The calorimeter mechanical structure	54
2.7.2.2	The undoped Cesium Iodide crystals	55
2.7.2.3	The calorimeter readout chain	56
2.7.2.3.1	Silicon PhotoMultipliers (SiPM)	56
2.7.2.3.2	The Front-End electronics	56
2.7.2.3.3	The Mezzanine board	57
2.7.2.3.4	The digitizer Board	58
2.7.2.3.5	The crate	58
2.7.2.4	Module-0 prototype	58
2.7.3	The Cosmic Ray Veto	58
2.7.4	The Muon Beam Stop and the Stopping Target Monitor	60
3	The straw tracker	63
3.1	Drift tubes	63
3.1.1	Gas ionization	64
3.1.2	Drift of ions and electrons	64
3.1.3	Avalanche multiplication	66
3.1.3.1	Avalanche gain	66

3.1.3.2	Quench gas	67
3.1.3.3	Operation modes of gaseous ionization detectors . . .	68
3.1.4	Signal creation and propagation	69
3.2	The Mu2e straw tracker	69
3.2.1	The sensitive unit: the straw	69
3.2.2	The building blocks: from panels to the station	71
3.2.3	The entire straw tracker detector	73
3.2.4	The drift tubes in magnetic field	73
3.2.5	The tracker front-end Electronics	75
3.2.5.1	The preamplifiers	77
3.2.5.2	The Digitizer Readout and Assembler Controller board	77
3.2.5.3	The tracker data format	78
3.2.5.4	The Data Acquisition System	79
3.2.6	Tracker performance requirements	80
4	Commissioning of the tracker DAQ and FEE	81
4.1	The test stand setup	81
4.2	Validation of ROC readout and buffering	84
4.2.1	Development of the ROC bit-level simulation	84
4.2.2	The <i>overflow</i> readout mode	85
4.2.2.1	Hit timing distribution and channels occupancy . . .	85
4.2.2.2	Number of hits	89
4.2.3	The <i>regular</i> readout mode	89
4.2.3.1	Time distribution and occupancy	89
4.2.3.2	Number of hits	90
4.3	Study of preamplifier performance	91
4.3.1	Test 1: channel occupancy versus channel ID	92
4.3.2	Test 2: analysis of the readout pulses waveforms	95
4.3.2.1	Estimate of the waveform baseline	96
4.3.2.2	Inverted waveforms	96
4.3.2.3	The waveform charge and pulse height	97
4.3.2.4	Channels response uniformity	102
5	First steps towards the station calibration	105
5.1	Overview of the timing calibration	105
5.1.1	Operational constraints of the station	106
5.2	Cosmic muons as a calibration source	107
5.2.1	Cosmic muons energy and angular distribution	108
5.2.2	Monte Carlo cosmic generation with CRY	109
5.2.3	Monte Carlo sample and coordinate system	109
5.3	Monte Carlo event selection	111
5.3.1	Panel illumination pattern	111
5.4	Reconstructing muon tracks	112
5.5	Results	115

5.6	Conclusions	118
6	Pre-pattern recognition studies	120
6.1	δ -electrons as source of background	120
6.1.1	Compton scattering	121
6.1.2	Pair production	123
6.1.3	Delta rays	124
6.2	Monte Carlo samples	125
6.3	δ -electrons in the Mu2e tracker	126
6.4	δ -electrons flagging algorithms	128
6.4.1	The <i>FlagBkgHits</i> algorithm	129
6.4.2	The <i>DeltaFinder</i> algorithm	131
6.4.2.1	Step 1: identifying δ -electron segments	131
6.4.2.2	Step 2: connecting the <i>seeds</i>	132
6.4.2.3	Step 3: identifying proton candidates	133
6.5	Performance analysis and comparison	133
6.5.1	Hit-level comparison	134
6.5.2	High-level comparison	139
6.6	Conclusions	142
7	Conclusions	144
7.1	Commissioning of the tracker DAQ and FEE	144
7.2	First steps towards the station calibration	146
7.3	Pre-pattern recognition studies	147
A	Charged particle in a magnetic field	149
B	Simulation and Analysis tools	151
B.1	art and FHiCL	151
B.2	STNTUPLE and ROOT	152
C	Mu2e Offline event reconstruction	153
C.1	Hits reconstruction and pre-filtering	153
C.1.1	Tracker hit reconstruction and pre-filtering	153
C.1.2	Calorimeter Hit Reconstruction	154
C.2	Helix search	154
C.2.1	Time Clustering	155
C.2.2	Pattern Recognition	156
C.2.2.1	TrkPatRec: Tracker-only Pattern Recognition	156
C.2.2.1.1	Reconstruction in the x - y plane	156
C.2.2.1.2	Reconstruction on $\phi - z$ plane	156
C.2.2.2	CalPatRec: Calorimeter+Tracker Pattern Recognition	158
C.3	Kalman Fitting	158
D	Panels illumination	160

List of Figures

1.1	Elementary particles of the Standard Model.	12
1.2	$\mu \rightarrow e\gamma$ and $\mu N \rightarrow eN$ Feynman diagrams.	17
1.3	SUSY contribution to $l_i \rightarrow l_j\gamma$	19
1.4	Sensitivity of $\mu \rightarrow e\gamma$ and $\mu N \rightarrow eN$ experiments.	23
1.5	Branching ratio limits in muon rare decay modes.	24
1.6	MEG detector.	25
1.7	MEG II detector (I).	25
1.8	MEG II detector (II).	26
1.9	SINDRUM experiment.	27
1.10	Mu3e experiment.	28
1.11	SINDRUM II experiment.	29
1.12	DeeMe experiment.	30
1.13	COMET experiment.	31
1.14	Tau lepton-flavour-violating branching ratio upper limits.	32
2.1	Decay-in-orbit spectrum on aluminium.	36
2.2	The electron energy spectrum near to the endpoint.	36
2.3	Mu2e simulated signal.	39
2.4	The Mu2e apparatus.	41
2.5	The pulsed proton beam delivery.	41
2.6	Proton beam profile.	42
2.7	The Mu2e beam timing	42
2.8	The Extinction Monitor location.	43
2.9	The Extinction Monitor.	44
2.10	The Production Target design.	44
2.11	The muon beamline.	45
2.12	The cross-section of the production solenoid.	46
2.13	The Transport Solenoid and the collimators.	47
2.14	The design of the collimators in the TS.	47
2.15	The structure of the Detector Solenoid coils and cryostat.	48
2.16	The structure of the inner and outer proton absorber.	48
2.17	The Stopping Target design.	49
2.18	The muonic atom mean lifetime and muonic free decay fraction. The dependence of the electron energy spectrum endpoint on the DIO. . .	50

2.19	The momentum resolution of the straw tracker.	51
2.20	The momentum calibration scheme.	53
2.21	The Mu2e calorimeter disks.	54
2.22	CAD representation of the calorimeter disks.	55
2.23	The calorimeter components.	55
2.24	The breakout of calorimeter mechanical components.	56
2.25	The design of SiPM cells.	57
2.26	The calorimeter FEE board.	57
2.27	The calorimeter Mezzanine Board and FEE.	58
2.28	The energy and time resolution of the calorimeter Module-0.	59
2.29	The CRV features and components.	60
2.30	The cross-section of a CRV module.	60
2.31	The Stopping Target Monitor geometry.	61
2.32	The muon stops and nuclear captures detection in Al ST.	62
2.33	The Stopping Target Monitor geometry.	62
3.1	Schematic view of a drift tube.	63
3.2	Electron drift velocity versus electric field in Ar:CO ₂ mixtures.	65
3.3	The model of an ionisation avalanche forming at the anode wire of a proportional tube.	66
3.4	The dependence of output signal of a counting tube on the applied voltage in gaseous ionisation detectors.	68
3.5	The straw tracker components.	70
3.6	A Mu2e straw tube.	70
3.7	The straw termination.	71
3.8	The straw arrangement within a panel.	72
3.9	One of the fully assembled panel.	72
3.10	One tracker plane fully assembled.	73
3.11	The total simulated drift time versus the initial radial distance from the sense wire.	74
3.12	Signal flow through front-end electronics.	75
3.13	Overview of the straw tracker front-end electronics.	76
3.14	The DRAC board schematics.	77
3.15	The Mu2e Data Acquisition system architecture.	80
4.1	Block diagram representation of the tracker DAQ test stand	82
4.2	The tracker DAQ test stand.	82
4.3	Graphic illustration of pulses in an Event Window.	83
4.4	Delay between channel 0 and the reference channel in digi-FPGA-1.	85
4.5	Delay between channel 44 and the reference channel in digi-FPGA-2.	85
4.6	The time distribution of hits.	86
4.7	The occupancy histogram (RUN281).	87
4.8	The distribution of the number of hits in channel 0.	88
4.9	The distribution of the total number of hits readout per event.	89

4.10	The hit time distribution.	90
4.11	The occupancy histogram (RUN105038).	91
4.12	The distribution of the number of hits in channel 0.	91
4.13	The distribution of the total number of hits per event in <i>regular</i> mode.	92
4.14	The regular occupancy plot.	93
4.15	The occupancy plot (dead channel and more hits than expected).	94
4.16	Δt distribution between hits in the 70th channel.	94
4.17	The occupancy plot and cross-talks.	95
4.18	The zoomed view of the Analog Mother Board with the preamp boards installed.	95
4.19	A regular waveform.	96
4.20	The fitted baseline distribution of channel 0.	97
4.21	The fitted baseline distribution of channel 8.	97
4.22	An example of waveform with dips.	98
4.23	An inverted waveform.	98
4.24	The (positive) charge distribution of the waveforms.	99
4.25	The pulse height distribution of the waveforms.	99
4.26	The first sample distribution of the waveforms.	100
4.27	The negative charge distribution of the waveforms.	100
4.28	The 2D distribution of pulse height versus (positive) charge.	101
4.29	The distribution of the waveform s_{mean}	101
4.30	2D distribution of (positive) charge versus s_{mean}	102
4.31	2D distribution of pulse height versus s_{mean}	102
4.32	A simulation of the charge and pulse height distribution behaviour.	103
4.33	Glitch waveform of channel 0.	103
4.34	Noisy waveform of the 36th channel.	104
4.35	A first example of Data Quality Monitoring.	104
5.1	Schematic view of the valves located in a tracker panel.	107
5.2	The relation of θ^* to θ	109
5.3	Schematic view of a tracker station and the coordinate system.	110
5.4	Monte Carlo illumination across the panel (panel 0, plane 0).	112
5.5	The $y - z$ direction distribution of muons satisfying 4/4 requirement.	113
5.6	Schematic view of a cosmic muon hitting the vertical oriented station.	114
5.7	The reconstructed longitudinal coordinate in the 0th panel frame.	115
5.8	The bias between the reconstructed and the true hit coordinate.	116
5.9	The reconstructed $y - z$ direction distribution of cosmic rays.	117
5.10	The 2D histogram of the longitudinal bias versus the true coordinate.	117
5.11	The profile of the longitudinal bias versus the true longitudinal coordinate.	118
6.1	The Compton effect.	121
6.2	The pair production.	123

6.3	Monte Carlo momentum distribution of particles producing hits in the Mu2e tracker ($CE - 1BB$ data sample).	128
6.4	Monte Carlo momentum distribution of particles producing hits in the Mu2e tracker ($PBAR - 0BB$ data sample).	129
6.5	Before and After background hits flagging.	130
6.6	δ -electrons and CE patterns in $r - z$ plane.	131
6.7	A δ candidate <i>seed</i>	132
6.8	Monte Carlo deposited energy distribution in the tracker.	133
6.9	The distribution of the total and flagged number of muon <i>ComboHits</i> as a function of the particle momentum.	135
6.10	The electron and positron energy distribution for $E < 2$ MeV.	137
6.11	The electrons (blue) and positrons (red) hit flagging efficiency versus particle momentum in the low-momentum range.	138
6.12	Momentum distribution of reconstructed tracks.	139
6.13	<i>TimeCluster</i> on $time - z$ plane.	140
6.14	<i>TimeCluster</i> on $time - z$ plane.	141
B.1	Mu2e simulation and data handling.	152
C.1	Procedure adopted to search for the center of the $x - y$ projection of the helix.	157
C.2	Sketch of the resolution of the 2π ambiguity.	157
C.3	Combinatorial background reduction achieved by exploiting the calorimeter clusters seeding.	158
C.4	The symmetry of the straw generates an ambiguity for the hits.	159
D.1	Plane 0 illumination pattern of cosmic hits.	160
D.2	Plane 1 illumination pattern of cosmic hits.	161

Chapter 1

Charged Lepton Flavour Violation

This Chapter offers a concise overview of the fundamental theoretical and experimental components essential to understand the objectives of the Mu2e experiment at Fermilab and the work conducted for this Thesis. The introduction to the Standard Model and certain extensions serves to justify the investigation of Charged Lepton Flavour Violation (CLFV), since it would be a clear signal of new physics beyond the current theories. Fundamental bibliography for this Chapter can be found in [4] and [5].

1.1 Theoretical introduction

1.1.1 The Standard Model

The Standard Model (SM) provides an excellent description of elementary particles and their interactions. It describes three out of four fundamental interactions known to this day: electromagnetism, weak and strong interactions. The SM is based on the gauge symmetry group $U(1)_Y \times SU(2)_L \times SU(3)_C$. The first two terms describe the electroweak interaction, Y indicates the hypercharge and L refers to the fact that this acts only on the left handed components of the fields. The last term describes the strong interaction, and C indicates the color charge. The SM describes the elementary fields shown in Fig. 1.1.

1.1.1.1 Bosons

A boson is a particle with zero or integer spin, which follows Bose-Einstein statistics. There are twelve fundamental bosons, that mediate interactions: γ , Z , W^\pm for the electroweak interaction and eight gluons mediating the strong interaction. The Higgs field is a complex scalar weak isospin doublet that is responsible for the mechanism through which fermions and bosons acquire mass and also explains the origin of




Figure 1.1: Elementary particles of the SM.

$U(1)_{EM}$ through the spontaneous symmetry breaking of the $SU(2)_L \times SU(1)_Y$ in the electroweak sector. Also mesons, that are composed of a quark and antiquark pair, are bosons. Bosons can be either massive, as Z , Higgs and W^\pm , or massless, as the photon and gluons.

1.1.1.2 Fermions

A fermion is a particle with a half-integer spin, i.e. $1/2, 3/2$, and follows the Fermi-Dirac statistics which is based on the Pauli exclusion principle. Fermions are divided in two categories, leptons and quarks, depending on the forces through which they interact. The arrangement of fermions into three generations is dictated by the mass property with more massive particles assigned to higher generations, as depicted in Fig.1.1.

Particles of second and third generations are not stable and decay into first-generation particles. Leptons do not interact strongly because they are not color charged, so they interact only via weak and electromagnetic interaction. They are categorized into two groups based on the electric charge: e, μ, τ are the charged leptons and ν_e, ν_μ, ν_τ are the neutral ones. These particles form doublets of flavour. Neutrinos, due to their neutral nature, interact only weakly so their detection is extremely challenging. The six quarks participate in all known interactions. The known quarks are, in ascending order of mass and generation: up (u) and down (d), strange (s) and charm (c), bottom (b), top (t), and their antiparticles. They primarily interact with each other through the strong force by gluons exchange. Free quarks have never been observed due to confinement, since they carry color charge. Confinement of quarks is a fundamental aspect of quantum chromodynamics (QCD),

the theory of strong interactions. Quarks combine to form color-neutral particles known as hadrons, classified into baryons and mesons. Baryons consist of an odd number of quarks, while mesons, as mentioned in Subsection 1.1.1.1, are composed of a quark and an antiquark. Since quarks have weak electric charge and isospin, they can interact with each other and other fermions through weak and electromagnetic interactions.

1.1.2 History of flavour

The concept of flavour, namely the presence of three duplicates for every family of elementary fermions, was implemented within the SM by introducing three copies of the gauge representations of fermion fields, without any fundamental requirement. This view began to take shape in the late 1940s. The origin can be traced back to the experiment conducted by Conversi, Pancini and Piccioni in 1947. Their experiment revealed that negative muons, called at that time *mesotrons*, did not undergo nuclear capture as expected. They decayed in electrons, similarly to positive muons, therefore they could not be Yukawa particles. In the same year, Powell and his group identified a two-step decay process ($\pi \rightarrow \mu \rightarrow e$), distinguishing the pion from the muon. Bruno Pontecorvo suggested that the muon could be a sort of *isomer* of the electron, leading to the idea of a second generation of elementary fermions. Rochester and Butler discovered unusual events in cosmic ray pictures, later identified as *V-particles* (later discovered that they originated from neutral kaons). This was the first hint of the existence of a second generation of quarks. In 1945 the search for decay $\mu \rightarrow e\gamma$ began. This decay was not found, leading to the principle of lepton conservation. The first negative result was reported by Hincks and Pontecorvo in 1948. On the hadron side, the second generation of quarks was established in the mid-70s, involving the GIM mechanism and the discovery of the charm quark.

Meanwhile, on the leptonic side, the upper limit on the branching ratio of $\mu^+ \rightarrow e^+\gamma$ decay was set in 1955. The discovery of parity violation in the late 1950s suggested the weak interaction is mediated by bosons. Feinberg hypothesized that if the bosons existed, $\mu^+ \rightarrow e^+\gamma$ could occur at a level of 10^{-4} , through a loop with neutrino and a boson. This led to the two-neutrino hypothesis, suggesting that the neutrino coupled to the muon differs from that coupled to the electron, thereby prohibiting $\mu^+ \rightarrow e^+\gamma$. The existence of two different neutrinos was verified at Brookhaven National Laboratory, by confirming muon production, and no electron production, from the scattering of neutrinos coming from pion decays that produced only muons.

After the observation of CP violation in neutral kaon decay, a third generation of quarks was hypothesized. After the discoveries of the τ (1976), the b quark (1977), t quark (1995) and ν_τ (2000), a complete picture was achieved and the concept of flavour was consolidated in the SM.

1.1.3 Lepton sector in the Standard Model

In the SM, only one Higgs field Φ exists. The fermions masses and the mixing term arise from the couplings of fermions with Higgs field. In the following, I will call the left-handed quark doublets as $Q_L^T = (u_L \ d_L)$, u_R and d_R the right-handed quarks up and down respectively, the left-handed lepton doublets as $L_L^T = (\nu_L \ e_L)$, and e_R the right-handed leptons. There is no right-handed neutrinos. The Yukawa coupling of fermions with the Higgs field \mathcal{L}_Y is the sum of two terms: \mathcal{L}_e that describes the leptonic component (Eq.1.1) and \mathcal{L}_q that describes the quark one (Eq.1.2).

$$-\mathcal{L}_e = (Y_e)_{ij} \bar{L}_{Li} e_{Rj} \Phi + \text{h.c.} \quad (1.1)$$

$$-\mathcal{L}_q = (Y_u)_{ij} \bar{Q}_{Li} u_{Rj} \tilde{\Phi} + (Y_d)_{ij} \bar{Q}_{Li} d_{Rj} \Phi + \text{h.c.} \quad (1.2)$$

where the term Y_f ($f = u, d, e$) describes the 3×3 Yukawa complex matrices. $\tilde{\Phi} \equiv i\tau_2 \Phi^*$ is the conjugate Higgs field. The mass terms of fermions, characterized by the $m_f \bar{f}_L f_R$ form, originate from the breaking of the $SU(2)_L \times U(1)$ symmetry caused by the non-zero vacuum expectation value of the Higgs field.

$$\langle \Phi \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix} \quad v \simeq 246 \text{ GeV} \quad (1.3)$$

The fermion mass is given by:

$$(m_f)_{ij} = \frac{v}{\sqrt{2}} (Y_f)_{ij} \quad f = u, d, e \quad (1.4)$$

Since a right-handed neutrino does not appear in the Lagrangian, neutrinos have no mass, as it is formulated in the SM. The Yukawa matrices can be diagonalized through unitary rotations of the fields, as it follows:

$$Y_f = V_f \hat{Y}_f U_f^\dagger \quad f = u, d, e \quad (1.5)$$

where \hat{Y}_f is the diagonal Yukawa matrix. It is possible to label fermions in the rotated basis as f' , so $f_L = V_f f'_L$ and $f_R = V_f f'_R$. Since V_f and U_f are unitary, the rotation will not affect the neutral interactions term and the kinetic terms, as we can see in $\bar{f}_L \gamma^\mu f_L = \bar{f}'_L \gamma^\mu (V_f^\dagger V_f) f_L = \bar{f}'_L \gamma^\mu f'_L$. The coupling between fermion and Higgs boson will be:

$$-\mathcal{L}_{h\bar{f}f} = \frac{(\hat{m}_f)_{ij}}{v} \bar{f}'_{Li} f'_{Rj} h + \text{h.c.} \quad f = u, d, e \quad (1.6)$$

where \hat{m}_f denotes the diagonalized mass matrix: it is clear that there are no flavour-violating terms. In the quark sector flavour-violation arises from the rotations in Eq.1.5 in the charged-current interactions with the W bosons:

$$\begin{aligned} \mathcal{L}_{CC} &= \frac{g}{\sqrt{2}} (\bar{u}_L \gamma^\mu d_L + \bar{\nu}_L \gamma^\mu e_L) W_\mu^+ + \text{h.c.} \\ &= \frac{g}{\sqrt{2}} (\bar{u}'_L \gamma^\mu (V_u^\dagger V_d) d'_L + \bar{\nu}'_L \gamma^\mu (V_\nu^\dagger V_e) e'_L) W_\mu^+ + \text{h.c.} \end{aligned} \quad (1.7)$$

The violation comes from the fact that $V_u \neq V_d$. The mixing is controlled by the Cabibbo-Kobayashi-Maskawa matrix $V_{CKM} \equiv V_u^\dagger V_d$, [6]. Meanwhile, as the lepton sector contains massless neutrinos, V_ν can be arbitrarily chosen as $V_\nu = V_e$. From the preceding discussions, it becomes evident that in the SM with massless neutrinos, there is no occurrence of LFV in any form. The Lagrangian \mathcal{L}_Y is invariant under three independent global $U(1)$ rotations, resulting in the conservation of three lepton family numbers: L_e , L_μ and L_τ . However, if the Yukawa coupling Lagrangian includes supplementary terms involving the lepton fields, flavour violation can occur in the lepton sector. Examples of such additions could be a neutrino mass term or a second Higgs doublet.

1.1.4 Overview of Charged Lepton Flavour Violation

There are three different lepton flavours: the electron-lepton L_e , the muon-lepton L_μ and the tau-lepton L_τ . In Table 1.1, the quantum numbers assigned to each lepton are displayed.

Lepton	L_e	L_μ	L_τ
e^-/e^+	+1 / -1	0	0
$\nu_e/\bar{\nu}_e$	+1 / -1	0	0
μ^-/μ^+	0	+1 / -1	0
$\nu_\mu/\bar{\nu}_\mu$	0	+1 / -1	0
τ^-/τ^+	0	0	+1 / -1
$\nu_\tau/\bar{\nu}_\tau$	0	0	+1 / -1

Table 1.1: Lepton numbers assigned to neutrinos and charged leptons.

In the SM defined with massless left-handed neutrinos, Lepton flavour (LF) is a conserved quantity, [7]. Experimental observations have demonstrated that, as they travel, neutrinos exhibit flavour oscillations, which implies that they must have non-zero masses and mixing angles. This phenomenon represents also a violation of the conservation of the lepton flavour. The SM, while successful in many aspects, fails to explain phenomena like neutrino masses and the consequent flavour oscillations. Since neutrinos get their masses through renormalizable Yukawa interactions with the Higgs, the predicted CLFV transitions are suppressed by sums over $(\Delta m_{ij}^2/M_W^2)^2$, as calculated in [8] and as shown in Section 1.1.5, where Δm_{ij}^2 is mass-squared difference between the neutrino mass eigenstates i, j and M_W is the W boson mass. The neutrino mass difference is very small ($\Delta m_{ij}^2 \leq 10^{-3} \text{ eV}^2$)

with respect to the W boson mass, and the expected branching ratios below 10^{-50} are well beyond the experimental reach. Experimental studies of the lepton flavour violating process could open a window to new physics. Moreover, lepton flavour constitutes an accidental symmetry within the SM, not related to the Gauge structure of the theory but coming from its particle content, especially from the absence of right-handed neutrinos.

Minor deviations from the SM can give rise to extra occurrences of lepton flavour violation, leading to notable rates of CLFV. There are various extensions of the SM that could potentially be examined in the upcoming experimental searches for CLFV, that I will report in the next sections.

1.1.5 CLFV in the SM with massive neutrinos

The first evidence against the hypothesis of massless neutrinos emerged with the solar neutrino problem. In the 1960s, the solar neutrino detection experiment at Homestake revealed that the observed number of solar neutrinos, generated by fusion in the Sun, was significantly lower than the anticipated value based on the standard solar model, given that the detector was only sensitive to ν_e , [9]. Consistent results were replicated in subsequent experiments using radiochemical and Cherenkov detectors, discovering neutrino oscillations. These oscillations firmly established non-zero neutrino masses. The lepton flavour-violating neutrino oscillations showed that the global $U(1)$ symmetries associated with the lepton family numbers are not fundamental symmetries. A correction to SM is needed to include neutrino mass terms. This is possible adding a right-handed neutrino singlet ν_R or some non-renormalizable operators.

In the first case, an additional term $\Delta\mathcal{L}_D$ to the Yukawa coupling should be introduced:

$$-\Delta\mathcal{L}_D = (Y_\nu)_{ij} \bar{\nu}_{Ri} \tilde{\Phi}^\dagger L_{Lj} + \text{h.c.} \quad (1.8)$$

Similarly to other fermions, a Dirac mass term $m_\nu \bar{\nu}_L \nu_R$ is generated through symmetry breaking:

$$(m_\nu^D)_{ij} = \frac{v}{\sqrt{2}} (Y_\nu)_{ij} \quad (1.9)$$

In this case, the small neutrino masses can be explained only if a very small term Y_ν is considered ($\leq 10^{-12}$), [5]. This brings us considering the second hypothesis: adding non-renormalizable operators can introduce Majorana masses for left-handed neutrinos alone. The corresponding $\Delta\mathcal{L}_M$ can be written as:

$$-\Delta\mathcal{L}_M = \frac{1}{2} (m_\nu^M)_{ij} \bar{\nu}_{Li}^C \nu_{Lj} + \text{h.c.} \quad (1.10)$$

where $\bar{\nu}_L^C$ is the charge-conjugated fields. This term violates lepton number and requires an operator of dimension 5, [10], to be consistent with SM symmetries. A minimal Lagrangian is given by:

$$-\Delta\mathcal{L}_{M \text{ eff}} = \frac{C_{ij}}{\Lambda} \left(\bar{L}_{Li}^C \tau_2 \Phi \right) \left(\Phi^T \tau_2 L_{Lj} \right) + \text{h.c.} \quad (1.11)$$

where the Λ term represents a mass scale characteristic of extra degrees of freedom and the C_{ij} is an antisymmetric charge conjugation matrix. The corresponding Majorana mass term is:

$$(m_\nu^M)_{ij} = \frac{C_{ij}v^2}{\Lambda} \quad (1.12)$$

In this case, the small neutrino masses can be explained only if $\Lambda \gg v$: this seems to appear more natural than the previous case. No matter how the extra neutrino mass factor is expressed precisely, a physical basis diagonalizing the mass matrix is determined, resulting in $V_\nu \neq V_e$ in Equation 1.7. Lepton mixing is described by $U_{PMNS} \equiv V_\nu^\dagger V_e$, which is similar to the CKM matrix. The Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix is the term that is typically used to refer to it. As it is in the basis diagonalizing charged lepton masses and it diagonalizes the neutrino mass matrix, U_{PMNS} also explains the mixing between neutrino flavour eigenstates ν_α and mass eigenstates ν_i :

$$\nu_\alpha = \sum_{i=1,2,3} (U_{PMNS})_{li} \nu_i \quad l = e, \mu, \tau \quad (1.13)$$

In addition to the neutral Lepton Flavour Violation (LFV) observed in neutrino oscillations, the mixing outlined by U_{PMNS} can in principle give rise to CLFV processes, i.e. LFV that involves charged leptons. The new Feynman diagrams are loops involving neutrinos and W bosons, as $\mu \rightarrow e\gamma$ in Fig.1.2a and $\mu N \rightarrow eN$ in Fig.1.2b.




Figure 1.2: $\mu \rightarrow e\gamma$ and $\mu N \rightarrow eN$ Feynman diagrams.

In the SM, each of these mechanisms is significantly suppressed. Using the example of $\mu \rightarrow e\gamma$, the branching ratio of this process may be computed as follows:

$$\begin{aligned} BR(\mu \rightarrow e\gamma) &= \frac{3\alpha}{32\pi} \left| \sum_{i=2,3} U_{\mu i}^* U_{e i} \frac{\Delta m_{1i}^2}{M_W^2} \right|^2 \\ &= \frac{3\alpha}{32\pi} \left(\frac{1}{4} \right) \sin^2 2\theta_{13} \sin^2 \theta_{23} \left| \frac{\Delta m_{13}^2}{M_W^2} \right|^2 \end{aligned} \quad (1.14)$$

where α is the fine structure constant, $U_{\mu i}$ and $U_{e i}$ are corresponding elements in the PMNS matrix, Δm_{1i}^2 are the neutrino squared mass differences, M_W is the W

boson mass and θ_{13} and θ_{23} are rotating angles in PMNS matrix parametrization, [8].

As mentioned in 1.1.4, the expression yields $BR(\mu \rightarrow e\gamma) \sim \mathcal{O}(10^{-54})$. The big discrepancy in mass between neutrinos and the W boson results in an extraordinarily small value for $|\Delta m_{13}^2/M_W|$.

The rates predicted by the SM are extremely small, making them impractical for detection in any experiment. On the other hand, numerous Beyond the Standard Model (BSM) theories incorporate mechanisms that substantially amplify CLFV rates, a topic to be addressed in the subsequent section. Any experimental observation of any CLFV process would unequivocally indicate physics beyond the SM.

1.1.6 Beyond the Standard Model

Numerous BSM theories propose mechanisms that could contribute to CLFV processes, potentially yielding experimentally detectable rates. Here, we highlight a selection of BSM theories known for their CLFV contributions. It is important to note that this list is not comprehensive; for further studies, additional reviews are available in [5] and [7].

1.1.6.1 CLFV in Supersymmetry

Supersymmetry (SUSY) is a theoretical framework that has oriented experiments in the CLFV research for many years. On one hand, models with SUSY broken at energies close to electro-weak scale have given solution to the hierarchy problem, i.e. how to maintain the Higgs mass significantly smaller than the Planck scale ($\sim 10^{19}$ GeV). In this framework, each elementary particle has a superpartner, with the same quantum numbers except for spin: a boson is the superpartner of a fermion and vice versa. A superpartner of a lepton is called *slepton*. If there is no common eigenstate base between lepton and *slepton*'s mass matrices then a physical *slepton* will be a superposition of flavours. On the other hand, the suppression of CLFV processes is due to the wide separation of the neutrinos and W masses, which can be mitigated by introducing SUSY partners of neutrinos and W bosons. This suggests that CLFV processes should have been observable earlier, unless SUSY breaking occurs at or near the electroweak scale ($\sim 10^2$ GeV), [5].

In this case a loop diagram can lead to CLFV, as shown in Fig.1.3. Despite the similar topology to that of the SM contribution (Fig.1.2a), the typical SUSY mass is expected to be much higher than that of the neutrinos. Predictions for the branching ratio of this process vary among different SUSY models, contingent upon specific mechanisms and particle masses. These rates can be significantly enhanced. For example, in an $SU(5)$ SUSY grand unified theory, the computed branching ratio could reach $\mathcal{O}(10^{-14})$ for a slepton mass on the order of $\mathcal{O}(10^2 \text{ GeV}/c^2)$, a value measurable by upcoming experiments, [5].




Figure 1.3: SUSY contribution to $l_i \rightarrow l_j \gamma$, through *sleptons* mass mixing, [7].

1.1.6.2 Two Higgs Doublet model

Although the SM incorporates only one Higgs boson, there are no constraints against the presence of additional Higgs fields. One straightforward example of a comprehensive theory featuring multiple Higgs fields is the two Higgs doublet model (2HDM), where two Higgs bosons exist, each interacting with fermions and possessing a vacuum expectation value, [11].

Generally, the Lagrangian incorporating extra Higgs fields post-electroweak symmetry breaking can be written as:

$$-\mathcal{L} = m_i \bar{f}_{Li} f_{Ri} + (Y^\alpha)_{ij} \bar{f}_{Li} f_{Rj} h^\alpha + \text{h.c.} + \dots \quad (1.15)$$

Non-renormalizable terms of higher dimensions are omitted. $(Y^\alpha)_{ij}$ represents couplings to a single scalar field and the contributions from different Higgses are summed over. The non-zero-off-diagonal terms in $(Y^\alpha)_{ij}$ give rise to flavour violating Yukawa couplings. From accelerator and precision experiments, constraints on the off-diagonal coupling of the 125 GeV Higgs boson can be obtained.

1.1.6.3 Leptoquark models

Leptoquarks (LQs) are hypothetical particles initially proposed within the Pati-Salam model, [12]. A LQ possesses both a baryon number (B) and a lepton number (L). In various LQ models the quark and lepton sectors are unified. This unification allows for direct coupling between quarks and leptons via the exchange of LQs. Consequently, specific CLFV processes like $K_L^0 \rightarrow e\mu$ and $\mu N \rightarrow eN$ are mediated by LQs. Constraints on LQ models arise from both collider experiments and rare decay searches. Direct searches at ATLAS and CMS have already excluded scalar LQs of the first and second generations with masses below ~ 1 TeV.

Additionally, LQ models must satisfy constraints related to proton stability, as some models involve LQs that could mediate proton decay. To mitigate this, the corresponding LQs must either have extremely high masses or their related couplings must be exceedingly small, [13].

1.1.6.4 Additional Neutral Gauge Boson

Grand unified theories (GUTs) are constructed based on extended Gauge groups in the pursuit of a more fundamental model. At lower energies, these extended

gauge groups break down to the direct product of the SM Gauge group $SU(3) \times SU(2) \times U(1)$ along with an additional $U(1)$ factor. The neutral Gauge boson associated with this $U(1)$ group can mix with the original SM neutral Gauge boson, resulting in two mass eigenstates, namely Z and Z' . Additionally, extended Gauge theories require the introduction of additional fermion fields to cancel anomaly-free currents beyond those of $SU(5)$. These *new* fermions can mix with the known SM fermions possessing the same electric and color charges, consequently affecting their couplings with Gauge bosons. The appearance of off-diagonal terms in neutral current couplings to fermions can lead to flavour-changing couplings to Z and Z' . Certain CLFV processes, such as $\mu \rightarrow eee$ and $\mu - e$ conversion, receive tree-level contributions through intermediate Z and Z' bosons. Further insights into the phenomenology of the Z' boson can be found in, [14]. The search for Z' bosons is conducted using channels like $Z' \rightarrow \bar{f}f$ at hadron colliders. Lower limits on the Z' mass from various models are listed in [15], primarily falling within the low TeV range. Particularly, the limits reported in CLFV final states $e\mu$, $e\tau$ and $\mu\tau$ range in between 3.5 TeV and 4.5 TeV. Upper limits on the $Z \rightarrow l_1 l_2$ couplings of the SM Z boson are also provided in Table 1.2.

1.2 Experiments searching for CLFV

CLFV has not been observed yet, despite the ongoing efforts to detect such violation in different channels in both dedicated and general-purpose experiments.

Some of these efforts are documented in Table 1.2, which presents their respective experimental upper limits. In addition to searches at collider experiments, such as observing Z and Higgs decays, rare decay experiments play a significant role as a complementary method in the search for CLFV. Collider experiments enable the exploration of various CLFV channels simultaneously, including those involving τ leptons. However, complexities in data selection and reconstruction present challenges to improve the limits in this field. Rare decay experiments, by focusing on specific processes or a set of similar processes, can achieve high statistics using an intense particle beam. By suppressing backgrounds, these experiments can significantly improve the sensitivity. Furthermore, this approach enables the investigation of high mass scales, as will be explained in next sections.

1.2.1 μ Channels

The most promising channel is currently represented by muons. There are three primary CLFV channels involving muons: $\mu^+ \rightarrow e^+\gamma$, $\mu^- N \rightarrow e^- N$ and $\mu^+ \rightarrow e^+e^+e^-$. As seen in Table 1.2, these channels have the lowest branching ratio limits: muon production is highly favored in pion and kaon decays, which are easily produced through hadronic interactions, and sufficiently long muon lifetime enables the formation of muon beams. Muon beams are produced fairly easily at proton accelerator facilities. However, muons have small mass with respect to τ s, that results in a limited number of decay modes available for studies.

Reaction	Present limit	CL	Experiment	Year	Ref.
$\mu^+ \rightarrow e^+ \gamma$	7.5×10^{-13}	90%	MEG II	2024	[16]
$\mu^+ \rightarrow e^+ e^+ e^-$	1.0×10^{-12}	90%	SINDRUM	1988	[17]
$\mu^- \text{ Ti} \rightarrow e^- \text{ Ti}$	6.1×10^{-13}	90%	SINDRUM II	1998	[18]
$\mu^- \text{ Au} \rightarrow e^- \text{ Au}$	7.0×10^{-13}	90%	SINDRUM II	2006	[2]
$\mu^+ e^- \rightarrow \mu^- e^+$	8.3×10^{-11}	90%	SINDRUM	1999	[19]
$\tau \rightarrow e \gamma$	3.3×10^{-8}	90%	BaBar	2010	[20]
$\tau \rightarrow \mu \gamma$	4.4×10^{-8}	90%	BaBar	2010	[20]
$\tau \rightarrow e e e$	2.7×10^{-8}	90%	Belle	2010	[21]
$\tau \rightarrow \mu \mu \mu$	2.1×10^{-8}	90%	Belle	2010	[21]
$B^0 \rightarrow \mu e$	2.8×10^{-9}	90%	LHCb	2013	[22]
$B^0 \rightarrow \tau e$	2.8×10^{-5}	90%	BaBar	2008	[23]
$B^0 \rightarrow \tau \mu$	2.2×10^{-5}	90%	BaBar	2008	[23]
$K_L^0 \rightarrow \mu e$	4.7×10^{-12}	90%	BNL E871	1998	[24]
$K^+ \rightarrow \pi^+ \mu^+ e^-$	2.1×10^{-10}	90%	BNL E865	2005	[25]
$K_L^0 \rightarrow \pi^0 \mu^+ e^-$	4.4×10^{-10}	90%	KTeV	2008	[26]
$\pi^0 \rightarrow \mu e$	8.6×10^{-9}	90%	KTeV	2008	[26]
$\Upsilon(1s) \rightarrow \mu \tau$	6.0×10^{-6}	95%	CLEO	2008	[27]
$Z^0 \rightarrow \mu e$	1.7×10^{-6}	95%	LHC ATLAS	2014	[28]
$Z^0 \rightarrow \tau e$	1.7×10^{-6}	95%	LEP OPAL	1995	[29]
$Z^0 \rightarrow \tau \mu$	9.8×10^{-6}	95%	LEP DELPHI	1997	[30]
$h \rightarrow e \mu$	3.5×10^{-3}	95%	LHC CMS	2016	[31]
$h \rightarrow \tau \mu$	2.5×10^{-3}	95%	LHC CMS	2017	[32]
$h \rightarrow e \tau$	6.1×10^{-3}	95%	LHC CMS	2017	[32]

Table 1.2: Experimental upper limits for a variety of CLFV processes of leptons, mesons and heavy bosons, [5].

Muon rare decay studies can also provide theory differentiation power combining results of the three channels. We can consider the following model-independent

Effective Field Theory (EFT) Lagrangian of the $e\mu$ CLFV mode, [33]:

$$\begin{aligned} \mathcal{L}_{CLFV} = & \frac{m_\mu}{(1+\kappa)\Lambda^2} \bar{\mu}_R \sigma_{\mu\nu} e_L F^{\mu\nu} + \text{h.c.} + \\ & \frac{\kappa}{(1+\kappa)\Lambda^2} \bar{\mu}_L \gamma_\mu e_L \left(\sum_{q=u,d} \bar{q}_L \gamma^\mu q_L \right) + \text{h.c.} \end{aligned} \quad (1.16)$$

m_μ is the muon mass and $F^{\mu\nu}$ is the electromagnetic field tensor.

Λ represents the effective mass scale of the new degree of freedom, and κ determines the relative strengths of the two operators. The first term in the Lagrangian, called the dipole term, represents a magnetic moment type operator which directly mediates $\mu \rightarrow e\gamma$. The second one corresponds to a four-fermion interaction term, which mediates $\mu N \rightarrow eN$ at a tree level and the other two processes at one-loop level. The Mu2e experiment can probe the effective mass scale up to $\mathcal{O}(10^4 \text{ TeV})$ with its designed sensitivity assuming $\kappa \gg 1$. On the other hand, $\mu \rightarrow e\gamma$ search experiments are more sensitive when $\kappa \ll 1$; the dominant magnetic moment type term determines the other two processes have lower rates in such a case. In order to learn more about the new physics, one needs to combine information involving rates of different CLFV processes, [34]. The corresponding parameter space (the Λ - κ plane) is shown in Figure 1.4, [33].

Figure 1.5 shows how the branching ratio limits of muon rare decays have improved over the last several decades. The following paragraphs will discuss the experimental challenges and future perspectives for each of these channels. The next generation experiments aim to improve the search sensitivity by many orders of magnitude.

1.2.1.1 $\mu^+ \rightarrow e^+\gamma$

A clear signal of CLFV, in $\mu^+ \rightarrow e^+\gamma$ channel, is given by back-to-back electron and photon, each one with energy of 52.8 MeV, both produced simultaneously. Positive muons are preferred since the negative ones may undergo the nuclear capture. Muons are stopped and decay at rest. There are two most significant sources of background: Radiative Muon Decay (RMD) $\mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu\gamma$ and the accidental coincidence of $\mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ with a random γ generated by annihilation or bremsstrahlung. The first one is an in-time process, where neutrinos carry off a small part of the energy and it is only a small part of the total background. The accidental background is dominant. Since both backgrounds depend directly on the muon stop rate Γ_μ , a continuous beam is preferred. Moreover, since a higher number of stopped muons correspond to a lower statistical error, but at the same time, to a lower signal to background ratio, an optimal Γ_μ exists. The stopping target thickness must be optimized: a thin target is needed to minimise the multiple Coulomb scattering which affects the angular resolution of the outgoing electron, but it must be thick enough to stop a significant portion of incoming muons.




Figure 1.4: Sensitivity of $\mu \rightarrow e\gamma$ and $\mu N \rightarrow eN$ experiments to the new physics scale Λ as a function of κ as defined in Equation 1.16, [35]. The blue region is the New Physics phase space excluded by SINDRUM-II, [2]. The red region represents the limit set by MEG, [36], while the dashed red line represents the region that is excluded by MEG-II, [16]. The solid (dashed) blue line is the expected limit that would be set by Mu2e, [37].

1.2.1.1.1 The MEG Experiment The MEG experiment, [36], has been designed around two concepts: exploiting a liquid xenon detector (LXe) for positron and photon tracking and an anti-bottle magnetic field, [5]. The apparatus is shown in Figure 1.6. A polyethylene target is used to stop muons in the center of the magnet. The measured quantities are the electron and photon energies (E_e and E_γ) and the relative positions (angles $\theta_{e\gamma}$, $\phi_{e\gamma}$ and time $t_{e\gamma}$). A combination of a drift chamber (DCH) and plastic scintillator timing counters (TC) measure the positron momentum.

The photon energy and direction are measured in a volume of LXe with more than 800 photo-multiplier tubes. An energy resolution of less than 1% for each particle is required to distinguish background from the signal. In MEG, the magnetic field decreases uniformly from centre to periphery, pushing particles away from the centre. The exact shape of the field has been chosen to have a track radius proportional to the absolute momentum. This allows low energy positrons to be discarded by simply placing the detector far enough away from the magnet axis. This feature is unique to the MEG magnetic system and justifies its name as *CO*nstant *B*ending




Figure 1.5: History and outlook of branching ratio limits in muon (and τ s) rare decay modes [38].

RAdius (COBRA) magnets. The DCH spectrometer is composed by 16 trapezoidal drift chambers oriented radially and filled with He-C₂H₆. The timing from the DCH and TC is used to assess the radial coordinate, whereas the z location is determined by measuring the induced charged on the zig-zag shaped pads on the side of the drift chambers. The core momentum resolution for e^+ is ~ 330 keV. A liquid xenon scintillating detector was used for photon reconstruction to reduce passive material and improve temporal resolution. This option provides higher light yield than NaI crystals and has a substantially shorter decay period, with a timing resolution below 100 ps.

MEG collected 7.5×10^{14} stopped muons in 2008-2013, setting a limit of $BR(\mu^+ \rightarrow e^+\gamma) < 4.2 \times 10^{-13}$ at 90% CL, [36].

1.2.1.1.2 The MEG II experiment The MEG II detector is the upgrade of the MEG detector, [39]. MEG II was proposed to reduce the accidental background that could not be further reduced in MEG. Figure 1.7 shows the apparatus. The muon flux was increased up to $7 \times 10^7 \mu^+/s$ and a thinner but more inclined stopping target was installed to reduce the multiple scattering and bremsstrahlung while keeping the same stopping power ($205 \rightarrow 140 \mu m$). The old drift chamber was replaced with a new cylindrical drift chamber (CDCH) designed to have higher granularity and transparency and made of 9 layers of drift cells to improve positron track reconstruction. CDCH is shown in Figure 1.8. A more segmented system (pixellated-TC) was adopted instead of plastic scintillator timing counters (TC). A Radiative Decay Counter was introduced, that is a target of scintillator and LYSO calorimeter positioned transversely to detect positron from RMD emitted at low





Figure 1.6: Schematic view of the MEG detector, [36].

angle. When combined with the final result of MEG, MEG II collaboration obtained the most stringent limit up to date, $BR(\mu^+ \rightarrow e^+\gamma) < 3.1 \times 10^{-13}$ 90% CL, [16]. The MEG II collaboration took data in 2022 and 2023, collecting much more statistics compared to 2021. By 2026, more than twenty-fold increase in statistics is foreseen, with the goal of reaching a $BR(\mu^+ \rightarrow e^+\gamma) \lesssim 6 \times 10^{-14}$ 90% CL.

Figure 1.7: A sketch of the MEG II detector with a simulated $\mu^+ \rightarrow e^+\gamma$ event, [39].

1.2.1.2 $\mu^+ \rightarrow e^+e^-e^+$

The signature of a μ^+ decay at rest consists of two e^+ and one e^- at the same time, with total energy equal to the muon mass and null vector sum of the particle momenta. Since this is a three-body decay, where the momenta of the particles range from a few MeV (depending on the μ^+ speed) to about half the mass of the muon, it is crucial to use a tracker that is both thin and low-mass, with an excellent resolution. It is also important to estimate the probability of having the three particles with momenta above the detection threshold.




Figure 1.8: Picture of the open CDCH equipped with all the wires, [39].

There are two main sources of background. The first one is the radiative muon decay with internal conversion $\mu^+ \rightarrow e^+e^+e^-\nu_e\bar{\nu}_\mu$. This process has a branching ratio of $BR \sim 3.4 \times 10^{-5}$, and it is indistinguishable from the signal when the neutrinos have very low energy.

The second source of background is due to the coincidence of one Michel decay with a e^+e^- pair (1-MD) or two Michel decays with a single e^- (2-MD). In this case, the e^+e^- pair can be produced by Bhabha scattering or photon conversion, while the e^+ can be produced by Compton scattering or mis-reconstructed e^+ and e^+e^- (with the e^- not reconstructed). As a consequence, this source of background depends on the muon rate and can be suppressed with precise vertex reconstruction, timing and track reconstruction. As in the previous channel, the use of a continuous beam is preferred.

1.2.1.2.1 SINDRUM I The current best limit, $BR(\mu^+ \rightarrow e^-e^+e^+) < 1.0 \times 10^{-12}$ at 90% CL, was set by the SINDRUM I experiment at PSI, [40]. A schematic view of the SINDRUM spectrometer is given in Figure 1.9. The spectrometer consisted of a double cone-shaped stopping target in the middle of five concentric multi-wire proportional chambers surrounded by an array of plastic scintillator counters inside a solenoidal magnetic field. Considering a 50 MeV e^-/e^+ , the detector apparatus had a momentum resolution at the level of ~ 1 MeV, a timing resolution ≤ 1 ns and a vertex resolution of ~ 1 cm. The trigger consisted of a charge filter, requiring one negative and two positive particles, within a time window of 7 ns, [7].

1.2.1.2.2 The Mu3e Experiment The goal of the Mu3e experiment is to achieve a single-event-sensitivity of the order of 10^{-16} on the $\mu^+ \rightarrow e^+e^-e^+$ decay, [41] and [42]. A schematic view of Mu3e is shown in Figure 1.10. The same muon beam with MEG II experiment will be used, stopping muons on a thin hollow double-cone Mylar target. The detector will be a 2 m cylinder placed inside a 1.5 T magnetic field and segmented in 5 sections.

The central station will consist of two double layers of pixel detectors and a scintillating fiber tracker. The other four stations will be made of two layers of pixel




Figure 1.9: Schematic view of the SINDRUM experiment, [40].

sensors and a hodoscope of scintillator. Since the Mu3e search relies heavily on accurate track reconstruction, multiple Coulomb scattering is a limiting factor and the technical choices adopted for the detector design have been taken to minimize this effect. The tracker consists of High Voltage Monolithic Active Pixels (HV-MAPS) and the design is such as to exploit the (partial) canceling of the multiple scattering in half of turn. The estimated time and vertex resolutions are $\sigma_t \sim 100$ ps and $\sigma_{xy} \sim 200 \mu\text{m}$ and the momentum resolution will be $100 \div 400$ keV for $10 \div 53$ MeV particles. The experiment is projected in three phases. During the first one, the beam will have an intensity of $\mathcal{O}(10^7) \mu^+/\text{s}$ and there will be only the tracker installed. After that, the beam with an intensity of $\mathcal{O}(10^8) \mu^+/\text{s}$ will be used and the scintillating fibers and two of the additional tracking stations will be added. During the third phase, the beam intensity will increase up to $\mathcal{O}(10^9) \mu^+/\text{s}$ and to reach the single-event-sensitivity of 10^{-16} , two more tracking stations will be added.

1.2.1.3 $\mu^- N \rightarrow e^- N$

In the process of neutrino-less coherent muon to electron conversion, $\mu^- N \rightarrow e^- N$, the energy of the muon is completely transferred to the electron, as the energy transferred to the nuclear recoil is almost negligible. Consequently, the signal manifests itself as a monochromatic electron with 105 MeV energy. A significant advantage of this conversion search, compared to the $\mu \rightarrow e\gamma$ search, is the higher momentum and better separation of the electron signal from the background. Thanks to the one particle final state, this channel is not affected by an accidental background and can benefit from a beam intensity increase. Moreover, $\mu^- N \rightarrow e^- N$ has the best sensitivity to CLFV in a large range of new physics scenarios (Eq. 1.16), as it is sensitive to both the four-fermion interaction and dipole terms.

To cancel out uncertainties due to the overlap of the nucleus and the muon wave functions, the quantity often used to represent the results of such searches is defined




Figure 1.10: Mu3e experimental setup layout (3D view). An example of a $\mu^+ \rightarrow e^+e^-e^+$ decay event is shown, [42].

as:

$$R_{\mu e} = \frac{N(\mu \rightarrow e)}{N(\text{muon capture})} \quad (1.17)$$

The muon capture process occurs extremely fast, on the order of $\mathcal{O}(10^{-13})$ s, and is accompanied by the emission of an X-ray, which can be detected. By measuring the X-rays emitted during muon capture, the denominator of the above ratio can be accurately determined. As no coincidence is required, the experiment's success depends on efficient electron reconstruction.

The primary sources of background from muon decay include:

- Electrons from muon decay-in-orbit (DIO), whose spectrum exhibits a long energy tail and can contaminate the signal region. The probability of contamination within the last MeV is less than 10^{-16} , and it can be controlled with precise momentum resolution;
- High-energy photons generated by radiative captures of pions (RPC) and muons (RMC), which can asymmetrically convert and produce high-energy electrons. RPC can be mitigated using a pulsed beam combined with a delayed time window gate, while RMC can be controlled with precise momentum resolution;
- Cosmic rays, which can either generate electrons or be misidentified as electrons within the signal region. The cosmic ray background can be reduced by employing a veto system around the detector.

In a *hollow*-shaped experimental setup, the background from low-momentum particles can be reduced to manageable levels.

A more detailed discussion of signals and background sources is provided in Section 2.2.

The use of a curved transport solenoid helps to suppress prompt background by a factor of 10^{10} , eliminates neutral particles, and selects particles based on momentum. To prevent spurious pion production, it is essential to maintain the extinction factor - defined as the fraction of protons that strike the production target outside the selected window - below 10^{-10} .

Among all available channels, this one offers the highest sensitivity.

Additionally, this experimental setup can be used to search for the process $\mu^- N \rightarrow e^+ N$, which would also violate lepton number conservation.

1.2.1.3.1 SINDRUM II PSI delivered a 1 MW 590 MeV proton beam that was extracted from the ring cyclotron and directed onto a 40 mm carbon target. The beam line transported secondary particles (π , μ , e) emitted in the backward direction to the SINDRUM II spectrometer (Figure 1.11). The overall structure of the experiment was cylindrical and the target (B), which had a radius of 20 mm, was positioned in the middle of the detector.

Two drift chambers (F and G) were adopted to measure the helical trajectories, with the ionization electrons drifting radially towards the amplification regions. The tracker used CO_2 -isobutane (70%/30%) as a drift gas while the second one He-isobutane (85%/15%). Two 3 mm thick plastic scintillator hodoscopes (D) and a 3 cm thick plexiglass Cherenkov hodoscope (E) provided trigger and time information. The device included two end-cap hodoscopes at opposite ends of the tracking zone to aid in triggering and resolving ambiguities during event reconstruction.

The number of muons stopped was monitored observing the characteristic muonic gold X-rays passing through the superconducting coil of the spectrometer. A Ge(Li) detector was used for this purpose. SINDRUM II set the limit on the muon conversion on Au target (currently the best limit) at 7×10^{-13} , [2].




Figure 1.11: Pictorial view of the SINDRUM II experiment, [2].

1.2.1.3.2 Mu2e Mu2e will make use of an 8 GeV, 25 kW pulsed proton beam, with 100 ns wide bunches spaced by $1.7 \mu\text{s}$. Figure 2.4 shows a schematic view of the experimental setup and of the three sections of the experiment: Production Solenoid, Transport Solenoid and Detector Solenoid. The S shape lowers the background due to neutral particles and selects the muon beam sign. The straw tube tracker and the crystal electromagnetic calorimeter are located downstream of the Al target. Both these detectors adopted a *hollow-cylinder* geometry. The expectation is to reach the level of single event sensitivity of $SES = 3 \times 10^{-17}$ in three years of data-taking [37]. Mu2e will be described extensively in the Chapter 2.

The Mu2e Collaboration is also performing preliminary studies for the upgraded detector, Mu2e II, [43]. The proton beam intensity will be improved by the PIP-II upgrade that will increase the rate of stopped muons on target from $10^{10} \mu^-/\text{s}$ to $10^{11} \mu^-/\text{s}$. New detector technologies are under study for the upgraded Mu2e II.

1.2.1.3.3 DeeMe A simpler setup will be adopted by the Direct emission of electron from Muon to electron conversion (DeeMe) experiment at J-PARC to search the muon-to-electron conversion, [44]. A pictorial view of DeeMe experiment is shown in Figure 1.12. The experiment will use the Muon Science Establishment’s new beam-line (H-line), which will provide 3 GeV protons (a pair of bunches separated by 600 ns at 25 Hz) that will hit on the production target. Some muons ($\mathcal{O}(10^{10})\mu/\text{s}$) will be stopped in the target itself, allowing for a search for $\mu \rightarrow e$ with just one target. The signal will be detected with multi-wire proportional chambers and a spectrometer. A dipole will remove low momentum background particles from the transport system. The experiment aims to achieve a single-event sensitivity of 10^{-13} using a graphite target, and later of $10^{-14} \div 10^{-15}$ (depending on running duration) using a silicon carbide (SiC) target with a higher capture rate.




Figure 1.12: Schematic view and timing structure of the DeeMe experiment, [44].

1.2.1.3.4 COMET The COherent Muon-to-Electron Transition (COMET) experiment is under construction at the Japanese Proton Accelerator Research Center (J-PARC), [45]. This experiment shares several characteristics with Mu2e. COMET will use a 8 GeV, 56 kW pulsed proton beam with $1.17 \mu\text{s}$ spaced bunches. The two primary differences between the COMET and Mu2e experiments can be observed in the schematic representations provided in Figure 1.13. First, COMET employs

a C-shaped transport solenoid, in contrast to the S-shaped one used in Mu2e. The presence of a C-shaped transport solenoid will allow a tighter muon momentum selection, with a reduced beam intensity ($\sim 30\%$ less).

Secondly, COMET introduces an additional curved solenoid after the stopping target. This solenoid serves to remove electrons, including those up to and beyond the Michel peak, before they reach the tracker.

COMET will have two stages: Phase-I and Phase-II.

Phase-I The first phase involves understanding the experimental techniques and studying the backgrounds, making an intermediate measurement of the $\mu - e$ conversion at a 90% CL sensitivity on $R_{\mu e} \sim 7 \times 10^{-15}$.

The proton power will be restricted to 3.2 kW and a single 90° bend will be utilized. The major challenge is the short distance between all the elements. To track electrons, a cylindrical drift chamber will be used. To trigger and time the tracker, scintillating hodoscopes will surround the tracker.

Phase-II A straw tube tracker and a LYSO crystal electromagnetic calorimeter will be adopted to deal with the increased particle rate. The entire magnetic system will be expanded.




Figure 1.13: Conceptual view of the staging approach in COMET experiment, [45].

1.2.2 τ Channels

The τ lepton is a very promising source of CLFV decays, [7]. The large τ mass ($m_\tau \sim 1.777$ GeV) allows to study multiple CLFV channels, with respect to muon decays, [5]. Some of these processes are $\tau \rightarrow l\gamma$ (radiative decay), $\tau \rightarrow 3l$ (three-body decay) and $\tau \rightarrow lh$, where h is a light hadron and l is an electron or muon.

The τ sector of a CLFV search takes advantage from a higher predicted branching ratio, compared to muons, according to $(\frac{m_\tau}{m_\mu})^\alpha$, where α depends on the model.

However, the search is affected by the short τ lifetime ($\tau_\tau \sim 2.9 \times 10^{-13}$ s), that makes not possible to produce τ beams.

To constrain the kinematics of decay, large detectors with good particle identification, tracking, calorimetry, and hermeticity are required. Because τ CLFV searches are performed with beams and detectors used for a broader physics programme, the increased sensitivity derived by its greater mass is partially decreased by the number of τ s that can be detected.

In these experiments, a pair of $\tau^+\tau^-$ is produced by the decay of $\Upsilon(4s)$ resonance at $\sqrt{s} = 10.58$ GeV. The branching ratio of this process is 90% of the one of $b\bar{b}$. One τ -lepton decays via SM process (tag side), while the signal side is determined based on the specific topology of each channel. At e^+e^- and pp colliders, the τ s are not created at rest. Because of the boost, the decay products could have energy of several GeV, posing an experimental challenge of providing wide-range calibrations for detectors (from a few hundred MeV to several GeV), [7].

Table 1.2 lists the current best limits on the τ CLFV searches and Figure 1.14 displays the results from the BaBar, [23], Belle, [46] and LHCb, [47]. In a short




Figure 1.14: Tau lepton-flavour-violating branching ratio 90% CL upper limits summary plot, [48]. The green and blue cones are respectively the Belle and Babar's best upper limits. The red dots are Belle II's expected upper limits for the 50 ab^{-1} run.

time, the Belle 2 experiment at Super KEKB will try to increase its limits below 5×10^{-9} or 10^{-9} for the radiative and three body decays respectively, with an integrated luminosity of 50 ab^{-1} . The main background, in radiative decay channel only, arises from $e^+e^- \rightarrow \mu^+\mu^-\gamma$ or $e^+e^- \rightarrow \tau^+\tau^-\gamma$ where one of the τ s decays via $\tau \rightarrow \nu\bar{\nu}$. To reduce the background, it is possible to use the τ polarization or to collect large samples of τ leptons at a lower center-of-mass energy, where initial state radiation is negligible in the signal region. This might be the case for a τ -charm factory operating slightly above the τ production threshold, [49].

Chapter 2

The Mu2e Experiment

The Mu2e experiment aims to investigate the phenomenon of Charged Lepton Flavour Violating (CLFV) neutrino-less conversion, where a negative muon transitions into an electron within the field of an aluminium nucleus. The experiment will use as variable for the results the ratio between the conversion and the nuclear muon capture rates:

$$R_{\mu e} = \frac{\mu^- + N(Z, A) \rightarrow e^- + N(Z, A)}{\mu^- + N(Z, A) \rightarrow \nu_\mu + N(Z - 1, A)} \quad (2.1)$$

The goal is to improve the current limit, set by the SINDRUM-II experiment [2], by four orders of magnitude and reach a SES (single-event-sensitivity) of 2.3×10^{-16} on the conversion rate, a 90% CL of 6.2×10^{-16} and a 5σ discovery reach at 1.2×10^{-15} during Mu2e Run I. Mu2e is presently undergoing commissioning and integration at the Fermilab Muon Campus, with contributions from an international collaboration. Data taking is planned to begin in 2026. This Chapter provides an overview of the Mu2e experimental techniques and infrastructure. Fundamental bibliography for this chapter can be found in [1], [3], [4], [50], [37].

2.1 Experiment concept

The Mu2e concept has been proposed back in 1989 by V. Lobashev and R. Dzhilkibaev [51], and it relies on the following principles. A beam of negative muons, μ^- , is generated by directing a proton beam at a production target, yielding negative pions along with other mesons and hadrons. Pions, with a decay rate of more than 99.9% due to helicity suppression, undergo $\pi^- \rightarrow \mu^- \bar{\nu}_\mu$ decay in flight.

Low-momentum secondary muons are trapped by a Stopping Target (ST) to form muonic atoms. Within approximately 10^{-13} s, a muonic atom transitions to the 1s state [52]. Given the brief cascading time compared to the mean muon lifetime in a muonic atom, instances of muon decay before reaching the ground state are negligible. The cascades emit photons, aiding in estimating the number of muons captured on the ST. Stopped muon lifetime depends on the ST material. In the Mu2e experiment, an Aluminum ST is adopted resulting in a lifetime of 864 ns.

Muons interacting with a nucleus mostly undergo two processes: nuclear capture $\mu^- N \rightarrow \nu_\mu N'^*$, where N'^* represents an excited Magnesium nucleus, and muon decay-in-orbit (DIO), that is the three-body decay with neutrinos $\mu^- N \rightarrow e^- N \nu_\mu \bar{\nu}_e$. As shown in Figure 2.18a, the ratio between these processes varies with the ST material. For ^{27}Al , approximately 60.9% of muonic atoms undergo muon nuclear capture. The remaining 39.1% undergo muon DIO. The experiment aims to identify a third process, neutrinoless muon-to-electron conversion $\mu^- N \rightarrow e^- N$.

The Mu2e experiment can also detect a lepton-number violating process $\mu^- N \rightarrow e^+ N'$. This process violates both the charged lepton flavour and the total lepton number conservation ($|\Delta L| = 2$).

2.2 Signal and backgrounds

2.2.1 Conversion Electron signal

The specific goal of Mu2e search is the neutrinoless coherent $\mu^- \rightarrow e^-$ conversion in the field of an aluminum nucleus: the electron recoils off the entire nucleus and follows the two-body decay kinematics [1]. The mass of the nucleus is large compared to the mass of the electron, hence the recoil term is minimal.

The outgoing nucleus remains in the ground state and as a result the signature of the process is a monochromatic conversion electron (CE) with the energy slightly lower than the muon rest mass:

$$E_{CE} = m_\mu - E_{recoil}(A) - E_{bind}(Z) \quad (2.2)$$

where m_μ is the muon mass, $E_{recoil} \simeq \frac{m_\mu^2}{2m_N}$ is the recoil energy of the target nucleus, with m_N the nucleus mass, and $E_{bind} \simeq \frac{Z^2 \alpha^2 m_\mu}{2}$ is the binding energy of the $1s$ state of the muonic atom [37]. For the Mu2e ST material, ^{27}Al , $E_{CE} = 104.97$ MeV [53].

2.2.2 Backgrounds

Five main background sources to the $\mu^- \rightarrow e^-$ conversion search are as follows [1]:

1. electrons or muons coming from cosmic rays;
2. muon Decay-in-Orbit (DIO) and Radiative Muon Capture (RMC);
3. antiproton annihilation in the beamline and in the ST;
4. radiative pion capture (RPC), decays in flight (DIF).

The detailed discussion of these background sources will be provided below, along with the corresponding strategies to reduce them.

2.2.2.1 Cosmic rays

Cosmic ray particles, predominantly muons, also contribute to the experiment background. Several processes that may occur are listed here:

1. muon decays occurring within or near the detectors;
2. muons interacting with nuclei within detectors or surrounding materials;
3. muons scattering within the detectors and being erroneously identified as electrons.

Detailed simulations estimated that in the Mu2e experiment, conversion-like signals generated by cosmic rays occur approximately once per day [54].

Mitigation involves making use of a combination of active veto and shielding. The Mu2e Cosmic Ray Veto (CRV) system is described in Section 2.7.3.

2.2.2.2 Intrinsic backgrounds

The intrinsic backgrounds originate from two physics processes: muon decay-in-orbit (DIO) and radiative muon capture (RMC). In this context, *intrinsic* denotes that this source of background is generated by the same muons used to perform the conversion signal search. It consequently scales with the stopped muon flux and the number of protons on target.

2.2.2.2.1 Muon Decay-in-Orbit Electrons energy coming from the three body decay of a free muon, $\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$, is described by the Michel spectrum. The differential decay rate can be computed [55]:

$$\frac{d\Gamma_{\text{free}}}{dx} = \frac{G_F^2 m_\mu^5}{192\pi^3} x^2 (6 - 4x + \frac{\alpha}{x} f(x)) \quad (2.3)$$

where $x = \frac{2E_e}{m_\mu}$, $0 \leq x \leq 1$, G_F is the Fermi constant, α is the fine-structure constant, m_μ is the muon mass, E_e is the electron energy and $f(x)$ represents a complicated radiative correction term, described in [56].

However, the presence of an Al nucleus allows the electron to interact with it, exchanging momentum and resulting in a maximum possible energy, called *endpoint energy*, being the same as the CE energy. As shown in Figure 2.1a, the spectrum (with the presence of the Al nucleus) peaks at 52.8 MeV, which is the half of the muon rest energy or the energy of the searched CE. The difference between the two energy spectra, one neglecting the nuclear recoil and the other one in the presence of the Al nucleus, is illustrated in Figure 2.2. Figure 2.1b shows a more detailed zoom on the high-energy range of the spectrum. When the electron energy $E_e \geq 85$ MeV, the dominant term at the leading order scales with $(E_{\mu e} - E_e)^5$ [53], resulting in a low rate within the energy range very close to the endpoint. However, due to the energy loss of particles interacting with ST or detector materials and radiative

corrections, the CEs tend to be reconstructed with a left-skewed energy distribution, with a full width at half maximum (FWHM) approximately on the order of 1 MeV [57], as shown in Figure 2.3.

A particle tracking detector with high momentum resolution also helps to reduce background, which will be explored in the next sections.

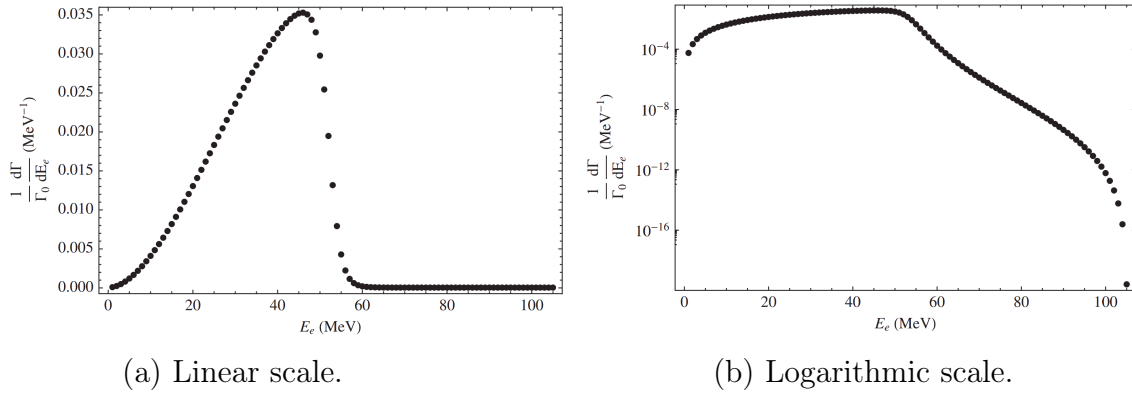


Figure 2.1: Decay-in-orbit electron spectrum on Al [53].




Figure 2.2: The electron energy spectrum near to the endpoint. The black line represents the DIO spectrum w/o the nuclear recoil, while the red dashed line takes into consideration the recoil of the Al nucleus [53]. The green arrow shows the spectrum end at 104.973 MeV.

2.2.2.2 Radiative Muon Capture The Radiative Muon Capture (RMC) is a muon capture with emission of an energetic photon.

In the process $\mu^- N \rightarrow \gamma \nu_\mu N'^*$, the photon can either be real or virtual. The photon, interacting with matter or undergoing pair production, can produce electrons with energies close to E_{CE} , introducing background to the experiment. The emitted photon's energy follows a spectrum, with its maximum energy, denoted as the kinematic limit k_{max} , determined by the equation [1]:

$$k_{max} = m_\mu c^2 - |E_b| - E_{rec} - \Delta M, \quad (2.4)$$

where E_b represents the muon binding energy on the initial nucleus, E_{rec} the recoil energy of the daughter nucleus and ΔM is the rest energy difference between the final and initial nuclei. This formula neglects higher order nuclear effects. RMC can be effectively mitigated in the Mu2e experiment by selecting Al as the ST material. The ST is selected so that the daughter nuclide of a muon-capture process of any kind is heavier than the original nuclide. For aluminum, the RMC photon endpoint energy is 101.9 MeV, approximately 3.1 MeV below the CE energy [1]. The FWHM of the conversion peak is around 1 MeV, therefore the RMC background will be outside the signal region. However, the RMC background might distort the DIO spectra around 100 MeV range, making it difficult to extrapolate the high momentum end of the DIO spectrum. The RMC background will be determined from the positron data.

2.2.2.3 Prompt processes

This type of background sources can generate electrons at roughly the same time as the entering beam particles. There are four primary sources: radiative pion capture (RPC), pion decay-in-flight (π -DIF), muon decay-in-flight (μ -DIF) and beam electrons.

2.2.2.3.1 Radiative Pion Capture Pions, reaching the ST and captured in Al can produce high-energy photons, i.e. $\pi^- N(A, Z) \rightarrow \gamma N^*(A, Z - 1)$. This process, called radiative pion capture, has a branching ratio of $\sim 2\%$ [58]. Similar to RMC, the photon can internally convert into an electron-positron pair or be radiated as an on-shell photon and convert in the detector material producing an e^+e^- pair. The resulting electrons contribute to the experiment background. Despite its similarity to RMC, RPC is a more challenging background to suppress due to the fact that the endpoint of the energy spectrum of photons, and consequently the resulting electrons, is not constrained by the rest energy of the muon. The pion mass is 139.6 MeV, above the CE energy, and the spectrum of the electrons resulting from RPC overlaps with the search region. The SINDRUM II results were limited by the pion-induced background and also by the low intensity of its muon beam [2]. SINDRUM II used a primary proton beam with a frequency of one pulse every 19.75 ns, lasting approximately 0.3 ns. This interval between pulses was shorter than the 26 ns lifetime of pions [15], ensuring a consistent pion flux. To mitigate RPC, SINDRUM II made use of a degrader to suppress pions and a veto counter in the beam, resulting in less than one out of 10^9 pions reaching their target. However, given the more intense beam, the Mu2e experiment has to change the approach. Mu2e will suppress the RPC background by delaying the search window in time with respect to the incoming beam by several hundred nanoseconds [37]. One important point to note is that the out-of-time protons could produce *late* pions not rejected by the delayed timing window. Consequently, it is important for the pulsed beam to achieve a high extinction level. Further elaboration on the pulsed beam used in Mu2e will be provided in Section 2.4.

2.2.2.3.2 π -DIF and μ -DIF The pion decay-in-flight (π -DIF) and the muon decay-in-flight (μ -DIF) exhibit quite similar characteristics. Free pions and muons can undergo electron decay while transitioning from the production target to the ST, through the processes $\mu^- \rightarrow e^- \nu_\mu \bar{\nu}_e$ and $\pi^- \rightarrow e^- \bar{\nu}_e$. In the center-of-mass frame of the initial particles, electrons originating from the first process exhibit an energy spectrum that reaches an endpoint of 52.8 MeV, while those from the second process have a consistent energy of approximately 70 MeV. Since the pions and muons move at relativistic velocities, the energies (and momenta) of the resultant electrons are boosted. For instance, a muon with momentum of around 79 MeV/c or a pion with momentum close to 70 MeV/c can generate an electron with an energy of 105 MeV [1]. Implementing a pulsed proton beam and using a delayed live window can help suppress background from π -DIF and μ -DIF events. Particles with sufficient momentum to boost the daughter electrons to the concerning energies move quickly along the muon beamline and are gone by the time the live search window begins [3].

2.2.2.3.3 Beam electrons Other mechanisms generate electrons, both at the production target and along the muon beamline. For instance, neutral pions formed at the production target can decay in two photons, after which the photons can either create electron-pairs or interact with nearby materials to generate electrons. The beam electron background is suppressed by the delayed search timing window.

2.2.2.4 Delayed processes from antiprotons

Protons, at a given threshold, can generate antiprotons within the production target. This occurs through the process of antiproton production: $pp \rightarrow ppp\bar{p}$. If we consider that all four particles in the final state are at rest in the center of mass frame, the minimum kinetic energy needed for $p\bar{p}$ production can be found, which is approximately $6m_p c^2 \sim 5.6$ GeV, where m_p is the mass of the proton. In an ideal scenario, maintaining the beam energy below this threshold would eliminate this background source. However, the Mu2e proton beam energy is the threshold for antiproton production. Antiprotons are long-lived and massive. Antiprotons with momenta below 100 MeV/c travel at speeds less than $0.1c$, requiring several μs to spiral from the production target to the ST [1]. They have the correct charge and momentum to pass through the collimators placed between the production target and the ST. Annihilating or undergoing interactions with other materials, they have the capability to release a substantial amount of energy, generating electrons.

The delayed timing of such interactions is much longer than the muon lifetime, leading to a continuous flow of antiprotons reaching the ST. The pulsed beam and the delayed live window fail to suppress the antiproton background. The best approach is to prevent the antiprotons from reaching the region where the ST is located. Two thin absorbers are positioned in the muon beamline to capture the antiprotons. Their design was developed to find a compromise between increasing the antiproton absorption and decreasing the muon beam loss.

2.2.3 Background estimates and signal sensitivity

Mu2e run plan is divided in two different phases. Run I will take place in 2027, before a 2 years shutdown due to the planned accelerator upgrade for the long baseline neutrino program. In Run I phase one, a low intensity proton beam, 1.6×10^7 protons/pulse, will be used. In Run I phase two, the mean intensity will be increased up to 3.9×10^7 protons/pulse. As discussed in [37], during Run I, the 5σ discovery sensitivity is $R_{\mu e}^{5\sigma} = 1.2 \times 10^{-15}$, with a total expected background of 0.11 ± 0.03 events. Reaching the 5σ significance level requires observing 5 $\mu \rightarrow e$ events in the two-dimensional search window $103.60 < p < 104.90$ MeV/c, $640 < T_0 < 1650$ ns. In the absence of a signal, the expected upper limit is $R_{\mu e} < 6.2 \times 10^{-16}$ at 90% CL. The single event sensitivity (SES) is defined as

$$SES \equiv \frac{1}{N_{POT} \cdot P_{\mu \text{ stop}} \cdot \epsilon_{CE} \cdot BR_{capture}} \quad (2.5)$$

where N_{POT} is the number of protons on target in the experiment, $P_{\mu \text{ stop}}$ is the number of muons stopped on target per proton, ϵ_{CE} is the CE acceptance, which is a product of detector efficiency (dependent on the momentum signal region) and fraction of muons interacting in the live time window and $BR_{capture}$ is the branching ratio of muon captures, which is 60.9%. The optimized Mu2e signal window corresponds to a SES of 2.3×10^{-16} . The background estimates after the sensitivity optimization are summarized in Table 2.1.

Figure 2.3 shows the momentum and time distributions for CE signal and background processes corresponding to the optimized signal window for Run I. A detailed analysis and an estimate of the Mu2e expected backgrounds for Run I can be found in [37].

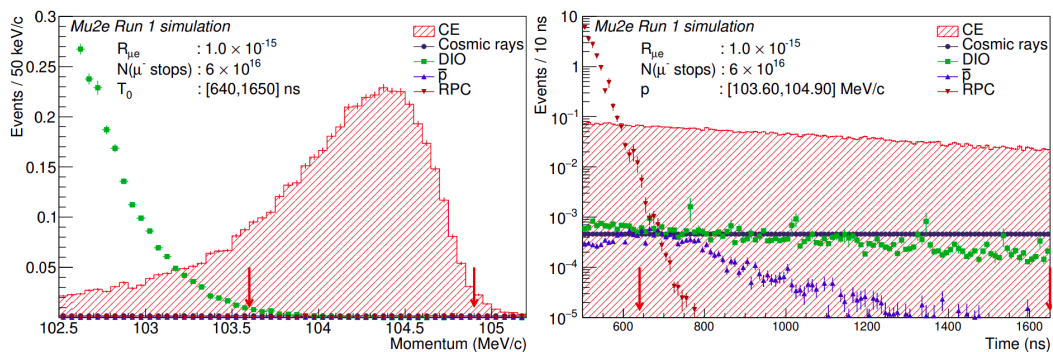


Figure 2.3: Left: Momentum distribution of the CE signal and expected backgrounds. Right: Time distribution of the CE signal and expected backgrounds. The arrows show the signal region selected for the analysis, $103.60 < p < 104.90$ MeV/c and $640 < T_0 < 1650$ ns, Section 2.4.1. The CE signal distributions correspond to $R_{\mu e} = 1 \times 10^{-15}$ [37].

Channel	Mu2e Run I
SES	2.4×10^{-16}
Cosmic rays	0.046 ± 0.010 (stat.) ± 0.009 (syst.)
DIO	0.038 ± 0.002 (stat.) $\pm_{-0.015}^{+0.025}$ (syst.)
Antiprotons	0.010 ± 0.003 (stat.) ± 0.010 (syst.)
RPC in-time	0.010 ± 0.002 (stat.) $\pm_{-0.003}^{+0.001}$ (syst.)
RPC out-of-time	$(1.2 \pm 0.1$ (stat.) $\pm_{-0.3}^{+0.1}$ (syst.)) $\times 10^{-3}$
RMC	$< 2.4 \times 10^{-3}$
Decays in flight	$< 2 \times 10^{-3}$
Beam electrons	$< 1 \times 10^{-3}$
Total	0.105 ± 0.032

Table 2.1: Summary of the several background sources to the CE search as expected in Mu2e Run I [37]. The Table also shows the corresponding SES.

2.3 Experimental setup

Figure 2.4 shows a schematic overview of the Mu2e experiment. The experiment uses a solenoid system to generate magnetic fields essential for its operations. The Production Solenoid (PS) surrounds the production target, while further downstream, the Transport Solenoid (TS) provides the magnetic field for the muon beamline. The TS, configured in an S-shape, incorporates collimators and proton absorbers. The ST is located at the beginning of the Detector Solenoid (DS). Proton absorbers surrounding the ST are not shown in the picture. The tracker and the calorimeter are housed in the DS, enabling momentum measurement and particle identification. Additionally, a Stopping Target Monitor (STM) positioned downstream at the DS's end, not shown, measures the rate of muons stopped in the ST. Also not shown in the picture is the Mu2e CRV system which surrounds the DS and half of the TS. In the next sections, a more detailed description of systems of the Mu2e experiment will be given.

2.4 Accelerator system and the proton beam

2.4.1 Pulsed proton beam

As previously discussed in Section 2.2.2, the Mu2e experiment uses a pulsed proton beam.

The 8 GeV, 8 kW beam originates from the Fermilab Booster [59]. The proton pulses are separated by 1695 ns. Figure 2.5a illustrates the Fermilab accelerator facilities involved in generating and delivering the pulsed proton beam. The Fermilab Booster delivers 8 GeV protons in 20 batches throughout a 1.4 s Main Injector cycle


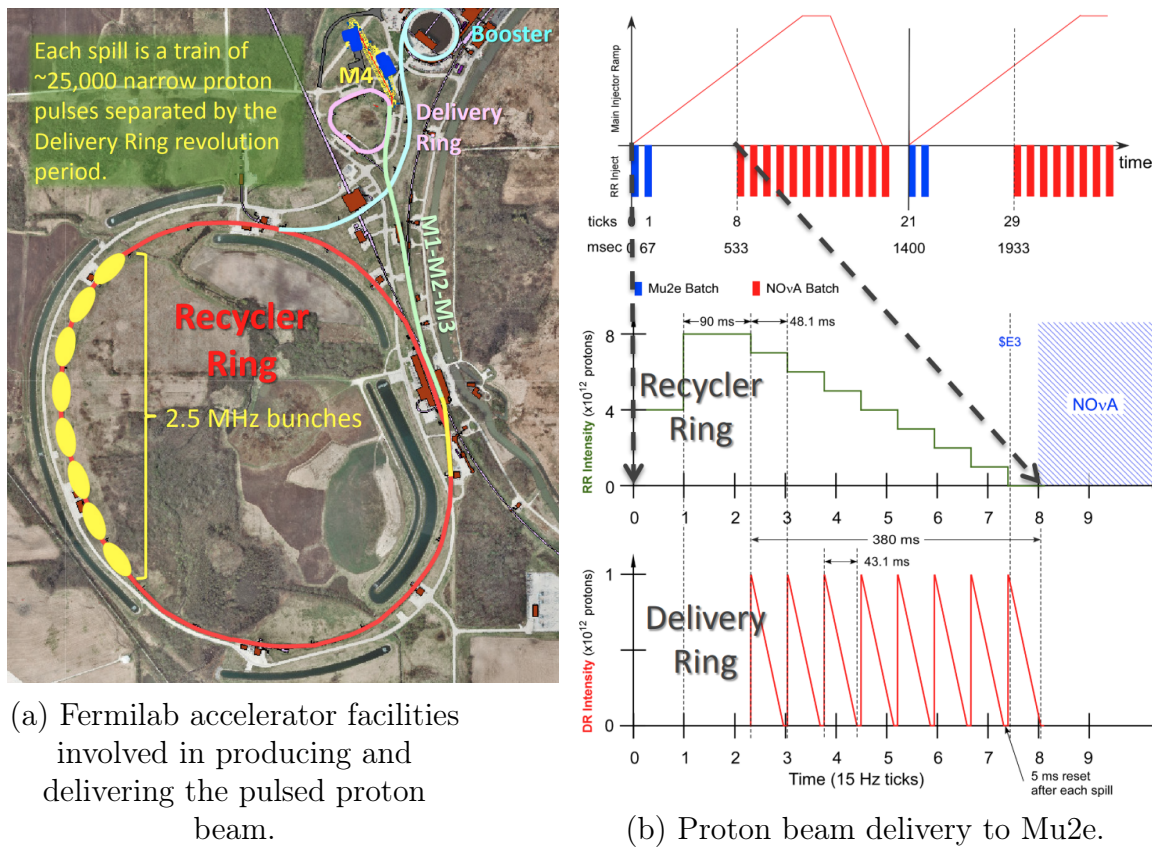


Figure 2.4: Schematic view of the Mu2e apparatus. The center of the Mu2e reference frame is located at the COL3 collimator center, its y -axis points upwards, the z -axis is parallel to the DS axis and points downstream, and the x -axis completes the right-handed reference frame [37].

at 15 Hz, as shown in Figure 2.5b. Thus, the accelerator timeline is described using a fundamental time unit of 1 tick that corresponds to 66.7 ms.



(a) Fermilab accelerator facilities involved in producing and delivering the pulsed proton beam.

(b) Proton beam delivery to Mu2e.

Figure 2.5: Pulsed proton beam delivery [60].

The two Mu2e batches, represented by the two blue bars at ticks 1 (1BB) and 2 (2BB), are injected into the Recycler Ring, each containing 4×10^{12} protons. The protons from these batches are reorganized within the Recycler using a 2.5 MHz radio frequency (RF) system into 8 bunches. These bunches are then extracted

individually from the Recycler and transported to the Delivery Ring every 48.1 ms, as shown in the middle part of Figure 2.5b. Once inside the Delivery Ring, a single bunch of 1×10^{12} protons undergoes gradual extraction. This process results in the extraction of a small fraction of the bunch per revolution, delivered to the Mu2e experiment. The complete bunch is extracted over a span of 43.1 ms, across ~ 25000 turns around the Delivery Ring, as shown in the bottom part of Figure 2.5b. Figure 2.6 illustrates the temporal profile of the beam at the PT. Consecutive

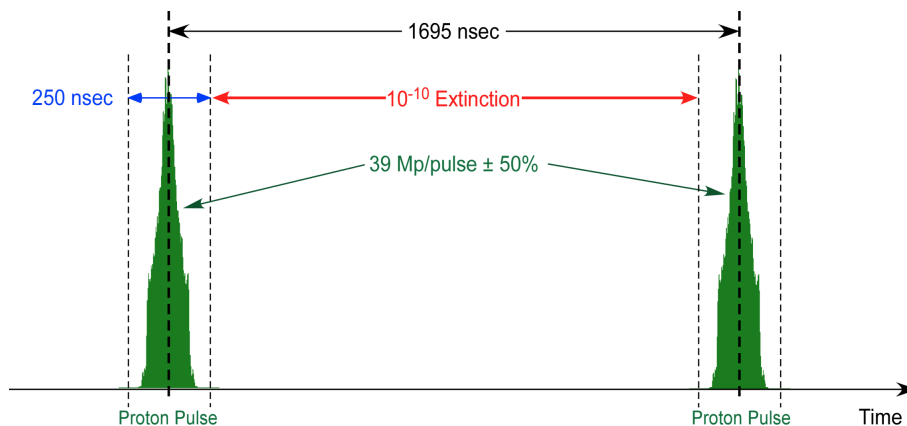


Figure 2.6: Proton beam profile at the Mu2e Proton Target [60].

proton pulses are spaced by 1695 ns. Each pulse lasts for 250 ns and contains $(3.9 \pm 2.0) \times 10^7$ protons. The 1695 ns pulse separation is highly advantageous for the Mu2e experiment. Figure 2.7 shows the beam pulse, the simulated pion flux, the muon capture rate on the ST and the muon decay rate. The active window for detecting CEs begins at ~ 640 ns and extends for more or less $1 \mu\text{s}$.




Figure 2.7: The Mu2e beam timing: the proton pulses arrive at the production solenoid every 1695 ns. A delayed live-time window can suppress the beam-related background [37].

2.4.2 Proton beam extinction and Extinction Monitor

As mentioned in the previous section, the Mu2e experiment needs an extinction level of the incoming proton beam, to reduce backgrounds caused by out-of-time protons. The extinction rate, defined as the ratio between the number of out-of-time protons and the number of the in-time protons, should be lower than 10^{-10} [1]. The structure of the beam leads to an extinction level 2.1×10^{-5} . To take into account the fact that some beam will leak out of two consecutive proton pulses, an Extinction Insert is deployed in the M4 beamline between the Delivery Ring and the Mu2e experiment. The out-of-time beam particles are swept into a collimator system by an oscillating dipole, called AC dipole. The AC dipole provides an additional extinction factor of 5×10^{-8} , improving the overall extinction to 1.1×10^{-12} , leaving a margin of 10^2 [60]. An Extinction Monitor is positioned downstream of the production target along the proton beamline, as shown in Figure 2.8. It monitors the extinction level of the proton beam and delivers a measurement with an accuracy of 10%.

The Extinction Monitor consists of a collimator and magnetic filter system, a pixel telescope, a system of trigger scintillators and a range-stack. The collimator and magnetic filter system transport a small number of particles generated at the production target to the Extinction Monitor. The pixel telescope tracks the trajectory of charged particles coming from the collimator. The pixel telescope consists of a permanent magnet and 8 scintillators, as shown in Figure 2.9. The system uses a permanent magnet to separate two sets of four scintillator planes, allowing for momentum measurements of entering particles. The range stack is located further downstream from the pixel telescope. Steel absorber plates separate scintillators, distinguishing between hadrons and muons based on their penetrating capacity.




Figure 2.8: The Extinction Monitor is located downstream of the production target [61].




Figure 2.9: The tracking spectrometer of the Mu2e experiment, consisting of eight planes of pixel detectors and a permanent magnet spectrometer [61].

2.4.3 Production Target

The Mu2e production target (PT) is made of tungsten. The material was chosen because of its high melting point (3422 °C) and excellent resistance to deformation and corrosion.

It is positioned in the middle of the PS bore [1]. The tungsten core is embedded in a *bicycle wheel* frame, suspended by the spokes. The current design of the PT, shown in Figure 2.10, has circular rings at the end and its segmented core provides a better temperature dissipation. This design is expected to last for more than one year. The PT provides a muon yield of 0.0016 μ /proton [62].




Figure 2.10: The design of the PT. Left: bicycle wheel structure with the PT at the center. Right: zoom on the tungsten target.

2.5 Solenoids

Mu2e uses a sophisticated magnetic system to form the high-intensity muon beam by collecting and filtering the particles emerging from the PT. It is composed of three parts: the **Production Solenoid** (PS), the **Transport Solenoid** (TS) and

the **Detector Solenoid (DS)**, shown in Figure 2.4. Each one is made of superconducting coils wound with aluminum stabilized Nb-Ti Rutherford cables.

The resulting magnetic field varies from 4.6 T at the upstream end of the PS to 1 T at the downstream end of the DS. The muons are guided towards the ST and the DS by lowering the magnetic field. Local magnetic field minima are avoided to avoid trapping particles in these areas. In the PS, the magnetic field decreases from 4.6 T to 2.5 T at the entrance of the TS. The large gradient helps to collect the secondary pions and muons and to direct them towards the DS. The magnetic field across the S-shaped TS changes its value only by a factor of 0.5 T. The shape of the TS allows to select charged particles and its dimension was set to avoid transmitting particles with large momentum. As particle drifts in the solenoid field is dependent on the particle charge, positively and negatively charged muons drift in opposite vertical directions and are separated, as explained in Appendix A. In the upstream part of the TS, Figure 2.13, the positive (blue) and negative (red) muons are deflected downwards and upwards respectively. Positive muons are stopped in collimators COL3u and COL3d. The TS is long enough for pion decay, suppressing the RPC background. The magnetic field in the first half of the DS is reduced from 2 T to 1 T. In a smoothly gradient field, the adiabatic invariance of the magnetic flux can be used. Assuming a constant p_{\perp}^2/B , there is:

$$v_{\parallel}^2 = v_0^2 - v_{\perp 0}^2 \frac{B(z)}{B_0} \quad (2.6)$$

Here, \perp is referred with respect to the magnetic field B and \parallel is referred to the z direction. The subscript 0's indicates the initial state. The gradient pitches electrons forward into the tracker's acceptance.

The second half of the DS containing the tracker and the calorimeter has an approximately uniform magnetic field of 1 T.




Figure 2.11: The muon beamline, composed of the Production Solenoid (PS), the Transport Solenoid (TS) and the Dector Solenoids (DS) [63].

2.5.1 The Production Solenoid

The Production Solenoid (PS) (4 m long), shown in Figure 2.12, collects pions, kaons generated by the interactions between the 8 GeV proton beam and the production target. Its coils are wound with Nb-Ti superconducting cables stabilised with Al, the magnet is cooled by liquid helium. To avoid damages to the superconducting cables and to the cooling system a bronze heat and radiation shield is built inside the solenoid. The position and shape of both the target and the magnetic field have been optimized to maximize the production and collection of the desired particles: backwards pions and muons. This choice was made to avoid the overwhelming flux of particles (neutrons, photons, electrons and positrons from photon conversions) produced in the forward direction and the leftover incoming protons. The magnetic field is 4.6 T at the end of the PS and decreases almost linearly to 2.5 T at the junction with the TS. The magnetic field gradient allows to collect also part of the particles emitted along the proton beam direction, improving the stopped muon and pion yield.




Figure 2.12: Cross-section of the production solenoid [64]. The production target is placed approximately at the center of the superconducting coils.

2.5.2 The Transport Solenoid

The S-shaped TS is composed of a series of wide aperture superconducting solenoid rings. It also contains a set of collimators and absorbers to provide charge and momentum selection and reduce the flux of antiprotons. A thin window assembly is installed at the beginning of the TS and also between the rotatable collimators to absorb antiprotons in the beam.

The TS is divided in five regions.

The first one is called TS1, and it contains the first collimator (COL1), Figure 2.14, made of copper wedges. It filters particles from their momentum and reduces the radiation damage for coils of the upstream part of the TS (TSu).

A toroidal B-field of the TS2 leads to a vertical displacement of the charged particles, according to their charge and momentum. Neutral particles, that are not sensitive to the magnetic field are not transported to the next section.

In TS3, negative muons are selected with two rotating collimators (COL3u and COL3d, shown in Figure 2.14). The rotating collimators allows selection of μ^- 's instead of μ^+ 's, which can be used for the detector calibration.

The TS4 is the second toroid section of the TS and its role is to collimate the beams on the TS axis.

TS5 connects TS with DS and it contains a collimator (COL5), Figure 2.14, made of polyethylene, which will serve as a shield from neutrons. It is linked to the DS and matches its field to provide the best beam transmission.




Figure 2.13: The schematic representation of the TS and the collimators COL3u and COL3d showing the offset apertures in those collimators. The upper spiraling negative muons (red) pass through the aperture while the positive muons (blue) are stopped by these collimators [65].




Figure 2.14: Design of the collimators in the TS. From left to right: COL1, COL3u and COL3d and COL5.

2.5.3 The Detector Solenoid

The DS, shown in Figure 2.15, is approximately 11 m in length and 2 m in radius. It houses the ST, the proton absorber, the tracker and the calorimeter. The system is divided into two sections: a 4 m gradient section following the TS, where the magnetic field decreases linearly from 2 T to 1 T, and a 6 m spectrometer section.

The gradient region drives the CE inside the tracker acceptance, while the low momentum particles are driven to the low radius region that is not covered by




Figure 2.15: Overall structure of the DS coils and cryostat [3].

detectors. This allows to detect both e^+ and e^- , with the disadvantage that the detectors are exposed to the beam flash and to protons and neutrons produced by the muon nuclear capture in the ST.

A series of materials placed inside the DS, shown in Figure 2.16, is used to suppress the rate of protons and neutrons that can generate spare hits in the tracker:

- **Inner Proton Absorber (IPA):** this is a thin (0.5 mm) conical frustum made of low Z -material. It aims to reduce the rate of protons produced in muon nuclear capture (around 0.03 per captured muon) while minimizing the energy loss of CEs;
- **Outer Proton Absorber (OPA):** another conical frustum with a thickness of 20 mm, serving the same purpose as the IPA.




Figure 2.16: The structure of the inner and outer proton absorber.

2.6 Stopping Target

The Mu2e ST, Figure 2.17, is composed of 37 annular aluminum foils of 75 mm of radius with a purity of above 99.99% [3]. The foils are 100 μm thick to minimize energy losses of the CEs. This design narrows the reconstructed CE momentum distribution and separates it from the DIO electron momentum distribution. The disks have a central hole of 21 mm of radius. The annular design minimizes interactions with the beam electrons, reducing the radiation load by $\sim 30\%$. While the

central hole does impact the target's capacity to halt muons, which move in helical patterns, this effect is relatively minor, representing a necessary compromise in the design.

The space between each disk is 22.2 mm, bringing the total length of the ST to ~ 80 cm. Muons passing through the hole of an upstream foil will stop in a downstream layer. The design of the ST has been chosen by considering two conflicting physical requirements. Firstly, the target must be thick enough to stop a significant fraction of the muons, with the current version achieving an indicative stopping fraction of 30%.

Secondly, the target must be thin enough to control the energy loss of the CEs, which is why the target has a segmented geometry.




Figure 2.17: The ST design. Each disk is supported by three cables not shown in the picture

There are various factors to consider while selecting aluminum as the ST material. First, as described in Section 2.2.2, the aluminum target has a lower RMC background. Moreover, the muonic aluminium atom has a quite long lifetime, as shown in Figure 2.18a. The long lifetime allows separation between prompt backgrounds and a live window with a good decay rate. The muon DIO endpoint energy, further, depends on the type of nucleus, as shown in Figure 2.18b. Aluminum has a high endpoint energy, so when muons are captured on other detector materials with a higher atomic number Z , they have lower endpoint energies and do not contribute to background. For these reasons, aluminum is a suitable ST material for muon-to-electron conversion searches. Moreover, the branching ratio (BR) of the conversion varies depending on the ST material due to differences in atomic number (Z) and mass number (A). If a $\mu^- \rightarrow e^-$ conversion is observed, comparison of conversion BR s on different nuclei normalized to aluminum could allow to identify the dominating operator type, such as scalar (S), dipole (D), vector of transition charge radius type and vector of effective Z -penguin type, ($V(\gamma)$) and ($V(Z)$) respectively. Despite challenges in separating prompt backgrounds from conversion signal due to short lifetimes of muonic atoms, materials with higher Z offer better model differ-

entiation. If the Mu2e experiment observes a conversion signal, a subsequent search could make use of titanium as the ST. A more detailed discussion can be found in [66], [67] and [68].

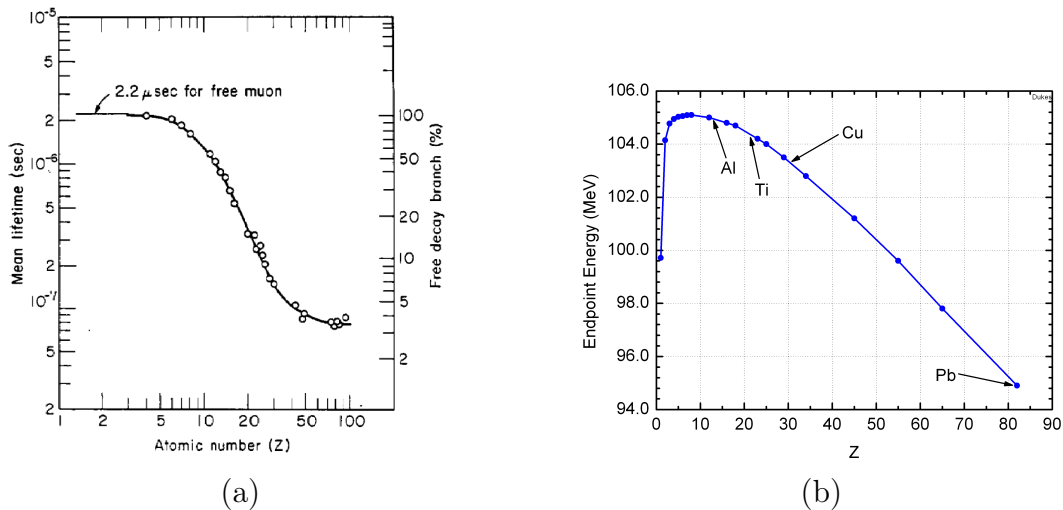


Figure 2.18: (a) The dependence of the mean lifetime and free decay fraction of the muonic $1s$ state on the atomic number of the nucleus to which the negative muon is bound [69], and (b) the dependence of the electron energy spectrum endpoint on the DIO [43].

2.7 The Mu2e detectors

Mu2e makes use of a set of complementary detectors to measure particles momentum and energy. These detectors are annular and situated within a solenoidal magnetic field of approximately 1 T along the z -axis, which aligns with the direction of the muon beam. This distinctive geometry is highly effective in reducing background, as it allows the passage of particles produced during muon capture, remnant beam, and electrons from the initial proton collision to the beam dump without striking the detector elements, that otherwise would cause excessive instantaneous detector occupancy and accumulated radiation damage. Most of the muons are typically too low momentum to exit the central region, preventing them from reaching the detector elements. As a result, the occupancy is kept at a manageable level. Due to acceptance limitations for particles originating from the ST, only particles with a momentum higher than ~ 80 MeV/c can be reconstructed. The Mu2e detection system comprises a straw tracker, followed by a calorimeter and a STM, all enclosed within the CRV.

2.7.1 The straw tracker

The Mu2e straw tracker is placed inside the DS downstream from the ST in a 1 T uniform magnetic field. The tracker is one of the most important Mu2e detectors: it must provide very good momentum resolution to disentangle the monochromatic CE signal from the background. Since the shape of the DIO spectrum near the endpoint decreases as $(E_{\mu e} - E_e)^5$, the corresponding background increases very quickly as the resolution degrades. To achieve the desired DIO suppression, the required momentum resolution should be below 1 MeV/c. Figure 2.19 shows the expected momentum resolution as determined from a full GEANT4 simulation of the detector, Front-End Electronics (FEE), and the DAQ. It is fundamental to have a thorough understanding of the tails of the momentum resolution. The right tail could push low-energy DIO electrons into the CE signal window. On the other hand, the left tail, which is primarily due to energy losses in the ST, in the IPA, in the OPA, and in the tracker itself, could push CEs below the signal window in the region dominated by DIOs thus reducing the signal efficiency. To minimise the probability




Figure 2.19: Momentum resolution of the straw tracker as determined for a sample of Monte Carlo CEs [3].

of scattering and the energy loss of the CE, the volume inside the DS is evacuated to 10^{-4} Torr. The Mu2e collaboration has selected the straw-tube technology [3] for its tracker, driven by the need for a detector with high core resolution, minimal high-side tails, and low energy loss. This choice led to straw tubes, which offer a combination of low mass, rapid drift times, good spatial resolution and they are in wide use in particle physics.

The detector has an annular shape with an active region covered by straws only between the internal radius of 380 mm and the external radius of 700 mm.

Chapter 3 will provide a more detailed description of the entire detector, including the mechanical structure, the FEE and DAQ system.

2.7.1.1 Momentum scale calibration

One of the most challenging aspects of the Mu2e experiment is achieving the required momentum resolution.

Mu2e requires at least one source of high-energy calibration electrons or positrons to independently measure the absolute momentum scale of the tracker. The momentum scale calibration relies on accurately reconstructing the $\pi^+ \rightarrow e^+\nu_e$ peak ($BR \sim 1.23 \times 10^{-4}$), with its success critically dependent on the suppression of $\mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ decays in flight (DIF). Several approaches are being explored:

1. A special run with adjusted TS collimators: operating with reversed TS collimators and a reduced DS field (70% of the nominal value), this setup is designed to capture the leptonic decay ($\pi^+ \rightarrow e^+\nu_e$) of stopping π^+ in the nominal ST;
2. Low intensity run with reduced DS magnetic field: running at low intensity and a further reduced DS field (50% of the nominal value), this approach aims to gather high-statistics data to accurately fit the Michel edge of the DIO momentum distribution.

Several important considerations include:

- Early time measurement: to reconstruct the $\pi^+ \rightarrow e^+\nu_e$ decays, measurements must be performed at early times, $T \sim 300$ ns. This necessitates operating at a reduced proton beam intensity to mitigate the pileup;
- Special instrumentation requirement: to minimize the background from muon decays in flight, a beam degrader in the DS is needed.

The momentum scale calibration described involves two key measurements:

1. The edge of the positron momentum spectrum from $\mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ decays at $B = 0.5$ T;
2. The $\pi^+ \rightarrow e^+\nu_e$ peak at $B = 0.7$ T.

Each of these measurements determines the momentum scale at their respective magnetic field values. By combining these measurements, we can extrapolate the momentum scale to the nominal magnetic field, $B = 1.0$ T. This concept is illustrated in Figure 2.20.

The calibration scheme assumes that the magnetic field is initially set based on readings from an NMR probe.

2.7.2 The electromagnetic calorimeter

The Mu2e electromagnetic calorimeter serves multiple purposes [70]:




Figure 2.20: Illustration of the momentum calibration scheme.

- measurement of the energy deposited in the calorimeter allows to separate electrons and muons of the same momentum. The energy-to-momentum ratio (E/p) will be used for particle identification and to suppress the background due to cosmic muons;
- the calorimeter information helps improving the pattern recognition and the quality of the track reconstruction;
- it provides triggers independent from the track-based triggers [71].

For these purposes, the calorimeter needs to fulfill the following physical requirements:

- an energy resolution better than $\sigma_E/E \sim 10\%$ allows to reach a rejection factor at the level of 200 between CE and the ~ 40 MeV energy deposit from 105 MeV/c cosmic ray muons mimicking the signal;
- a timing resolution better than 500 ps ensures that the energy depositions in the calorimeter are in time with the CEs reconstructed by the tracker and also improves the PID;
- a position resolution of ~ 6 mm allows to match the position of the energy deposit with the extrapolated trajectory of a reconstructed track;
- it should perform a fast enough response in order to handle the experimental high rate ($\tau < 40$ ns);
- a temperature and gain stability within $\pm 0.5\%$ are required not to deteriorate the energy resolution;
- it must be capable of operating in a magnetic field of 1 T, a pressure of 10^{-4} torr, a neutron flux equivalent to 10^{12} n/cm²/year, and in a high radiation environment (exposures up to 15 krad/year);

- reliability and redundancy are needed to operate in vacuum for one year without any interruptions.

The calorimeter disks are being assembled at Fermilab's SiDet facility, as shown in Figure 2.21. The components will be described in the next sections.




Figure 2.21: Mu2e calorimeter disks being assembled at the Fermilab SiDet facility.

2.7.2.1 The calorimeter mechanical structure

The calorimeter is located in the DS downstream of the tracker (Figure 2.15) It consists of two annular disks, with an internal radius of 35 cm and an external radius of 66 cm, located at a relative distance of 70 cm (Figure 2.22). The separation was chosen to match half of the distance between two periods of the helical trajectory of a typical 105 MeV CE. This maximises the detection probability of CEs: electrons which pass through the hole of the first disk will be detector by the second one [72].

The calorimeter mechanical structure was designed to support the layout of the crystals by piling them up in a self-standing array organized in consecutive staggered rows. Each crystal array is supported by two coaxial cylinders. The inner cylinder must be as thin and light as possible in order to minimize the passive material in the region where spiraling background electrons are concentrated. The outer cylinder is as robust as required to support the load of the crystals (700 kg). Each disk has two cover plates. The plate facing the beam is made of carbon fiber to minimize the degradation of the electron energy, while the back plate should also be robust to support the SiPMs, the FEE and the SiPM cooling lines and it is therefore made of polyether ether ketone (PEEK). The crystal arrangement is self-supporting, with the load carried primarily by the outer ring. The heat generated by SiPMs, FEE and read out electronics must be removed within temperature values acceptable for the correct operation of each device. Furthermore, the difficult access to components requires a cooling system free of fault and maintenance needs for at least one year. The cooling system has to maintain SiPM temperature at approximately -10°C to minimize the dark current: this is obtained by circulating a refrigerating fluid at approximately -15°C .




Figure 2.22: Calorimeter disks installed in the Mu2e experimental area (CAD representation) [72].

2.7.2.2 The undoped Cesium Iodide crystals

Each calorimeter disk is filled with 674 undoped cesium iodide (CsI) crystals [71], wrapped in a $150 \mu\text{m}$ Tyvek material.

Undoped CsI represents the best compromise between cost, reliability, performance and radiation hardness, providing a fast emission time and a sufficiently high light yield. Pure CsI has a wavelength peak emission at 310 nm, a scintillation time of 20 ns, and a light yield of about 2000 photons/MeV. Figure 2.23 (top right) shows some CsI crystals: the dimensions are $34 \times 34 \times 200 \text{ mm}^3$. The crystal length is approximately $10X_0$ which is sufficient to contain the 105 MeV CE shower since the average incident angle is 50° . Undoped CsI crystals also have the radiation hardness adequate for the Mu2e operational conditions.




Figure 2.23: (Left): Schematic CAD representation of the two calorimeter disks. (Top right): cesium iodided crystal used in the calorimeter. (Botto, right): Two SiPM arrays glued onto the copper holder (left) and one fully assembled readout unit made of two SiPM arrays and two Front-End boards mounted on the copper holder (right) [70].

2.7.2.3 The calorimeter readout chain

The readout of the calorimeter is performed through a chain of components: Silicon PhotoMultipliers (SiPM), FEE, Mezzanine board and Digitizer board. In the following, a brief description of each component will be provided.

2.7.2.3.1 Silicon PhotoMultipliers (SiPM) Since the calorimeter will operate in a 1 T magnetic field, SiPMs were chosen over PMTs (Figure 2.25) [73]. To match the scintillation emission of 310 nm to the SiPM photon detection efficiency, UV extended Hamamatsu SiPMs with a front window made of a silicon resin were selected. Each SiPM is composed of a 2×3 array array of individual $6 \text{ mm} \times 6 \text{ mm}$ cells. The array can be seen as the parallel of two series of 3 cells, that are powered by the same source. The mixed configuration is motivated by two competing requirements: a parallel connection has greater capacitance and a slower response time but requires a lower bias voltage, whereas a series connection has the opposite characteristics. To improve reliability, light collection efficiency and resolution, each single crystal is readout by two custom SiPMs arrays. Magnetic field resistance was the factor in choosing between SiPMs and PMTs. The good efficiency of the SiPMs allows to collect approximately 20 p.e./MeV per SiPM. To operate in vacuum, minimise outgassing contributions, reduce thermal coupling between crystals and FEE, the crystal-SiPM coupling is done without any optical grease and a 2 mm air gap is maintained between the crystal and the readout sensor.

Each crystal is coupled with a readout module consisting of two ultraviolet (UV)-extended SiPM arrays and the corresponding analog FEE boards, as illustrated in Figure 2.23 (bottom right) [74], [75] and [76]. As shown in Figure 2.24, each board is positioned within the electronic crates that surround the calorimeter disks.




Figure 2.24: Left: the breakout of calorimeter mechanical components. Top right: the breakdown of front panel plate. Bottom right: the Mezzanine and DiRAC boards [70].

2.7.2.3.2 The Front-End electronics Each SiPM is read by one Front-End board directly linked to the pins on the back of the SiPMs and, via cable, to the Mez-






Figure 2.25: The design of a 2×3 array of SiPM cells.

zantine board (Figure 2.26a). The Front-End board provides signal preamplification and shaping, local linear regulation of the SiPM bias voltage, and monitoring of the SiPM current and temperature. The two Front-End boards connected to the two SiPMs reading the same CsI crystal are housed in a copper holder (Figure 2.26b), which serves as a mechanical support and ensures the SiPM cooling. An external Faraday cage reduces the noise. The holder also holds an optical fiber used to pulse the crystal with a laser light to calibrate the energy and time response.



(a)



(b)

Figure 2.26: (a): The FEE board, with its connectors to SiPMs back and to the cable coming from Mezzanine. (b): The assembled FEE holder.

2.7.2.3.3 The Mezzanine board The Mezzanine Board serves as the interface between 20 different Front-End boards and one Digitizer (Figure 2.27) The Mezzanine sets all FEE parameters (such as High Voltage (HV)) and monitors all Front-End boards and SiPM parameters (such as HV, temperature, and current). It also provides low voltage to the Front-End boards. The HV of each SiPM can be independently set. These tasks are performed by an ARM processor located on the Mezzanine board. The cables chosen to link the Front-End boards to the Mezzanine board must be radiation-resistant, shielded, and have low outgassing.




Figure 2.27: Left: The Mezzanine Board with connectors highlighted. Right: The Cable linking 4 FEEs to one Mezzanine.

2.7.2.3.4 The digitizer Board The Digitizer Readout Controller (DiRAC) digitizes the SiPM signals received from the Mezzanine Board and manages data transmission to the experiment’s DAQ system. Zero suppressed data are sent from the DiRAC board to the DAQ via optical fibers. A sampling rate of 200 MHz (one sample every 5 ns) and a 12-bit resolution satisfy the calorimeter requirements. The two SiPMs which read the same crystal are always connected to different DiRAC boards hosted in separate crates to prevent a complete loss of a crystal in an event of a single crate fault.

2.7.2.3.5 The crate The Mezzanine and DiRAC boards reside in crates. Each crate contains up to eight Mezzanine and DiRAC boards and provides cooling to the boards via a network of cooling lines grooved in the crate walls. The thermal contact between the electronics components on the Mezzanine and the DiRAC and the crate walls is ensured by a copper plate covering the entire board.

2.7.2.4 Module-0 prototype

To validate the chosen design before commencing the mass production of the calorimeter components, a large-scale prototype was constructed and tested with an electron beam at the Beam Test Facility of the National Laboratories of Frascati. The Module-0 consists of 52 undoped CsI crystals, read by 102 SiPMs connected to FEE boards. Its mechanics was designed

to closely resemble the final calorimeter, allowing for the testing of the assembly procedures and cooling. The energy and time resolution are shown in Figure 2.28 [3] and [77].

2.7.3 The Cosmic Ray Veto

As reported in Section 2.2.2, during Mu2e operations, cosmic rays will strike the detector area at a rate of approximately 15,000/s, resulting in about one fake CE signal per day.




Figure 2.28: The energy and time resolution of the Module-0 for the Mu2e calorimeter for electrons at the conversion energy. (Top left) The resolution for electrons striking the array at normal incidence. (Top right) Resolution for electrons striking the array at 50° with respect to the face. (Bottom) The time resolution. All measurements are compared with simulations (red line) [3] and [77].

To minimize the backgrounds, Mu2e exploits a combination of passive shielding and an active cosmic ray veto system (CRV). Figure 2.29 shows a schematic representation of the CRV covering the region of the DS and a half of the Transport Solenoid almost entirely (there is no veto on the floor at the bottom), for a total detector area of 327 m^2 . The CRV modules are manufactured from plastic scintillator extrusions. In different regions of the detector the extrusions may have different lengths, but they all have the same cross-sectional area of $5 \times 2 \text{ cm}^2$. The extrusions are coated with Titanium dioxide to improve internal reflections and thus, the light yield. Two holes are extruded inside the scintillator bar throughout its length to contain two 1.4 mm diameter wavelength shifting fibers. The fibers transport light to the extrusion ends where each fiber is instrumented with a $2 \times 2 \text{ mm}^2$ SiPM at each end. Figure 2.30 shows the cross-section of a CRV module. Each module is made of four overlapping staggered layers of plastic scintillator counters to reduce the effects of gaps. The layers are separated by 10 mm aluminium plates used as absorbers. When three of the four layers of the CRV are activated, an offline veto window is opened. A conversion-like event observed during this time window is assumed to be produced by a cosmic-ray muon and will be disregarded. The estimated total dead time is $\sim 5\%$. Based on the simulation, the required CRV veto efficiency is 99.99% within the detector fiducial.




Figure 2.29: The CRV covers the DS entirely and half of the Transport Solenoid. It is made of several sectors, including the top (CRV-T), the right (CRV-R), the left (CRV-L), the upstream (CRV-U), the downstream (CRV-D), and the hole where the Transport Solenoid enters the enclosure (TS-hole) [1].




Figure 2.30: Cross-section of a CRV module, including geometry and nomenclature. There are internal gaps between counters in a di-counter [78].

2.7.4 The Muon Beam Stop and the Stopping Target Monitor

A Muon Beam Stop (MBS) is installed at the downstream end of the DS to absorb muons that are not stopped in the ST, shown in Figure 2.31. The MBS has been designed to limit the effects of muon decays and captures. The magnetic field gradient prevents the majority of the low-energy charged particles created in the MBS from moving upstream towards the detectors. The MBS is made of high- Z minerals and polyethylene. Muons have an extremely short lifetime in high- Z materials, as shown in Figure 2.18a, which minimizes the effects of muon decays and captures much earlier than the Mu2e live window starting at 640 ns. Moreover, polyethylene absorbs protons and neutrons emitted by the excited nuclei during muon captures.

The Mu2e STM has been designed to determine the number of muons in the ST within 10% precision over the course of the experiment. It is made of two solid state High Purity Germanium (HPGe) crystal detectors and one scintillating lanthanum

bromide (LaBr₃) detectors. The STM is placed downstream from the MBS, as shown in Figure 2.31, to measure the number of muons stopped on the ST.





Figure 2.31: The STM geometry showing the DS region (left), the End Cap Shielding, sweeper magnet and STM field-of-view collimator. At the far end of the hall (right) is the final *spot – size* collimator and the STM detector.

The STM measurement technique exploits the 347 keV X-ray line produced in the $2p \rightarrow 1s$ radiative transition of the muon moving to the ground state in the muonic atom (Figure 2.32), which is present for 80% of the stopped muons. Although the $2p \rightarrow 1s$ transition has the largest yield, there are also additional X-ray transitions with substantial yields, including $3p \rightarrow 1s$ and $4p \rightarrow 1s$. In addition, the $3d \rightarrow 2p$ transition that populates the $2p$ state appears in the energy spectrum.


An alternative option is to measure gamma rays associated to other types of muon interactions in the ST (Figure 2.32). These are unambiguous signals that muons stopped in the aluminum target. The main signatures are the following ones:

- 1809 keV gamma-ray that is emitted in the nuclear capture chain:



The ${}^{26}\text{Mg}^*$ mean de-excitation time is negligible (476 fs) with respect to the muonic atom lifetime (864 ns). This process occurs for 51% of the nuclear captures.

- A 844 keV gamma-ray coming from the decay of long-lived ($\tau \sim 9.5$ min) isotopes produced in muon nuclear captures:



This process occurs for the 9.2% of nuclear captures.

The HPGe detectors have the energy resolution (FWHM) of 2.5 keV in the energy range of 300-2000 keV and are used to measure photons produced by the secondary




Figure 2.32: Reactions that can be exploited to detect the muon stops and nuclear captures in the Aluminum ST.

muonic aluminium orbital transitions (347 keV) and nuclear captures (884 keV). The LaBr_3 detector has an order of magnitude worse energy resolution than the HPGe detector, but it provides a much higher rate capability. It is used to detect prompt photons produced in the nuclear captures (1809 keV). The STM geometry is shown




Figure 2.33: The STM geometry.

in Figure 2.33.

Chapter 3

The straw tracker

This Chapter provides a brief introduction on the working principles of a drift tube and then a detailed description of the Mu2e straw tracker. This includes a description of the tracker mechanical structures (straw, panels, planes and stations) and the front-end and DAQ electronics. Most of the discussion is based on [79] and [3].

3.1 Drift tubes

Gaseous detectors provide spatial resolution which could be as good as $50 \mu\text{m}$ and high detection efficiency at a low cost [79]. There are many different gas ionization detectors, one of which is the drift tube. The basic configuration of a drift tube is shown in Figure 3.1. A hollow cylindrical conducting tube is grounded and serves as the cathode. The tube is filled with a combination of a noble gas, often Argon, and a quench gas. A thin sensing wire is positioned along the cylindrical cathode axis. The wire, called anode, receives a high voltage.




Figure 3.1: Schematic view of a drift tube [79].

Assuming an anode radius a , a cathode radius b and using Gauss theorem, the electric field is:

$$E(r) = \frac{1}{r} \frac{\lambda}{2\pi\epsilon} = \frac{1}{r} \frac{V}{\ln(b/a)} \quad (a < r < b) \quad (3.1)$$

Where λ is linear charge density on the wire and ϵ is the dielectric constant of the gas.

Lower diameter wires are typically preferred in drift tubes. A higher electric field near the wire increases the amplification factor of the drift tube at the same voltage. Additionally, smaller diameter of a signal wires improves the spatial resolution [79].

3.1.1 Gas ionization

There are two types of interactions that can deposit energy as particles traverse the gas volume: ionization and excitation. Collisions between a charged particle C and an atom A can result in the ejection of one or more electrons: $A C \rightarrow A^+ e^- C$, or $A C \rightarrow A^{++} e^- e^- C$, if more than one electron is released. This process is called primary ionisation. The mean energy loss per path length can be determined using the Bethe-Bloch formula. A noble gas atom A can turn also into an excited state A^* through the interaction $A C \rightarrow A^* C$. In case of a compound gas (A+B), if the excitation energy of A^* is higher than the ionization potential of B , the quencher, the Penning Effect can produce ionization through $A^* B \rightarrow AB^+ e^-$. In addition, noble gases can also form molecular ions through processes such as $A^* A \rightarrow A_2^* \rightarrow A_2^+ e^-$. A secondary ionisation can occur through these processes or through electrons that have sufficient energy for generating more ions. To compute the average number of electron-ion pairs produced by the initial particle, divide its total energy loss by the average energy required to make an electron-ion pair. Due to the energy lost during excitation, this average does not match the gas ionisation potential. Measurements showed an average of one electron-ion pair every 30 eV, with variations depending on gas composition and starting particle. Except for very slow particles, this value remains constant regardless of their initial energy. Without an electric field, electrons and ions created during ionisations spread uniformly. Collisions cause them to lose energy and eventually reach thermal equilibrium with the surrounding gas. They eventually recombine. An electron maximal range during ionisation is correlated with its initial kinetic energy. In a normal temperature and pressure gas, a 10 keV electron may be stopped in approximately 1 mm. Ionisation electrons often have lower kinetic energy, leading to a shorter range.

3.1.2 Drift of ions and electrons

An electric field accelerates free electrons and ions towards the anode and cathode along the field lines. As these particles accelerate, they scatter on other particles in the gas, losing energy. The directions of motion are randomised, and maximum speeds are set. As a result, these particles move uniformly along the electric field. This is referred to as the drift velocity of the charge. It is superimposed with the thermal motion. Drift velocities for ions and electrons varies based on several parameters. Ions have a larger mass than electrons and their masses are comparable to those of gas molecules. During collisions, gas molecules absorb a significant portion of the energy gained from ions during acceleration. In a drift tube detector just a little amount of electric field energy enters the energy associated with ion motion, making it comparable to the initial thermal energy before acceleration. The

ion drift velocity v_i is proportional to the reduced electric field E/N where N is the number density of the gas and it is typically $\mathcal{O}(10^{-2}) \mu\text{m}/\text{ns}$, except in the region near to the anode wire where a stronger electric field is present. The ion thermal velocity is typically $\mathcal{O}(10^{-1}) \mu\text{m}/\text{ns}$ at room temperature. On the other end, only a small fraction of the energy is released during elastic collision from electrons, so they acquire more energy from the electric field than their thermal energy. Many different factors impact electron drift velocity. Some gas molecules, such as H_2O or CO_2 , can interact with electrons to produce negative ions due to their great electron affinity. In rare situations, electrons gather enough energy to exceed gas molecule excitation threshold, resulting in inelastic collisions. Electron drift velocity is a complex function of electric field intensity due to several variables influencing electron collisions across a large energy range. Figure 3.2 shows the electron drift velocity at different electric field strengths in argon-carbon dioxide mixtures (Ar- CO_2) of different proportions [80]. Electron drift velocities in drift tubes are typically




Figure 3.2: Electron drift velocity versus electric field in Ar: CO_2 mixtures of different proportions [80]. 80%:20% Ar: CO_2 mixture is the gas used in the Mu2e tracker.

$\mathcal{O}(10) \mu\text{m}/\text{ns}$. This is significantly higher than ion drift velocities and comparable to electron thermal velocity under the same conditions. The radial coordinate of an ionisation can be computed using the electron drift velocity and time. Since drifting electrons and ions are scattered on gas molecules, they also diffuse along their trajectory. Electrons diffuse significantly quicker than ions because of their high velocity. Electron diffusion limits the intrinsic resolution of drift tubes used to

measure incoming particle coordinates. CO_2 has internal degrees of freedom at low collision energies, preventing electron energies from exceeding thermal energy until field strengths above about 2 kV/cm. This improves intrinsic spatial resolution.

3.1.3 Avalanche multiplication

Electrons can ionise when they face a high electric field near the anode wire. The released secondary electrons form tertiary electrons, and so on. The number of free electrons rapidly increases, resulting in an avalanche. In a drift tube, where the electron mean free path is about the order of μm , an avalanche develops when the electric field approaches $\mathcal{O}(10)$ kV/cm. According to Equation 3.1, using $a \sim \mathcal{O}(10^{-3})$ cm, $b \sim \mathcal{O}(1)$ cm and a normal voltage of 1-2 kV, the avalanches can occur within $\mathcal{O}(100)$ μm from the anode wire. Electrons from the avalanche are collected on the anode wire within 1 ns, while positively charged ions move towards the cathode. Drifting ions mostly generate signals in electrodes via induction. Figure 3.3 shows the model of an ionisation avalanche.




Figure 3.3: The model of an ionisation avalanche forming at the anode wire of a proportional tube or chamber [79]. (a) In the drift volume, electrons and ions are generated and drift to their corresponding electrodes. (b) Near the wire, the electron achieves a high enough field to induce secondary ionisation, resulting in an avalanche. (c) The electric field separates charges created during an avalanche. (d) Electrons have higher lateral diffusion than ions, causing the avalanche to expand around the wire and produce a positive charge cloud in the shape of a drop. (e) Electrons from the avalanche reach the anode within nanoseconds, but ions take longer, up to ms, to reach the cathode.

3.1.3.1 Avalanche gain

When an avalanche develops, the amplification factor is around $10^4 - 10^6$. The number of electrons freed per unit path length is given by the first Townsend ionization coefficient $\alpha = \sigma_{ion}n = 1/\lambda_{ion}$ and this depends on the electric field E , as a higher electric field corresponds to a higher kinetic energy of the electron, that increases the ionization cross-section. The increase dN of the number of electron-ion pairs

over a path length ds is [79]:

$$dN = \alpha(E)Nds \quad (3.2)$$

Solving this equation, we can easily obtain the gas amplification G :

$$G = \frac{N(s_a)}{N_0} = \exp\left(\int_{s_0}^{s_a} \alpha(E(s))ds\right) = \exp\left(\int_{E_{min}}^{E(a)} \frac{\alpha(E(s))}{dE/ds} dE\right) \quad (3.3)$$

where N_0 corresponds to unamplified electrons in $s = s_0$ and E_{min} corresponds to the minimum energy for ionisation to occur. The energy distribution depends on the electric field which is position dependent. Since the free path is inversely proportional to the particle density in the gas, $E_{min}(\rho) = E_{min}(\rho_0)\rho/\rho_0$. It is reasonable to say that the coefficient is proportional to the field strength, $\alpha = \beta E$, in the low field region. Adding this relation with Equation 3.1 and 3.3:

$$\ln(G) = \beta a E(a) \ln\left(\frac{E(a)}{E_{min}}\right) \quad (3.4)$$

where β can be related to w_i , that is the energy spent for one ionisation and its value is equal to $e\Delta V$. As the voltage drop per unit path length is $dV = E(s)ds = (\alpha/\beta)ds$, we obtain $dN = N\beta dV$. Integrating, we can see that $\beta = \ln(2)/\Delta V$, so the gain in a drift tube is:

$$\ln(G) = \frac{\ln(2)}{\Delta V} a E(a) \ln\left(\frac{E(a)\rho_0}{E_{min}(\rho_0)\rho}\right) \quad E(a) = \frac{V}{a \cdot \ln(b/a)} \quad (3.5)$$

which is the Diethorn's formula. Gain measurements with variable ρ/ρ_0 , a , and $E(a)$ can provide the parameters $E_{min}(\rho_0)$ and ΔV . The gas temperature T , pressure P and operating voltage V significantly impact the gain of a drift tube and gas mixture.

3.1.3.2 Quench gas

To avoid subsequent avalanches, the drift tube gas combination may contain a quench gas, such as CO₂, methane, or other hydrocarbons. During an avalanche, photons are created by gas deexcitation and electron attachment to electronegative species, resulting in negative ions. Photons can generate ionisations outside the primary avalanche zone or create free electrons on the cathode surface, resulting in secondary avalanches. The difficulty arises when the signal created is not proportional to the deposited energy by the original particle and is no longer localised to the energy deposition point. Enough intense photons can induce a chain reaction of secondary avalanches, leading to a continuous discharge. The use of quench gas prevents subsequent avalanches by absorbing ionising photons before they travel far. A tiny quantity of quench gas during normal operation can significantly decrease secondary avalanches and breakdowns.

3.1.3.3 Operation modes of gaseous ionization detectors

In drift tube detectors, the number of electron-ion pairs formed during an avalanche is proportional to the starting number of electrons, as shown in the gain computation. To operate in a proportional mode, an appropriate voltage is needed to reduce the effects of avalanche charges on the electric field. Figure 3.4 illustrates how a gaseous ionisation detector may work in multiple modes based on the operating voltage. Higher operating voltage leads to higher charges on the electrodes. Low voltage power supply causes ionisation charges to recombine before reaching electrodes, leading to no signal collection. At higher voltages, in ionisation chamber region, charges can drift to electrodes, but the electric field is insufficient for avalanches to occur. Increasing the operational voltage leads to drift tubes and proportional counters. When the voltage becomes high enough, proportionality is lost. When electrons from an avalanche are collected, the high density of positive ions near the anode might affect the electric field. Electrons in future avalanches that enter the area between the positive ion cloud and the wire face a decreased electric field, resulting in lower amplification. The electric field becomes greater in the tail of the avalanche, which is far from the wire than the ion cloud. This range of voltage is called region of limited proportionality. When the operating voltage reaches high values, avalanches create photons of sufficient energy to cause secondary avalanches that propagate across the detector, independently from the quench gas. This results in detector saturating the output. This way of operating is called breakdown mode, commonly known as the Geiger-Muller mode.




Figure 3.4: The dependence of output signal of a counting tube on the applied voltage in gaseous ionisation detectors. The numbers on the axes are for orders of magnitude only and they depend on the device geometry and gas concentration. Drift tubes operate in the proportional mode [79].

3.1.4 Signal creation and propagation

Drift tube signals do not come directly from avalanche charges. If they did, the anode wire would receive the entire signal in just a few ns. Instead, signals are generated due to the movement of charges on the electrodes, influenced by the mobility of electrons and ions. The Shockley-Ramo theorem [79], can be used to determine the induced charge and current. One of the key results of this theorem is that the total induced charge of a moving charge is determined by the initial and final positions only. A charge pair induces the same amount of charge on an electrode as the charge collected on it. Furthermore, if all electrodes are treated as an unity, their weighted potential will be one. If one electrode completely encloses the others, the weighted field in the contained region is always zero. This means that the total induced current across all electrodes is always zero. Applying this to the drift tube, the theorem helps us analyze the induced current signal produced by an avalanche of electron-ion pairs on the anode wire. When compared to the total charge induced $Q_{tot}^{ind} = -eN$, electron mobility in the avalanche accounts for just 1-2%. Positive ion drift from the avalanche accounts for the majority of the signal. A signal propagates to both ends of the drift tube from the avalanche location. Signals with distinct frequency components can propagate at varying velocities. This causes the signal to disperse. For more details, refer to [79].

3.2 The Mu2e straw tracker

The sensitive elements of the Mu2e tracker are the straw tubes with a diameter of 5 mm and a length varying in between 40 cm and 110 cm, filled with a 80%:20% Ar:CO₂ mixture at 1 atm. The detector has a modular design, and it is made of basic elements named *Panel*, *Face*, *Plane* and *Station*, shown in Figure 3.5. The panel is the fundamental unit of the detector: 96 straw tubes arranged like harp cords in two staggered layers to form one harp-shaped panel spanning 120° along the azimuthal angle [1], shown in Figure 3.5 (Top left). Three rotated panels form a face, and two faces rotated of 30° form one plane, Figure 3.5 (Top right-left). Two planes are coupled to make a station, mounting the second rotated on the vertical axes by 180°. Figure 3.5 (Top right-right). One station is thus made of twelve panels. The entire tracker is made of 18 stations, Figure 3.5 (bottom) and has the total length of 3.2 m and the diameter of approximately 1.7 m. Each panel is equipped with its own front-end and DAQ electronics placed within the mechanical structure (colored in red).

3.2.1 The sensitive unit: the straw

The Mu2e tracker straw tubes use the same detection principles as the gaseous ionisation detectors [79], to meet the experiment requirements. Figure 3.6 shows an end of a straw tube, compared to a pencil.




Figure 3.5: (Top left) Schematic representation of one straw tracker panel. (Top right) Straw tracker plane (left) and station (right). (Bottom) Full assembly of the 18 stations and support structure [1].



Figure 3.6: One of the Mu2e straw tube (compared to a pencil) [81].

All straws have the same diameter of 5 mm, and the straw length varies between a minimum of 334 mm and a maximum of 1174 mm, depending on their position on the panel, as shown in Figure 3.9.

The straw is wound with two layers of 6 μm -thick metallized Mylar separated by one 3 μm layer of glue. The straw wall is thus 15 μm thick: this minimizes the amount of detector material and thus the total energy loss of electrons in the detector. Moreover, this minimises the probability of significant deflections of the electron trajectory which makes pattern recognition and track reconstruction much simpler both at trigger and offline level, and allows to achieve the required excellent momentum resolution. The straw tube anodes are made of gold-plated tungsten wires with a diameter of 25 μm . The straw and the anode wires are tensioned and work-hardened to minimise sagging effects. Figure 3.7 shows the straw termination mechanical structure that allows to hold the anode wire in place. To increase the straw mechanical strength, two brass tubes are joined to both ends of the straw using

a silver epoxy. To ensure electrical insulation of the anode wire, a kapton sleeve is




Figure 3.7: The straw termination, depicted both exploded and assembled, features a brass tube connected to the straw using silver epoxy. An insulator (green) is inserted into a brass tube (red) to prevent breakdown near the tube’s end. The sense wire is soldered into the brass pin and epoxied to the injection-molded plastic. Post-assembly, the brass tube facilitates connection to the cathode, while the brass pin enables connection to the anode.

inserted inside the brass tube. The kapton sleeve holds an injection-molded plastic insert which contains a semicylindrical duct that allows the gas flow into and out of the tube. The insert has a groove along its axis and a U-shaped brass anode pin inserted at the end. To avoid slippage, the anode wire is epoxied into the groove and soldered to the anode pin. A T-shaped pin protects the anode pin from breaking by covering the groove in the plastic insert. The pin protector is epoxied to the insert with an extra brass ring connecting them. A ground clip is silver epoxied to two adjacent straws on the brass tubes and rings to provide a shared ground connection. To optimise the shape of the electric field within the straw, assembly procedures have been developed to align the anode wires to the panels with a precision of at least $25\ \mu\text{m}$ in the radial direction and $50\ \mu\text{m}$ in the perpendicular direction. From the straws, signals are sent to a common preamp board via the anode pin pair and grounding clip. The tracker front end electronics will be described in Section 3.2.5.

3.2.2 The building blocks: from panels to the station

Groups of 96 straws are assembled in two staggered layers of 48 straws and make a panel (Figure 3.8a).

Each panel spans a 120° arc (Figure 3.9). The dual-layer geometry was chosen to maximize the total panel coverage. This simplifies pattern recognition and increases tracking robustness and efficiency. A 1.25 mm separation between two consecutive straws accommodates manufacturing tolerances and allows for the straw expansion due to the gas pressure. This requires each straw to be self-supporting across its length. Channels within a panel are numbered from 0, corresponding to the radially most internal and longest straw, to 95, corresponding to the most external and shortest straw. An expanded view of the panel edge is shown in Figure 3.8b.

The straws are filled with a gas mixture of 80%:20% Ar:CO₂ at 1 atm. Since the

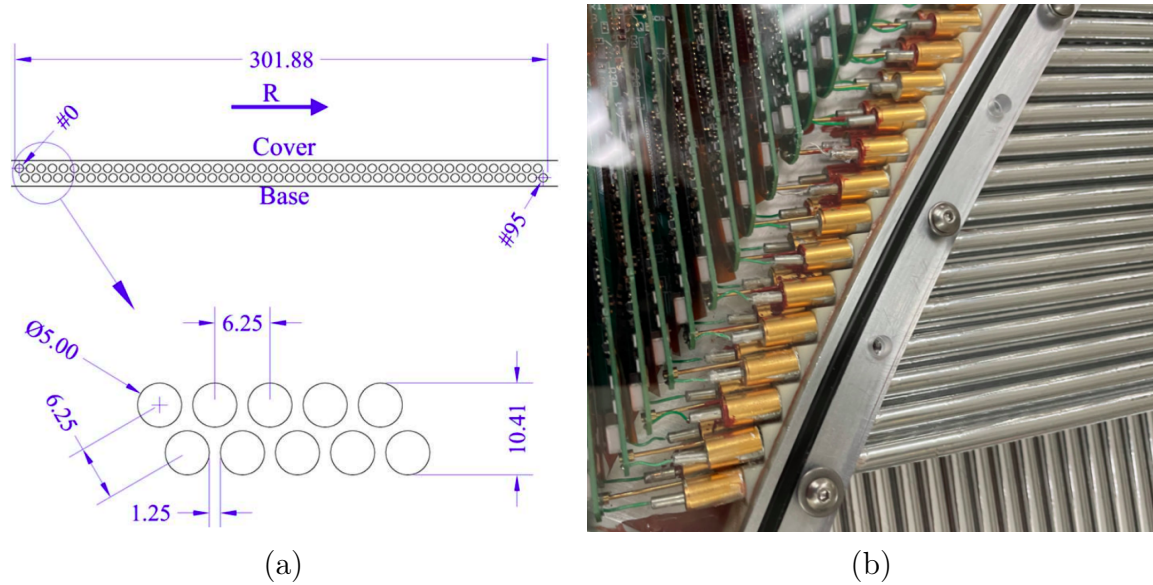


Figure 3.8: (a): The straw arrangement within a panel [81]. (b): Expanded view of the panel edge [81].




Figure 3.9: One of the fully assembled panel (Lab3 Fermilab Facility).

DS internal volume is evacuated to 10^{-4} Torr, the straws must endure such pressure difference. The overpressure is essential for the straw mechanical stability. Under normal temperature and pressure conditions, the panel must have an average leak rate below $0.014 \text{ cm}^3/\text{min}$. The nominal operating voltage of the straw is 1450 V. Preliminary tests showed that the straw tube gain is 1.25×10^4 at 1250 V and 7×10^4 at 1425 V. According to Diethorn's calculation, Equation 3.5, the gas gain at 1450 V is around 1×10^5 .

All panels are x-ray scanned to accurately measure and document the wire locations of straw tube channels.

Three panels rotated of 120° make a face and two faces rotated of 30° make a plane (Figure 3.5 (Upper right-left) and 3.10).

Once a plane has been assembled, a cooling ring is fitted around its outer circum-

ference. The front-end electronics is hosted in a ring around the outer circumference of the panel.

Two identical planes rotated of 180° around a vertical axis make one station (Figure 3.5 (Upper right-right)). During assembly, the second plane is rotated of 180° around the vertical axis.

At the moment of writing this Thesis, almost all tracker planes have already been assembled at Fermilab Lab3 Facility and the station assembly is progressing, as shown in Figure 3.10.




Figure 3.10: One tracker plane fully assembled at Fermilab Lab3 Facility.

3.2.3 The entire straw tracker detector

The entire detector is made of 18 stations (Figure 3.5 (bottom)) assembled together with a complex mechanical structure. Horizontal beams maintain longitudinal alignment of the rings. A thicker ring and two thinner rings placed at the downstream end of the detector stiffen the structure. The beams and the stiffening rings are made of stainless steel.

The detector rests on four bearing blocks attached to the stiffening rings and placed near the horizontal beams. The connection between the detector and the bearing blocks is kinetic, to avoid over-constraining and distorting the frame. The only constraint of the four points is along the vertical direction, and vertical adjustment screws allow to level and center the frame. Once fully assembled, a thorough mechanical survey of the detector will be performed before moving to the Mu2e experimental hall.

3.2.4 The drift tubes in magnetic field

By particle deflection in a magnetic field, it is possible to determine its momentum. In a drift tube, the sum of the electric drift field and magnetic field can produce several effects:

- change of the drift direction;
- change of the drift velocity (typically, the drift velocity decreases as the magnetic field strength increases);
- reduction of diffusion transverse to the magnetic field.

The magnitude of these effects depends on the type of drift gas used and the relative orientation of the magnetic field with respect to the electric field [79].

The drift velocity vector can be decomposed in three linearly independent components parallel to \vec{E} , \vec{B} and $\vec{E} \times \vec{B}$:

$$\vec{v}_D^B = -\frac{\mu^{B=0}}{1 + \omega^2 \tau^2} \left(\vec{E} + \frac{\vec{E} \times \vec{B}}{B} \omega \tau + \frac{(\vec{E} \vec{B}) \vec{B}}{B^2} \omega^2 \tau^2 \right) \quad (3.6)$$

where $\mu_B = 0$ is the mobility considering a zero magnetic field, $\omega = \frac{qB}{m}$ is the cyclotron frequency, and τ is the electron average collision time in the gas. Straws of the Mu2e tracker are orthogonal to the magnetic field and the angle between the drift direction and the ~ 1 T B -field varies from 0° to 90° . For this reason drift calculation is a complicated numerical problem. Moreover, as shown in Figure 3.2, in our straw gas, Ar:CO₂, v_d depends on the electric field. Numerous simulations have been conducted to account for this effect, revealing that the magnetic field significantly influences the drift time as a function of the particle radial distance from the anode wire (Figure 3.11). The difference is about 8 ns at a radial distance of 2.5 mm.




Figure 3.11: The total simulated drift time as a function of initial radial distance from the sense wire and initial angle of 90° . The black dotted line corresponds to the case with zero magnetic field, while the red one represents the case with $B = 1$ T.

3.2.5 The tracker front-end Electronics

The front-end electronics performs amplification, digitization and data packaging for transmission to the Data Acquisition System (DAQ). Each straw provides two hit times and one waveform, which are necessary to reconstruct the hit position along the straw and the hit charge, thereby allowing the use of constant fraction discrimination and mitigating time walk effects. A schematic representation of the entire system logic is reported in Figure 3.12. The plane front-end electronics is

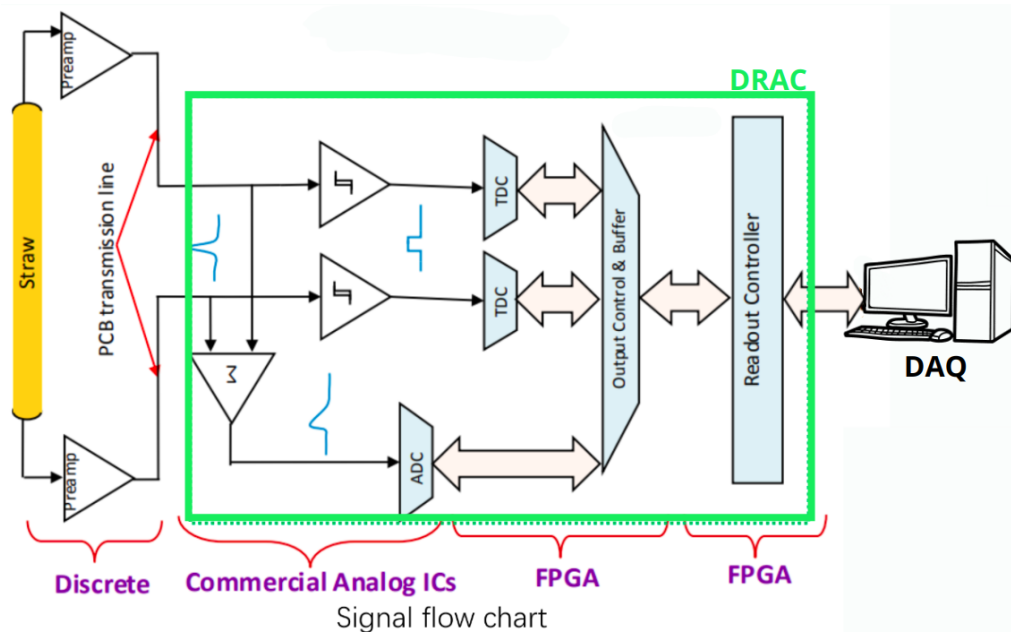


Figure 3.12: Signal flow through front-end electronics [1].

installed on several Printed Circuit Boards (PCBs) mounted on the outer section of the panel. The PCBs are shown in Figure 3.13 [82]. All PCBs are located in the outer section of the panel. The panel has two sides, the high-voltage side (HV-side), which handles the high-voltage supply distribution to the straws, and the calibration-side (CAL-side), whose FPGA is programmed to manage the generation of the test pulses.

Signals from the straw tube are first read out from both ends by the pre-amplifiers (preamps).

Each side of the panel has one Analog Mother Board (AMB) and one Jumper board, whose task consist of directing signals from the preamps towards the Digital Motherboard (DMB) positioned at the center, then to the Digitizer Readout and Assembler Controller (DRAC) (mounted on top of the DMB) to be processed and temporarily stored.

The AMBs and the DMB also handle the low voltage distribution, and the AMB on the HV side distributes the high voltage to the wires. The low voltage power supply is connected to the panel through the KEY board which contains also an optical fiber link and a JTAG connector. The AMB on the CAL side can inject

a calibration pulse into the wires. The boards are also equipped with sensors to monitor environmental parameters, including temperature, pressure and humidity.

In addition, the frontend components were qualified to sustain high level of radiation.

The first step of signal processing is the signal pre-amplification performed by the preamps mounted on two AMBs placed on the two lateral sides of the panel. The analog signals are transmitted from the preamps via microstrip transmission lines and through a Jumper board to the DRAC. Most of the data processing functions are performed by the DRAC board. The analog signals are processed by the Time-to-Digital Converters (TDCs) and by the Analog-to-Digital Converters (ADCs) installed on the DRAC. Digitized data are then transferred to the Readout Controller (ROC) for data packaging and transmission to the Mu2e DAQ (Figure 3.12).

To summarize, each straw tube is an independent detector channel instrumented with:

- two preamp channels, one for each end;
- two TDC channels, one for each end (192 in total);
- one ADC channel, measuring sum of both ends (96 in total);
- one high voltage feed.




Figure 3.13: Overview of the straw tracker front-end electronics [82]. The CAL-Side of the panel is not shown.

3.2.5.1 The preamplifiers

All 96 straw tubes are read out from both ends. Two adjacent straws are connected to the same preamp PCB board. These boards are mounted vertically on the AMBs. Thus, each preamp PCB board contains two preamps, each one connected to one single straw end. Each panel is thus equipped with 48 preamp PCB boards on the HV side and 48 on the CAL side, 192 preamps in total. To minimise signal reflections, the preamps have a matching $300\ \Omega$ input impedance. The function of the preamps is to convert the straw tube current signals into voltage signals which are amplified and shaped.

3.2.5.2 The Digitizer Readout and Assembler Controller board

The *brain* of the panel is called Digitizer Readout and Assembler Controller (DRAC) board. The DRAC performs digitization, packaging, temporary storage and data transfer to the Mu2e DAQ system. It also controls all panel operations. The schematics of the Figure 3.14 shows a picture of the board: comparators and ADCs, three DDR3 memory chips and three Field-Programmable Gate Arrays (FPGAs) are clearly visible. The two FPGAs on the left and on the right are named digi-FPGAs: each of them receives data from 48 straw channels, performs data monitoring, buffering and assembles the data packets which are then transferred to the central FPGA named Readout Controller (ROC) 3.2.5.2. The ROC handles communication, monitors slow control variables and controls all panel operations.




Figure 3.14: DRAC board schematics [83]. The DRAC board is the brain of the tracker panel. ADCs, FPGAs, DDR3 memories and comparators are also shown.

Figure 3.12 shows the data flow through the panel front-end and the DRAC:

- The two signals from the two sides of a straw tube are transmitted from the preamps to the DRAC;

- In the DRAC, the two biased signals are fed to zero-crossing comparators, which generate square waveforms if the signals are above their respective thresholds;
- The squared waveforms are transmitted to 16-bit TDCs implemented in the firmware in the digi-FPGAs. Timing digitization bin is 20 ps, including the determination of the arrival time and the time over threshold. With the intrinsic TDC resolution of 25 ps, taking into account also the comparator jitter, the noise and other external effects, the resulting time resolution is of the order of 70 ps. At run time, a hit on a straw is considered only if both ends of the same straw simultaneously have a signal above threshold;
- Signals from the two ends of the same straw are also summed up, the sum is digitized by a 10-bit ADC at 40 MHz and transmitted to the digi-FPGA;
- The digi-FPGA creates one data packet for each hit containing the TDC and ADC information;
- The data packets are transferred to the ROC 3.2.5.2, and temporarily stored in the DDR3 memory for later access by the DAQ system. There is a total of 8 Gb of memory space available on each DRAC. The main function of the ROC is to collect data from the digitizer boards digi-FPGAs, buffer data and transfer them to DAQ. ROCs continuously stream out the zero-suppressed data collected between two proton pulses to the Data Transfer Controllers (DTCs) (Section 3.2.5.4) [84]. Buffering is necessary, since the DAQ is not able to handle the instantaneous on-spill data rate and the off-spill time will be used to level the rate. The communication is flexible, thanks to the fact that ROC and DAQ are programmable.

3.2.5.3 The tracker data format

A hit data packet has a fixed length of 256 bits (32 bytes). The packet structure is as follows:

- 16 bit header: the header contains information to uniquely specify this is a packet header, a channel identifier to specify the channel so the ROC can assign the hit to a wire number (straw index);
- 16 bit for the TDC left straw end;
- 16 bit for the TDC right straw end;
- The ToT (time-over-threshold) values for the two ends of the straw are each stored using 8 bits;
- The ADC samples require 12 bits each. For each hit, a fixed number of samples (15) is read out so we pack the samples tightly in order to fit the payload into two data per hit;
- 12 bits are set aside for preprocessing flags.

3.2.5.4 The Data Acquisition System

The Mu2e Data Acquisition system (DAQ) is based on a *streaming* readout architecture: all detector data are digitised, zero-suppressed in the front-end electronics, and then transferred from the detectors for further processing and storage. This architecture results in a large data throughput in the DAQ but offers a significant flexibility for data selection and analysis. Figure 3.15[84], shows the global architecture of the Mu2e DAQ. The Readout Controllers of all detectors are shown on the left side. The central box includes the main DAQ components: the Run Control Host, 40 DAQ servers, the Detector Control System (DCS), and the Event Building Switch. The Mu2e data rates will be substantial: for efficient data handling this segment of the DAQ heavily relies on firmware implementation. The right box includes the offline components, which perform data storage and offline processing. DAQ operation during data taking will be coordinated by the Run Control Host, which will manage a predefined Run Plan. During an active spill (approximately the first 43 ms in the 48 ms bunch extraction cycle described in Chapter 2), the experiment receives RF Zero-Crossing Markers from the Accelerator synchronised to the 1695 ns proton pulse cycles. This defines the Event Window¹. Based on these markers, the Command Fan-Out (CFO) module generates a 40 MHz system clock and embeds Event Window Markers (EWMs) into this clock to denote the start of event windows. The CFO distributes this encoded system clock and run control packets to the DTCs in the DAQ servers. DTCs then pass the encoded clock to the ROCs within the detectors, where the EWMs are recovered with a fixed latency relative to the initial RF Zero-Crossing Markers. The local ROCs use the EWMs to distinguish data collected during consecutive Event Windows. For the tracker, this involves generating a DDR3 memory address at the start of each Event Window to store tracker hits recorded during that period. Data readouts from the ROCs to the DAQ system are triggered by Data Requests. The Data Requests can be configured via the CFO as per the Run Plan and are typically issued to the tracker and calorimeter ROCs following each event window through the DTCs. Conversely, the CRV Data Requests are issued via software. The ROCs respond to these Data Requests by transmitting the corresponding data to the DTCs. Data from multiple DTC sources is routed through the Event Building Switch to a designated DTC, pre-processed, and then transmitted to the Online Processing module within the same DAQ server². If a trigger decision is made, the corresponding CRV data is also read out from the ROCs through software-generated Data Requests. Events which passed the trigger are sent to the Data Logger and transferred to the long-term storage. Calculations estimate a data rate of 35 GBps from the ROCs and an annual storage requirement for the entire experiment of approximately 7 PB.

In addition to the detector data, the DTCs also handle the acquisition of slow control variables from the detectors. For the tracker, these variables include pan-

¹Event Windows are also assigned during off-spill periods when no beam batches are received. These Event Windows are longer due to the reduced data rate.

²The DTCs are implemented as commercial PCIe cards located within the DAQ Servers.

els temperature, pressure, humidity, voltages, and currents at the specified PCB boards, high voltages, and channel threshold settings across all panels. These data are managed by the DCS Host and recorded in long-term storage. The data are also accessible in real-time in the Control Room for monitoring purposes. Other detector control tasks and operations are managed through DCS commands via the DTCs. For the tracker, this includes panel configuration, calibration using preamp-generated pulses, and disconnection of channel high voltage when necessary.

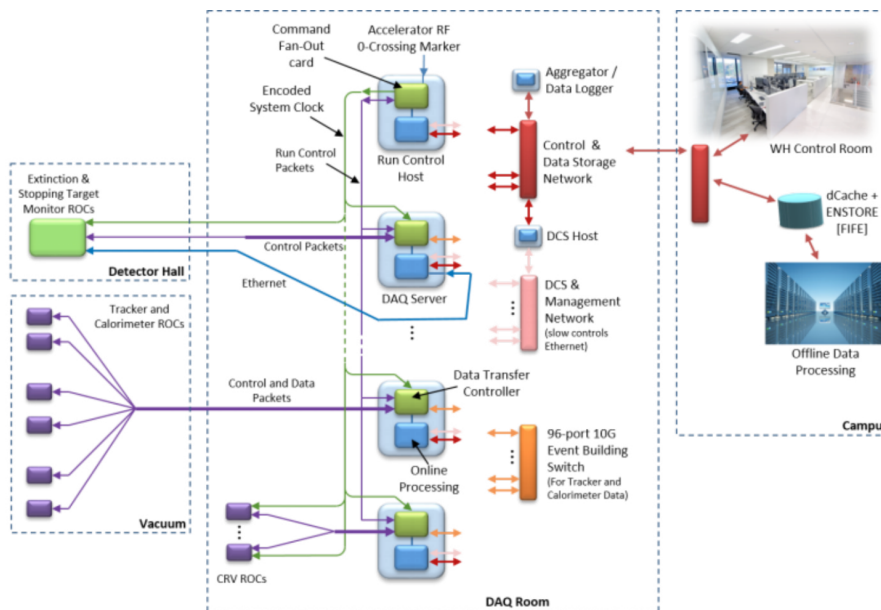


Figure 3.15: Mu2e Data Acquisition system architecture [84].

3.2.6 Tracker performance requirements

Here a summary of the requirements that the tracker must satisfy to ensure the success of the experiment is presented [85]. A momentum resolution better than $180 \text{ keV}/c$ for a $105 \text{ MeV}/c$ electron is needed in the nominal 1 T solenoidal field, as measured at the front face of the tracker volume (before passing through any tracker related material). Non-Gaussian tails, particularly any high-side tail, must be controlled such that the DIO background results in much less than one event at design sensitivity. To reach this, simulation results indicate that a single straw requires around 4 cm of longitudinal and $200 \mu\text{m}$ of transverse resolution for drift path lengths. It must have an acceptance of approximately 20% for conversion electrons.

The tracker must operate in an ambient vacuum ($< 10^{-4} \text{ Torr}$) and be able to withstand a rate of 5 MHz per straw (highest rate straw) 500 ns after the spill.

Chapter 4

Commissioning of the tracker DAQ and FEE

In this Chapter, I present the initial results of the tracker DAQ commissioning. Understanding the readout process is essential before reading out the detector. The first part focuses on the validation of the ROC readout through Monte Carlo simulation. The second one describes the initial data quality monitoring of the tracker preamps.

4.1 The test stand setup

Figure 4.1 and 4.2 respectively show a schematic representation and a picture of the tracker DAQ test stand installed at the IERC facility at Fermilab.

The test stand includes the entire readout chain of a 96 channels tracker panel: 48 channels connected to the digi-FPGA-1 and 48 channels to the digi-FPGA-2. Each channel corresponds to a tracker straw. At the test stand, the straws were not installed.

The DRAC board was connected via optical fiber to one DTC installed in DAQ computer (mu2edaq09). The DTC was programmed, for most of the tests, to emulate the CFO functionality to send the request to the ROC to perform the readout of one event. The event request was immediately followed by the event readout.

Depending on the test to perform, it was possible to select one between the two ROC operational modes:

- **ROC internal mode:** the ROC was emulating the data itself (user-defined) without reading the digi-FPGAs;
- **ROC external mode:** the ROC receives data from the digi-FPGAs. This is the mode to be used during the data-taking.

For most of the tests performed, the ROC was operated in the external mode: in some cases we simply used the digi-FPGAs internal pulser, in some other cases we injected calibration pulses directly in the CAL-side of the preamps at a frequency

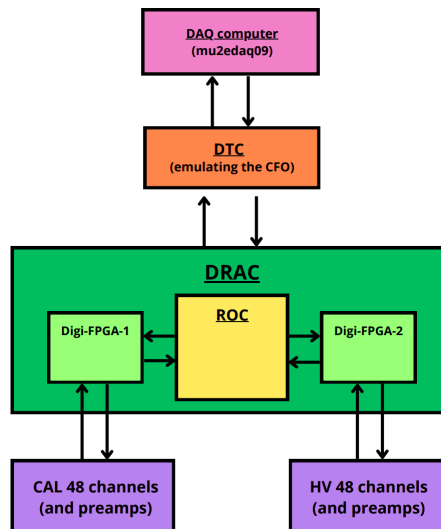


Figure 4.1: Block diagram representation of the tracker DAQ test stand. The two purple blocks represent the two channel batteries (48 CAL channels on the left, and 48 HV channels on the right), connected to their respective digi-FPGA (light green boxes), located in the DRAC board (dark green box). The digi-FPGAs are connected to the ROC (yellow box) that manages the communication with the DTC (orange box). The DTC is connected to the DAQ computer (pink box - mu2edaq09).




Figure 4.2: The tracker DAQ test stand (TS1) installed at IERC facility at Fermilab. The test stand includes the full readout chain of one entire tracker panel: 48 channels on the HV-side and 48 channels on the CAL-side, 1 DRAC board connected via optical fiber to one DTC installed in the DAQ computer (not shown).

we could customise. Figure 4.3 shows the readout timing diagram for one detector channel. There are two important parameters: the Event Window width (T_{EW}) shown in red, that represents the distance between the two consecutive proton pulses (called Event Window Markers-EWM), separated by $\sim 1.7 \mu\text{s}$ during Mu2e data taking, and the separation $T_{gen} = 1/f_{gen}$ between two consecutive hits (represented by grey triangles), where f_{gen} is the generator frequency. Each hit consists of two 16 bytes packets. In the process of testing, the T_{EW} has been varied between 700

ns to $50 \mu\text{s}$.

The ROC firmware has an internal hit buffer which stores up to 255 hits per event, adequate for the expected run-time occupancy. Depending on T_{gen} and T_{EW} , the data taking can proceed in two different modes:

- **regular mode:** the total number of hits within the Event Window is less than 255. In this case the ROC hit buffer doesn't get filled up and the total number of hits may vary from one event to another. It is expected that the number of hits will not exceed the threshold of 255, with an average of approximately 12 hits;
- **overflow mode:** the Event Window is large enough, so the total number of generated hits is greater than 255. In this case the ROC hit buffer always gets filled up, and only the first 255 hits are readout. The hits readout after the first 255 are lost.

Since the timing of the readout (i.e. the Event Window) is independent from the timing sequence of the pulse generator, the number of pulses contained in the Event Window can be different between two different windows. For example, Figure 4.3 shows that for two different relative timing offsets between the Event Window and the sequence of generated pulses, the number of pulses can be either three (Figure 4.3 Left) or four (Figure 4.3 Right).




Figure 4.3: Graphic illustration of pulses in an Event Window.

The relative timing offsets between different channels of the same digi-FPGA are of the order of few ns and can be measured. The pulse sequences from the two digi-FPGAs are offset relative to each other by a time interval Δt , which is constant for as long as the DRAC board is powered up and varies randomly between 0 and T_{gen} when the DRAC is powercycled.

The readout process is based on the sequential reading of channels. Depending on the Event Window and the start of the pulse sequence, the first channel in the readout sequence records as many hits as can fit within the Event Window (Figure 4.3). Once all hits for this channel are recorded, its readout phase concludes, and the system moves to the next channel in the sequence. This procedure is repeated for all channels, in a fixed order.

4.2 Validation of ROC readout and buffering

The first test performed was to verify the correct performance of ROC buffering. During the test, a single ROC was connected to a DTC. Data were collected with digi-FPGAs pulsed by their internal pulsers, and the ROC set in the external mode. The digi-FPGAs pulse generators operated at two specific frequencies, $31.29 \text{ MHz}/(2^7+1)$, or approximately 250 kHz, and $31.29 \text{ MHz}/(2^9+1)$, or approximately 60 kHz.

4.2.1 Development of the ROC bit-level simulation

ROC's readout logic can be emulated with a bit-level C++ simulation, which I contributed to develop. For each event, the simulated parameters are the number of hits in each channel and the total number of readout hits per event, which cannot exceed 255.

The main steps of the simulation are as follows:

- The start of the Event Window is set at $t = 0$;
- In each digi-FPGA, the timing of the first generated pulse is randomly sampled from a uniform distribution between 0 and T_{gen} ;
- Once the first pulse has been generated, the following pulses are added at the relative distance of T_{gen} , until the absolute time of the next pulse is above T_{EW} ;
- In the readout part of the simulation, the pulses are readout in the same order as in the firmware;
- the readout *continues* until all simulated hits have been *readout*, or the threshold of 255 hits has been reached.

The simulation allows to introduce the offset between digi-FPGA-1 and digi-FPGA-2 timing sequences and, internally to each digi-FPGA, the channel-to-channel offsets, as user-defined parameters. To fine-tune the simulation, we measured the channel-to-channel offsets from the data. For each digi-FPGA, we selected a reference channel and measured the offset of the remaining 47 channels relative to it. The reference channel is the first channel in the readout sequence: channel 91 for digi-FPGA-1 and channel 94 for digi-FPGA-2. Figure 4.4(4.5) shows an example of distribution of the time residuals between the pulses in a reference channel 91(94) and a randomly selected channel, 0(44) for digi-FPGA-1(digi-FPGA-2). The mean offset was determined with a Gaussian fit of the distribution. This procedure was repeated for all channels relative to each digi-FPGA. All measured offsets were of the order of a few ns and were included in the simulation. It was important to take into account all possible effects in the emulation of the ROC hit buffering to perform an accurate comparison between data and simulation. However, we were aware of the fact that given the order of magnitude of these offsets, the impact was negligible.




Figure 4.4: digi-FPGA-1: histogram of the delay between channel 0 and the reference channel 91, the first to be readout.




Figure 4.5: digi-FPGA-2: histogram of the delay between channel 44 and the reference channel 94, the first to be readout.

The comparison between the data and the simulation will be presented in the following sections, where we will refer to *occupancy* as the total number of hits recorded in a given channel during the test run.

4.2.2 The *overflow* readout mode

In the *overflow* mode (referred to as RUN281 in the following), the T_{EW} was set to $50 \mu\text{s}$ and f_{gen} to $\sim 60 \text{ kHz}$, that corresponds to the time distance between two consecutive pulses of approximately $16 \mu\text{s}$. This implies that the number of hits in a readout window could be either 3 or 4, depending on the offset between the start of the readout window and the generated pulses sequence.

4.2.2.1 Hit timing distribution and channels occupancy

We first checked the distribution of the hit time for each readout channel, and the distribution of the total number of hits (occupancy) in one specific readout channel

as a function of the channel number. Figure 4.6 shows the hit time distributions for two selected channels, channel 0 of digi-FPGA-1 and channel 2 of digi-FPGA-2. The top distribution is, as expected, uniform, however the bottom one shows some non-trivial features. The explanation for the two different distributions comes from

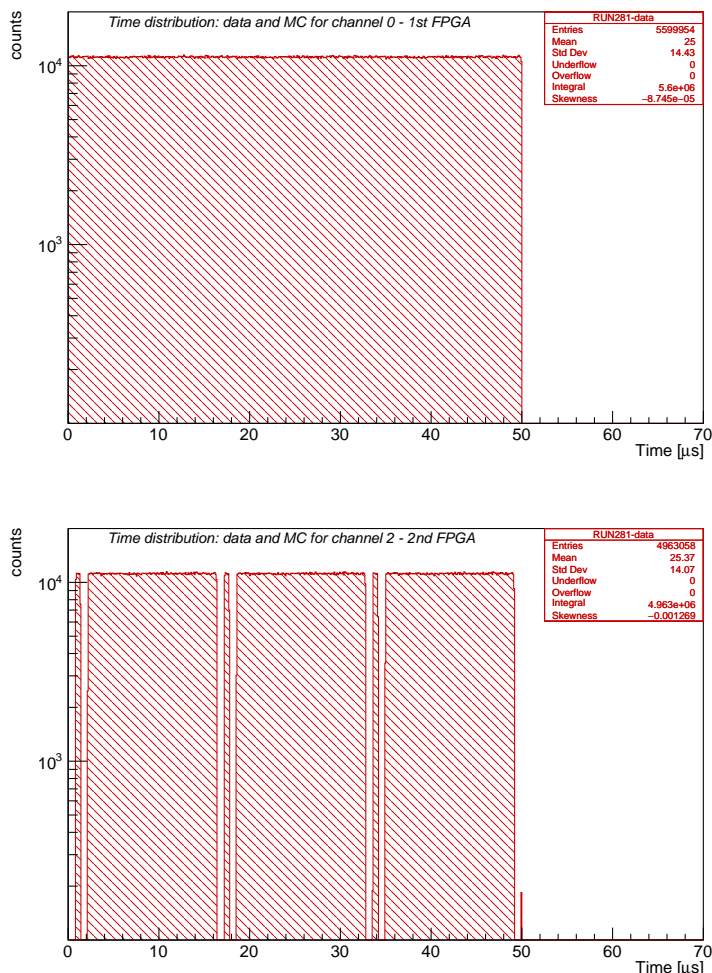


Figure 4.6: (Top): the time distribution of hits in the channel 0 in the digi-FPGA-1. (Bottom): the time distribution of hits in the channel 2 in the digi-FPGA-2.

Figure 4.7 (Top), which shows the channel occupancy as a function of the channel number (red is for data, blue for simulation).

The bin ordering in the histogram corresponds to the channel readout ordering. That means that the first bin on the left corresponds to the first channel and the last one on the right corresponds to the last channel in the readout sequence. In other words, channels on the left are at the beginning of the readout sequence and all hits from those channels are always fully readout. On the other hand, channels on the right are in the later part of the readout sequence and may not have all the hits readout, since the ROC buffer may get filled up.

Figure 4.8 shows the distribution of the number of hits in channel 0 of the digi-




Figure 4.7: (Top): number of hits versus channel number (data in red, Monte Carlo in blue). The two distributions are normalized to the same number of events. The channels are numbered in the readout order. Not all 96 channels are present in the histogram because the maximum threshold of 255 hits was reached with fewer channels. (Bottom): zoom on the last channels in the readout sequence. The data and MC distributions differ from each other by $\sim 10^{-3}$.

FPGA-1. As expected, the majority of events has 3 hits, with a tail of 4 hits. This provides an explanation of Figure 4.7 (Top). The first 69 bins of the histogram (Figure 4.7 (Top)) show a flat distribution with the maximum observed occupancy. Additionally, there are two more plateaus with slightly lower occupancy (less than 10% lower): one over the following 7 bins (with a small *dent* at the end), and the other spanning over the last 9 bins. This pattern is generated by the three following types of events added together.

- **channels 0-68**: the 48 channels of the digi-FPGA-1, the first one to be readout, are those with 4 hits per channel, which makes a total of 192 hits stored in the ROC buffer. In these conditions, 63 hits can still be stored. Because of the




Figure 4.8: The distribution of the number of hits in channel 0 of the digi-FPGA-1 (RUN281).

delay between the first and second digi-FPGA, only three hits can be readout by each channel of the digi-FPGA-2. The first 21 channels of the digi-FPGA-2 will be those with 3 hits per channel, resulting in a total of 255 hits. All hits sent to the following channels are lost.

- **channels 0-75:** the 48 channels of the digi-FPGA-1 are those with 3 hits per channel, which makes a total of 144 hits stored in the ROC buffer. The *dent* at the end of the second plateau is due to the fact that the digi-FPGA-2 contributes 111 hits and this number is not an integer of 4. The first 27 channels of the digi-FPGA-2 contribute 4 hits per channel each, but the three hits from channel 28 in the readout sequence fill up the total ROC buffer of 255 hits.
- **channels 0-85:** the 48 channels of the digi-FPGA-1 and the first 37 channels to be readout in the digi-FPGA-2 are those with 3 hits, 255 hits in total.

Not all 96 channels are present in the histogram because the maximum threshold of 255 hits was reached with fewer channels, in particular the number of 85 readout channels corresponds to the sum of 48 readout channels in the first digi-FPGA and 37 readout channel in the digi-FPGA-2.

Figure 4.7 (Bottom) zooms in on the rightmost channels of the distribution. The relative difference between the data and the MC distributions is at a level of 10^{-3} , which is a very good agreement.

Coming back to Figure 4.6, the first channels in the readout sequence always have all their hits readout, while the channels in the end of the readout sequence do not, as the ROC hit buffer gets filled up after the first 255 hits are readout. This results in a uniform time distribution for the first channels readout and in a non-uniform time distribution for the last readout channels, depending on T_{gen} and T_{EW} . The

dips in the hit timing distribution for channel 2 are defined by the timing offset between the two digi-FPGA pulsers.

4.2.2.2 Number of hits

Figure 4.9 shows that, in the *overflow* mode, the number of readout hits is exactly 255 for all events.




Figure 4.9: The distribution of the total number of hits readout per event (data in red, Monte Carlo in blue). The two distributions are normalized to the same number of events.

4.2.3 The *regular* readout mode

In the *regular* configuration, which will be referred to as RUN105038, the Event Window T_{EW} was chosen to be $25 \mu\text{s}$. The pulser rate f_{gen} of 60 kHz corresponds to $T_{gen} = 16 \mu\text{s}$.

4.2.3.1 Time distribution and occupancy

Figure 4.10 shows the distributions of the number of hits in two channels, one from the digi-FPGA-1 and another one, from the digi-FPGA-2 and Figure 4.11 shows the occupancy as a function of the channel number. In this readout configuration, the expected number of hits in a given channel within the Event Window is one or two, and the total number of pulses is always below 255. There is no overflow of the hit buffer and no pulses are lost.

In this mode, the readout of a given channel is not affected by the readout of previous channels and the channel *occupancy* distribution shown in Figure 4.11 is, as expected, uniform.

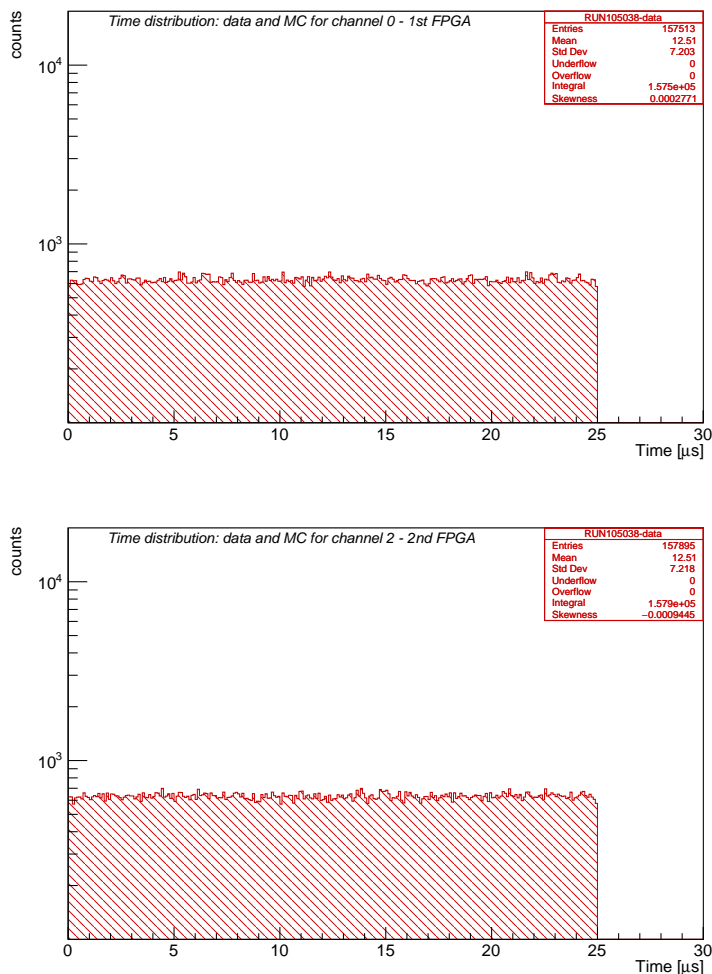


Figure 4.10: (Top): the hit time distribution for hits in channel 2, the digi-FPGA-2. (Bottom): the hit time distribution for hits in channel 0, the digi-FPGA-1.

Given the choice of f_{gen} and T_{EW} , the maximum number of hits in one Event Window is 2. Figure 4.12 shows the distribution of the number of hits in the channel 0 (FPGA-1).

4.2.3.2 Number of hits

Compared to RUN281, the Event Window in RUN105038 was twice shorter and the ROC readout buffer wasn't getting filled up.

Figure 4.13 shows the distribution of the total number of hits per event in the *regular* mode. There are two higher peaks, one at 144 and one at 192. The peak at 192 occurs when all 48 channels of digi-FPGA-1 and all 48 channels of digi-FPGA-2 have 2 hits each. The peak at 144 occurs when all channels of digi-FPGA-1 have 2 hits, and all channels of digi-FPGA-2 have 1 hit, or vice versa. The total number of hits within the event window depends on the relative offset of the Event Window




Figure 4.11: The number of hits versus the channel number for RUN105038 (data in red, Monte Carlo in blue). The two distributions are normalized to the same number of events. All 96 channels are readout.




Figure 4.12: The distribution of the number of hits in channel 0, digi-FPGA-1, for RUN105038. Entries in the $n(\text{hits})=0$ bin are due to the readout errors.

with respect to the digi-FPGA pulsers, and it varies from 144 to 192. As before, the relative difference between the data and the MC distributions is at a level of 10^{-3} , which indicates very good agreement.

4.3 Study of preamplifier performance

After having validated the correct ROC buffering, a set of different tests was accomplished to check the performance of the preamplifiers along with the readout chain. During the test described in this section, the same test stand as in Section 4.1 was




Figure 4.13: The distribution of the total number of hits per event in *regular* mode (data in red, Monte Carlo in blue). The two distributions are normalized to the same number of events.

used and one or two ROCs were connected to the same DTC, so one more test stand (TS2) was used. Preamps on the CAL side were added to perform the test. As mentioned in Section 3.2.5.1, the CAL side FPGA is programmed to manage the generation of the calibration pulses.

These pulses can be sent at an arbitrary rate and a frequency of 50 kHz was chosen. The data acquisition Event Window was set to 50 μ s. In this case, the number of pulses within one Event Window could be 2 or 3. To perform this test, I developed the initial step to the real-time monitoring and diagnostic tools: a set of real-time histograms to allow an easy check of the signal uniformity among channels within the same ROC or across multiple ROCs, and among different events. The diagnostics tools allowed the identification of errors, such as the presence of an invalid channel ID, more hits than the maximum allowed in a given channel and an undefined ROC link ID between ROCs.

4.3.1 Test 1: channel occupancy versus channel ID

The first test examined the channel occupancy. This was a crucial test, as it was important to determine if a channel was inactive, which would result in zero readout hits. It was also essential to detect the cross-talks (unwanted coupling between channels) which could cause the erroneous readout of hits in the incorrect channels. It was also possible to check if the involved channels showed more hits than expected.

The implemented preamplifier pulsing scheme allows to pulse every 8-th channel. Scanning all 96 channels of the panel, requires eight runs taken with the different first pulsed channel offset. Figure 4.14 shows the distribution of the occupancy for one channel sequence, which shows no anomalies. Figure 4.15 shows a run with one dead channel (ID=94) and a channel (ID=70) with more than three hits,




Figure 4.14: Regular distribution of number of hits versus channel ID. The 0th channel is the first to be pulsed. As a consequence, the active channels are the 8th, 16th, 24th, 32nd, 40th, 48th, 56th, 64th, 72nd, 80th and 88th. The number of hits is 3 or 4 for all channels. No cross-talks were observed in any neighbour channel.

which was not expected. In case of the dead channel, the problematic preamp was replaced. To investigate the second case, a distribution of the time difference Δt between consecutive pulses (Figure 4.16) has been made. Since the pulse frequency is 50 kHz, the Δt distribution should show one single peak at 20 μs . On the other hand, there are also two pronounced peaks at 16 μs and 4 μs . Analysis of the hit waveforms in the 4 μs peak showed an inversion of the waveforms: this has been further investigated and will be discussed in more detail in Section 4.3.2.2.

Figure 4.17 shows a clear example of a cross-talk between the two channels. We tried to characterize this phenomenon and it was noticed that it occurred only when the ID of the first channel to be pulsed is odd. The cross-talk was also asymmetric (the cross-talk from 3rd channel to 5th was observed, but no 3rd \rightarrow 1st). As described in Chapter 3 and shown in Figure 4.18, the preamps PBC boards are mounted on vertical boards. Each board contains two preamps, each of one connected to one of the two straw-tubes ends: odd numbers are associated to the preamp closer to the horizontal board, while even numbers are associated to the preamp farther from the horizontal board. For this reason, it is not surprising that the cross-talk appears to be between channels placed closer to the horizontal board. Moreover, we observed the cross-talk only in the first 20 channels IDs. This is due to the fact that the distance between consecutive preamp boards is slightly lower for the first channels (Figure 4.18). At the time of writing this Thesis, the possible solutions are still being investigated.




Figure 4.15: Distribution of number of hits versus channel ID. The 6th channel is the first to be pulsed. As a consequence, the channels that should be active are the 14th, 22nd, 30th, 38th, 46th, 54th, 62nd, 70th, 78th, 86th and 94th. The 94th channel was not responsive. In this case, the preamp was substituted. The number of hits is not the same for all channels. The 14th, 22nd, 38th, 46th, 54th, 62nd, 70th channels have more hit than expected. To address this issue, the distribution of Δt between hits of one of this channels is reported in Figure 4.16. No cross-talks were observed in any neighbour channel.




Figure 4.16: Δt distribution between hits in the 70th channel. The Δt distribution between hits should peak at 20 μs , as the pulser operates at a frequency of 50 kHz. However, peaks at approximately 16 μs and 4 μs are observed. All waveforms at around 4 μs are inverted compared to the regular ones. The waveforms are shown in Section 4.3.2, and the reason for this behaviour will be explained in the same section.




Figure 4.17: Distribution of the number of hits versus channel ID. The 7th channel is the first to be pulsed. As a consequence, the only channels that should be active are the 15th, 23rd, 31st, 39th, 47th, 55th, 63rd, 71st, 79th, 87th and 95th. Channels 9th, 17th and 25th were also observed to be active, indicating cross-talks.




Figure 4.18: The zoomed view of the Analog Mother Board with the preamp PCB boards installed. The distance between two consecutive boards decreases from left to right. The first preamps in the occupancy plot are the ones on the right.

4.3.2 Test 2: analysis of the readout pulses waveforms

This section describes the study of the reconstructed waveforms from the charge injection. The sampling frequency of the ADC is 40 MHz, which corresponds to the waveform bin width of 25 ns (Section 3.2.5.2).

Figure 4.19 shows an example of a reconstructed waveform (RUN105421). The reconstructed waveforms shown below have the baseline subtracted. In Section 4.3.2.1, I will provide a deeper analysis of the estimation of the waveform baseline. In first approximation, a good waveform is characterized by having a flat distribution in the first 10 samples, indicating the absence of noise at the start. It should feature a high positive charge peak with a sharp leading edge, which is the rising slope of the waveform. Additionally, the waveform should have a negative tail. The tail is generated by a differentiating circuit that shapes the signal using a high-pass filter.

This circuit was adopted because at high count rates two pulses can overlap: a new pulse can arrive before the previous one has returned to zero, thus leading to an overlap with the initial pulse's undershoot. This overlap may reduce the apparent amplitude of the subsequent pulse and generated the undesired broadening of peaks in the energy spectrum. The filter does not affect the leading edge of the pulse because the time constant of the differentiating circuit is large compared to the rise time.




Figure 4.19: Regular waveform (50th channel).

4.3.2.1 Estimate of the waveform baseline

A straightforward procedure was adopted to estimate the baseline: we simply take the mean value of the first 10 samples of the waveform. In each channel, the stability of the baseline is an indicator of the level of noise within the electronic chain. The baseline of each channel varies according to the tolerances of the electrical components on the board. Figures 4.20 and 4.21 show the baseline distributions for two different channels (RUN105421). In both cases, the baseline is close to 210 ADC counts, with a $\text{FWHM} = 2\sqrt{2\ln 2}\sigma \sim 4.5$ ADC counts ($\sigma \sim 1.9$ ADC counts).

In some cases, a lower baseline value was observed. We also observed dips of specific depths, for example 64, 128, or 192, which may suggest the possibility of malfunctioning 6th and 7th bit of the ADC (Figure 4.22). To address this issue and to correctly reconstruct the waveform, the problematic samples have been identified and excluded from the baseline estimate.

4.3.2.2 Inverted waveforms

As discussed in Section 4.3.1, a larger number of hits than 4 in channel 70 was observed. That led to a different from expected Δt distribution between the hits. While the Δt peak was expected at $20 \mu\text{s}$, in some channels two additional peaks




Figure 4.20: The fitted baseline distribution of channel 0.




Figure 4.21: The fitted baseline distribution of channel 8.

were identified: one at $16 \mu\text{s}$ and another at $4 \mu\text{s}$. The $4 \mu\text{s}$ peak is characterized by inverted waveforms (Figure 4.23). The $16 \mu\text{s}$ peak is characterized by regular waveforms, as the $20 \mu\text{s}$ peak. This phenomenon arises from the triggering on the trailing edge of $4 \mu\text{s}$ long input pulse. In fact, the sum of $4 \mu\text{s}$ and $16 \mu\text{s}$ is the inverse of the pulser frequency. The chosen solution was to replace the preamp.

4.3.2.3 The waveform charge and pulse height

To determine the integral of the positive and negative parts of the waveform, the pulse height and the first sample, a constant threshold discriminator, which we set at 5 ADC counts, was used. The first sample is defined as the first one to be above the threshold, and the positive charge is the integral of the region above the threshold. The negative charge is the integral of the waveform's negative tail. The pulse height




Figure 4.22: Waveform of the 58th channel. Before subtracting the baseline from the ADC counts, the dips were identified. The depth of the dips is 128.




Figure 4.23: Inverted waveform of the channel 50.

is the highest ADC count reached by the waveform. The distributions of charge, pulse height, first sample and negative charge are shown in Figures 4.24, 4.25, 4.26, and 4.27.

We first analysed the distributions of the positive charge and pulse height. The charge and pulse height distribution exhibit a non-trivial shape. Three and two peaks were observed in both distributions; it was verified that these peaks were not correlated with the number of hits. Moreover, a correlation between the positive charge and pulse height peaks has been observed (Figure 4.28). To understand the origin of this pattern, we computed s_{mean} , that is the waveform mean sample weighted with the charge:

$$s_{mean} = \frac{\sum_i \text{sample}_i \cdot q_i}{\sum_i q_i} \quad (4.1)$$




Figure 4.24: The (positive) charge distribution of the waveforms (channel 66).




Figure 4.25: The pulse height distribution of the waveforms (channel 66).

Figure 4.29 shows the s_{mean} distribution for channel 66. This distribution shows that the time difference Δs between consecutive pulses (approximately 2 ns), is well below the bin width, which is 25 ns. The next step was to check if s_{mean} could be correlated with positive charge and pulse height peaks, as demonstrated in Figure 4.30 and Figure 4.31. The pulse height and charge peaks are perfectly correlated with the s_{mean} peaks. In particular, few ns of delay between waveforms, correlated with the ADC clock, result in three or two peaks in the positive charge and pulse height distributions. These peaks are merely artifacts of the pulser timing shifted with respect to the ADC clock and are not problematic. This behaviour was also confirmed with a simple simulation. The simulation generates two series of triangular pulses, each delayed with respect to the other and models the behaviour of the corresponding ADC outputs. The ADC is simulated by partitioning signals into bins and calculating the signal mean within each bin. The delay between the




Figure 4.26: The first sample distribution of the waveforms (channel 66).




Figure 4.27: The negative charge distribution of the waveforms (channel 66).

two series is lower than the width of the ADC bins. The output of the simulated ADC consists of square waves for each series of triangular pulses. The integrals and maximum values of the square waves over defined intervals are then computed. The result is shown in Figure 4.32. The simulation is not intended to replicate reality in its entirety, but rather serves as a tool for understanding the correlation between the ADC clock and pulser timing. While the simulated positive charge and pulse height distributions effectively exhibit peaks, they do not precisely match the number of peaks seen in Figure 4.24 and 4.25. Since the ADC begins integration at the same time for both pulse series, any temporal shift between waves results in lower or greater waveform values on the y -axis at identical times. This discrepancy translates into distinct ADC outputs and consequently affects the observed pulse height. The same principle applies to the charge. Moreover, the pulse height peaks exhibit clear distinctions with respect to the charge ones, a characteristic also evident




Figure 4.28: 2D distribution of pulse height versus (positive) charge.






Figure 4.29: The distribution of the waveform mean sample weighted with the (positive) charge (channel 66).

in Figures 4.24 and 4.25. Lower values of pulse height or positive charge correspond to inverted waveforms, as mentioned in Sections 4.3.1 and 4.3.2, or to glitches. In the case of inverted waveforms, the positive charge value was ~ 100 and the pulse height ~ 20 , while for glitches, $\lesssim 70$ and ~ 35 . An example of a glitch is shown in Figure 4.33. The origin of the glitches is not yet understood. Special channels with excessive noise were identified through the charge distribution. Figure 4.34 shows an example of a noisy waveform. Since the threshold mentioned in Section 4.3.2.3 is 5 ADC counts, charge values for these channels were consistently lower than expected (approximately 10). This is due to the fact that only the first sample above 5 was considered as the entire waveform. To address the issue, the channel voltage threshold was either adjusted or the preamp was replaced.

Figure 4.30: 2D distribution of (positive) charge versus s_{mean} .Figure 4.31: 2D distribution of pulse height versus s_{mean} .

4.3.2.4 Channels response uniformity

From the perspective of data taking, it is important to find an easy way to identify as soon as possible problematic channels and check the response uniformity across the channels. For this reason, plots of positive charge, first sample, pulse height and baseline means versus channel ID were produced. These plots are shown in Figures 4.35. The outliers in the plots are due to crosstalk, noisy channels, inverted waveforms and glitches. These histograms are the first examples that will be used during the online Data Quality Monitoring.

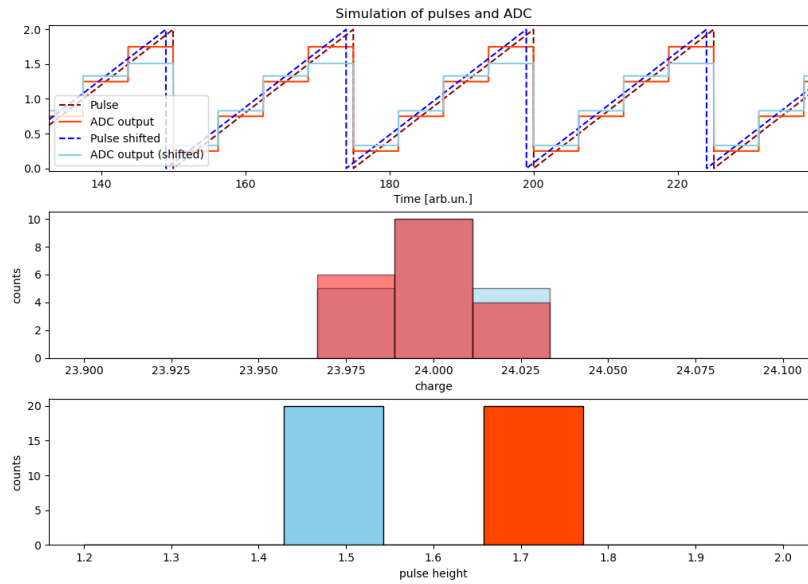


Figure 4.32: Simulation of the charge and pulse height distribution behaviour. In the upper plot, triangular pulses representing charge injection waveforms are plotted alongside the square waves from the simulated ADC. The red wave corresponds to the dark red pulses, while the light blue wave corresponds to the dark blue pulses, shifted from the red one. In the central plot, the distribution of charge is displayed, and in the bottom plot, the distribution of pulse height is shown. The light blue distribution corresponds to the light blue ADC, and the same goes for the red one.




Figure 4.33: Glitch waveform of channel 0.




Figure 4.34: Noisy waveform of the 36th channel.

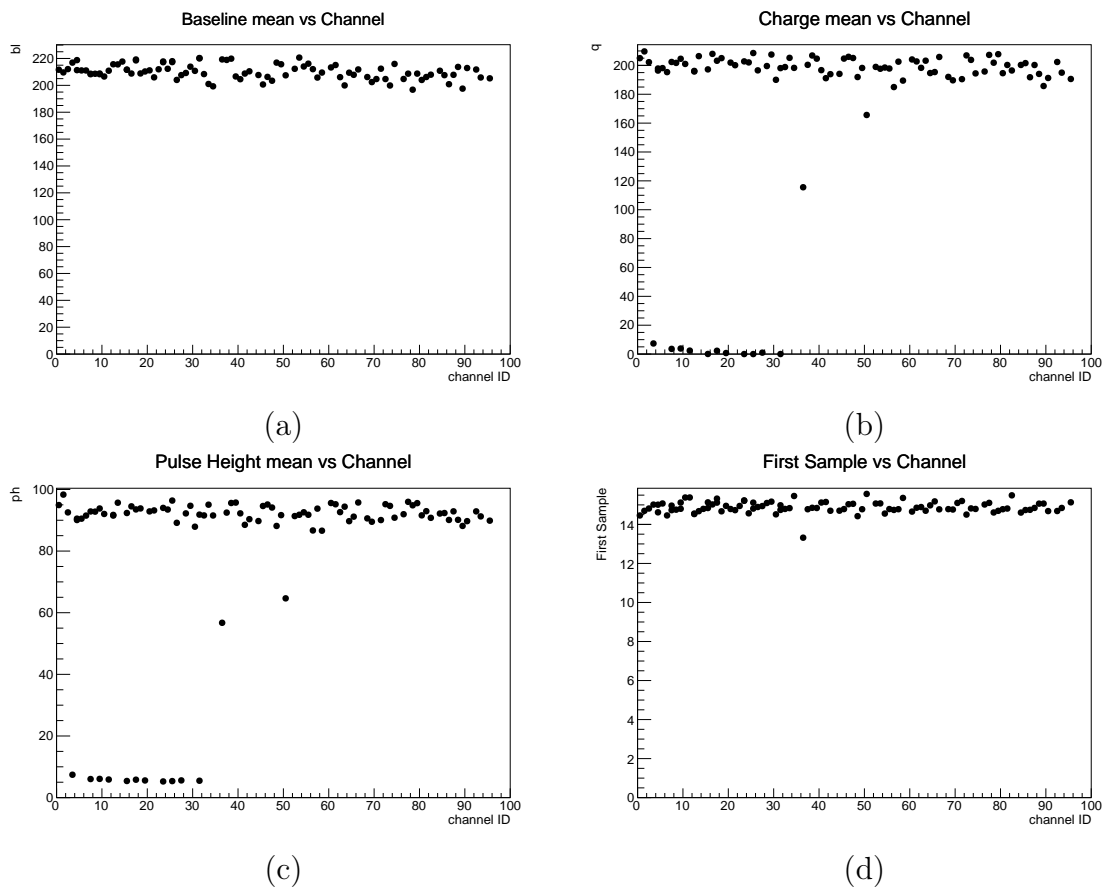


Figure 4.35: Baseline (a), charge (b), pulse height (c), first sample (d) mean value versus channel ID.

Chapter 5

First steps towards the station calibration

Following the comprehensive testing of the DAQ and FEE, the subsequent crucial step towards achieving a fully operational tracker is the calibration phase. This calibration will be carried out exploiting cosmic muons. The primary objective of this calibration process is to accurately determine signal propagation times and channel-to-channel delays within each straw. To accomplish this, it is essential to perform an unbiased reconstruction of the longitudinal position of the hits within the straws. This would be easily achieved using the station orientated horizontally. However, certain technical and mechanical contingencies prevent horizontal calibration, necessitating the exclusive focus on vertical orientation. This Chapter focuses on reconstructing the trajectories of cosmic muons within a vertically oriented station, as well as analyzing potential biases and systematic errors that could arise from this particular orientation.

5.1 Overview of the timing calibration

When a charged particle crosses the straw tube gas volume, the resulting ionization charge generates an electric signal along the anode wire which propagates towards the two ends of the straw and to the front-end electronics, where TDCs measure the signal arrival times t_1 and t_2 . The primary goal of the timing calibration is to enable the determination of the track position along the straw from the measurement of the arrival times t_1 and t_2 . The dependence of t_1 and t_2 on the track position x_{track} is given by the following equations:

$$\begin{aligned} t_1 &= t_0 + \frac{x_{\text{track}}}{v} + t_d + d_1 \\ t_2 &= t_0 + \frac{L - x_{\text{track}}}{v} + t_d + d_2 \end{aligned} \tag{5.1}$$

where t_0 is the particle's crossing time, t_d is the drift time in the straw, L is the length of the straw, x_{track} is the reconstructed track position on the wire, v is the

propagation velocity of the electric signal along the anode wire and d_i are the delays introduced by the front end electronics. t_1 and t_2 are measured by the TDCs on the front-end HV side and on the CAL side. The measurement of the time difference $\Delta t_{12} = t_1 - t_2$ allows determining the coordinate x_{track} , while $(t_1 + t_2)/2$ allows to measure the drift time up to an offset common to all channels. If we subtract and add the two equations:

$$\begin{aligned}\Delta t_{12} &= \frac{2x_{\text{track}} - L}{v} + (d_1 - d_2) \\ \frac{(t_1 + t_2)}{2} &= t_d + t_0 + \frac{d_1 + d_2}{2} - \frac{L}{2v}\end{aligned}\tag{5.2}$$

The first equation contains the time difference Δt_{12} between the signals at the two straw ends as measured by the TDCs, the reconstructed track coordinate x_{track} along the wire, the propagation velocity v of the signal along the wire and the difference $d_1 - d_2$ between the delays introduced by the front-end electronics. Performing the timing calibration means determining v and $d_1 - d_2$ from the Δt_{12} measured by the TDCs and from an independent and unbiased reconstruction of the track projected on the wire (x_{track}). This is the first calibration of the detector and will be performed using cosmic muons. At this level, the only available spatial information is whether a straw has been crossed by a muon or not. The only way to obtain x_{track} from a "yes or no" information is to rely on the tracker station geometry (Section 5.4).

5.1.1 Operational constraints of the station

The detector, its infrastructures and services, and, consequently, the stations have been designed to operate in a *vertical orientation*. This has now an impact on several aspects of detection operation: from the storage and handling tools of the fully assembled and equipped station, which is a very fragile system, to the available services, for example the gas distribution system, which have been designed for managing, moving and operating a complex detector with a specific orientation, and may not allow to move and operate even one single station, in a different orientation. The most reasonable orientation of the station for a test with cosmic muons is the *horizontal orientation*. That facilitates unbiased reconstruction, since cosmic muons are distributed according to $\cos^2\theta$ and they are predominantly vertical (Section 5.2.1), thus crossing the straws mostly perpendicularly. Initially, the operational constraints, which may prevent the station from being operated horizontally, were assessed. Apparently, the main challenge with the horizontal orientation is the operation of the gas distribution system. The gas distribution system of the panel is equipped with two check valves: the first one is located on the supply side, the second one on the exhaust side. These valves operate using a simple mechanism where a ball is pressed onto an O-ring to form a seal. The only operational mode for these valves is to be oriented vertically to ensure proper sealing. In our specific case, this is not an issue for the check valves on the supply side, since they should be kept open during calibration to allow gas flowing. On the other hand, the check

valves on the exhaust side of the panel must be closed at all times to prevent gas from leaking before reaching the panel. Figure 5.1 shows a schematic representation of the operational mode of the supply and exhaust valves. The red arrows represent the direction of the gas flow. On the supply side (Figure 5.1 (Left)), the ball has to rest on the O-ring, or the gas flows downwards towards the atmosphere, rather than flowing upwards towards the panel. According

to these specifications, it seems that the only operational mode of the station is in the vertical position. Additionally, other reasons include space constraints, as placing a station horizontally takes significantly more space than vertically. Moreover, the station is quite fragile, and excessive movement could lead to potential issues.

The ultimate goal is to perform a timing calibration of the first assembled station of the tracker using cosmic muons, with the aim of achieving a longitudinal hit position resolution better than 4 cm. This would enable a better distinction between different particle trajectories while reconstructing helices (Appendix C).

In the following section, we will describe the simulation performed to reconstruct cosmic tracks using a vertically oriented station, with the aim of understanding potential biases in determining the longitudinal position.




Figure 5.1: Schematic view of the valves located in a tracker panel.

To ensure optimal performance, it is essential to take data with a vertical oriented station to achieve conditions that are as close as possible to those of the real detector. However, if the vertical orientation proves ineffective for calibration, alternative mechanical solutions will be considered.

This Chapter reports on a simulation study which shows that the vertical orientation introduces a significant bias on the x_{track} which significantly complicates the station calibration.

5.2 Cosmic muons as a calibration source

Cosmic muons will play a fundamental role in the calibration of the entire Mu2e detector system and, in particular, the tracker. They have unique characteristics that make calibration with cosmic rays complementary to other techniques:

- calibration samples with cosmic muons can be taken during standard detector operations, with the same detector conditions as the physics samples;
- the cosmic muons flux ($\sim 1 \text{ cm}^{-2}\text{min}^{-1}$ for horizontal detectors and a mean muon energy of $\sim 4 \text{ GeV}$ [86]) is sufficient to allow collecting a substantial amount of calibration data in a relatively short time and thus monitoring continuously the detector response;
- a cosmic muon is a minimum ionizing particle (MIP) and its energy loss is almost independent of its energy;
- the speed of the cosmic muons is equal to the speed of light c , and the time they take to traverse a detector, in our case one tracker station, can be used to align the time offsets of all the channels without any external time reference.

5.2.1 Cosmic muons energy and angular distribution

The muon flux at sea level is usually described by the Gaisser formula [87]:

$$\frac{dI}{dE_\mu d\Omega dt dS} = \frac{0.14}{\text{cm}^2 \text{ s sr}} \left(\frac{E_\mu}{\text{GeV}} \right)^{-2.7} \left[\frac{1}{1 + \frac{1.1E_\mu \cos \theta}{115\text{GeV}}} + \frac{0.054}{1 + \frac{1.1E_\mu \cos \theta}{850\text{GeV}}} \right] \quad (5.3)$$

where E_μ is the muon energy and θ is the muon polar angle. The two terms in brackets correspond to the contribution of the charged pions and kaons, while the small contribution from charm and heavier flavors is neglected. This simplified formula doesn't take into account muon decays and the curvature of the Earth, thus it is only valid for zenith angles $\theta < 70^\circ$ and for energies $E > \frac{100}{\cos \theta} \text{ GeV}$.

A modified version of the standard Gaisser formula, called Gaisser-Tang model, is used to account for low energy and large zenith angle effects [87]:

$$\frac{dI}{dE_\mu d\Omega dt dS} = \frac{0.14}{\text{cm}^2 \text{ s sr}} \left(\frac{E_\mu}{\text{GeV}} \left(1 + \frac{3.64\text{GeV}}{E_\mu (\cos \theta^*)^{1.29}} \right) \right)^{-2.7} \left[\frac{1}{1 + \frac{1.1E_\mu \cos \theta^*}{115\text{GeV}}} + \frac{0.054}{1 + \frac{1.1E_\mu \cos \theta^*}{850\text{GeV}}} \right] \quad (5.4)$$

where the second term in the bracket is the same as in the standard formula, except that the zenith angle θ is substituted by the angle θ^* . The relation between $\cos \theta$ and $\cos \theta^*$ is given by:

$$\cos \theta^* = \sqrt{\frac{(\cos \theta)^2 + P_1^2 + P_2(\cos \theta)^{P_3} + P_4(\cos \theta)^{P_5}}{1 + P_1^2 + P_2 + P_4}} \quad (5.5)$$

The parameters $P_1 = 0.102573$, $P_2 = -0.068287$, $P_3 = 0.958633$, $P_4 = 0.0407253$, and $P_5 = 0.817285$ were calculated using a dedicated simulation of muon production in the atmosphere. A representation of the angles is shown in Figure 5.2. Another factor is included in Equation 5.4 to account for the possibility of muon decays, which are more significant at low energies. The numerical constants are derived by fitting experimental data from various cosmic muon studies.




Figure 5.2: The relation of the observed zenith angle of muons, θ^* , to the zenith angle at the muon production point in the atmosphere, θ . R is the radius of the Earth [87].

5.2.2 Monte Carlo cosmic generation with CRY

Several Monte Carlo programs allow to simulate sea-level cosmic ray muons: the most widely used is CRY, developed at Los Alamos National Laboratory (LANL). CRY [88] functions as a generator for air showers induced by primary cosmic rays. The CRY package uses precomputed input tables derived from comprehensive MC-NPX 2.5.0 simulation of protons in the energy range between 1 GeV and 100 TeV, at the top of the atmosphere. The generation of muons and other secondary particles is governed by the pion and kaon decays. The package generates the cosmic muon flux within a zenith angle range of $0-90^\circ$, following a $\cos^2\theta$ distribution, and an energy range between 1-100 GeV, following the Gaisser-Tang parameterization. The CRY package accounts for the dependence of cosmic muon flux on various parameters, including altitude, latitude, and solar activity. CRY provides muon flux data at three different altitudes: sea level, 2100 m, and 11300 m, with sea level being selected for this study. The latitude is set at 41.8° N, corresponding to Fermilab. To avoid variations of the primary cosmic muon flux due to solar activity, a common day (6-21-2021) was chosen outside the maximum and minimum sunspot cycle for the simulation. While there are other packages designed to produce cosmic muons with specific energy and zenith angles, such as CORSIKA [89], CRY was chosen for its straightforward setup in Mu2e Offline.

5.2.3 Monte Carlo sample and coordinate system

A sample of cosmic muons was generated with a uniform distribution within a horizontal plane approximately 11 m above the muon beam axis. The GEANT4 simulation also incorporates effects of the external neutron shield surrounding the DS and the concrete ceiling of the Mu2e experimental hall. Starting from the generation

plane, cosmic muons pass through the concrete ceiling and walls situated outside the experimental hall. As they propagate downwards to the detector, they interact with the building materials. The external neutron shield surrounding the DS consists of concrete ($\rho=2.3 \text{ g/cm}^3$) of roughly 0.9 m thickness, while the ceiling above the Mu2e experimental area comprises 1.8 m of concrete. The entire structure is enclosed within a volume of air at standard temperature and pressure.

The simulation of the station was performed in what in Mu2e jargon is called "extracted position", which means the station was placed outside the solenoid. Given that the Earth's magnetic field is approximately $4 \div 5 \times 10^{-5} \text{ T}$, and considering that the dimensions of the station are on the order of 1-2 m, along with the fact that the majority of particles under consideration have energies around $\sim 1 \text{ GeV}$, the curvature radius of the particles is approximately $R \sim p[\text{GeV}]/(B[\text{T}] \cdot 0.3) \sim 70 \text{ km}$. This radius is significantly larger than the station's dimensions, allowing us to reconstruct particle trajectories as straight lines. For this reason, in the simulation the magnetic field was set to zero.

Before describing the event selection and the reconstruction, it is important to define the coordinate system. As shown in Figure 5.3, the z axis is parallel to the DS axis, while the y axis is directed vertically.




Figure 5.3: Schematic view of a tracker station and definition of the coordinate system used for the analysis. Different shades of gray indicate different numbers of overlapping panel.

5.3 Monte Carlo event selection

Since in the Mu2e simulation the tracker is made of 18 stations, to perform this study, we selected those events where a cosmic muon crosses the first station.

Although the station has four layers, the panel geometry does not provide a uniform coverage: there are regions of the detector where a particle with a trajectory orthogonal to the panel crosses four, three, two or even only one detector layer (Figure 5.3). Given the rotation angles of the panels within the same face and between different faces, if a particle crosses four layers, the straws of the four panels are not parallel to each other. Each straw can be parameterized as a straight line in a plane $z = z_i$, where z_i is the z coordinate of the panel which contains the straw:

$$(D_{x,i}t + M_{x,i}, D_{y,i}t + M_{y,i}, z_i) \quad (5.6)$$

where $D_{x,i}$, $D_{y,i}$ are the x and y i -th straws' direction cosines and sines respectively, while $M_{x,i}$, $M_{y,i}$ are the x , y straws' midpoints.

To perform this study, it is necessary to reconstruct the 3D trajectory of the cosmic muon, which requires at least two 3D points along the muon's trajectory. The x and y coordinates of the each required 3D point can be determined from the intersection of one pair of straws. Since all straws within a single panel are parallel to one another, it is necessary to use at least two straws from different panels. The intersection of two straws from different panels on the same face lies outside the tracker circle (Figure 5.3), which is not physically possible. Therefore, at least one hit per face is required.

The z coordinate of the 3D intersection is determined as the average of the straws' z coordinates. Thus, to avoid selecting parallel straws in the same plane, the first selection cut was to require at least one hit per face. This first requirement will be denoted as 4/4 requirement in the following sections.

Since one particle can hit more than one straw within a panel, another selection criterion was to choose events with fewer than three hit straws per panel. This criterion was selected to minimize errors in track reconstruction.

5.3.1 Panel illumination pattern

To achieve precise calibration, the hits should be uniformly spread over the entire surface of the panel. On the other hand, this is not what is obtained, due to the station geometry and the selection cuts. Figure 5.4 shows the resulting hit distribution for panel 0 of plane 0, which corresponds to the panel in the upper right part of Figure 5.3, after applying the selection cuts outlined in Section 5.3. This panel was picked as one possible example, the distributions for all panels in both planes are similar, as shown in Appendix D.

Figure 5.4 shows a non-uniform, or spotty, panel illumination. This is due to the 4/4 requirement: given the station geometry, the 4/4 overlap regions are limited to the edges of the panels (darkest areas in Figure 5.3).




Figure 5.4: The Monte Carlo illumination across the panel (panel 0, plane 0). The selection requires cosmic muons which satisfy the 4/4 requirement and have less than three hits per panel.

Moreover, the 4/4 requirement selects specific muon directions in the y - z plane. Figure 5.5 shows the $m_{yz} = \Delta y / \Delta z$ distribution for the selected muons. As expected, there are no particles corresponding to the $m_{yz} \sim 0$, which corresponds to the horizontal direction and $m_{yz} \rightarrow \infty$ (vertical), since the 4/4 requirement is not satisfied by vertical muons. The selected muon directions tend to have $|m_{yz}| \sim 1$, which correspond to an angle of approximately 45° .

Figure 5.4 shows that there are almost no hits in the central region of the panel. This could lead to waveform shape variations, resulting in potential time walk. The signals generated by particles at one end of a straw must propagate to the other end as well. The leading edge decreases in amplitude and broadens, while moving to the other end, which may introduce timing systematics.

5.4 Reconstructing muon tracks

To determine x_{track} in Equation 5.2, it is necessary to reconstruct the 3D muon track starting from the *StereoHit* coordinates. As mentioned relatively to Equation 5.6, what is known about each single straw is:

- the straw direction $(D_{x,i}, D_{y,i})$;
- the straw midpoint $(M_{x,i}, M_{y,i})$;
- the straw z_i coordinate.

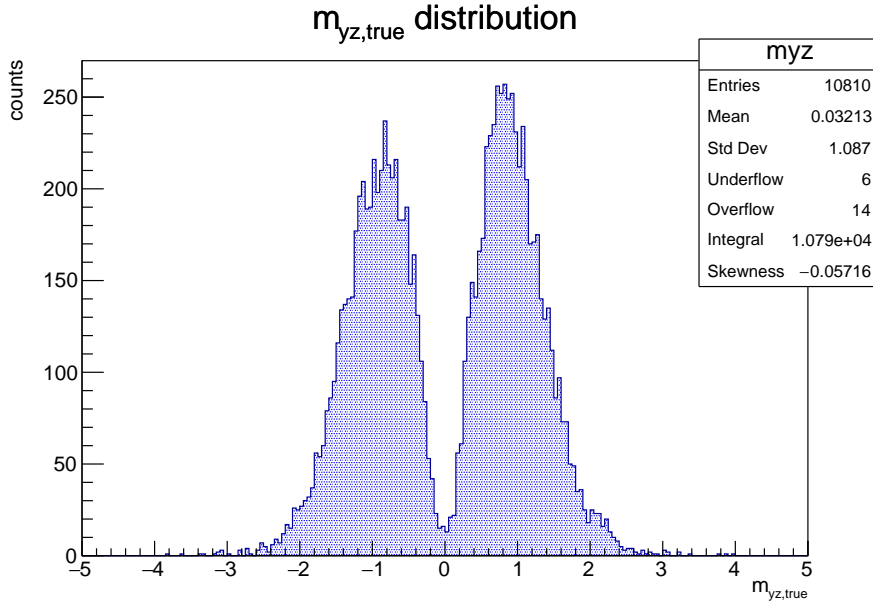


Figure 5.5: $m_{yz} = \Delta y / \Delta z$ distribution of muons satisfying 4/4 requirement.

Track reconstruction begins with the sole information of whether a straw has been hit or not. To improve the hit spatial resolution, adjacent straws simultaneously hit, which have been most likely crossed by the same muon, are combined and form the object called *ClusterHit*. A maximum of three straws are used to reconstruct a *ClusterHit*, to minimize the uncertainty on the determination of the 3D intersection points. This new object represents a straw with the same direction and the midpoint computed as the average of the midpoints of the straws which make it. Since muons with hits in four panels in four different faces are selected, four *ClusterHits* are reconstructed within a station. The new *ClusterHit* midpoints will be referred to as $(x_{m,i}, y_{m,i}, z_{m,i})$. Two *ClusterHits* in the two faces of the same plane are combined to determine one *StereoHit* on the $x - y$ plane. The x and y coordinates of the *StereoHit* are determined by the intersection point of the two *ClusterHits*. The z coordinate of the *StereoHit* is the average of the z coordinates of the *ClusterHits*. This is reproduced in Figure 5.6. Considering only one plane and the following equations as the $x - y$ *ClusterHits* equations:

$$\begin{aligned}
 y &= \frac{v_1}{u_1}(x - x_{m,1}) + y_{m,1} \\
 y &= \frac{v_2}{u_2}(x - x_{m,2}) + y_{m,2}
 \end{aligned}
 \tag{5.7}$$

where u_i is the *ClusterHit* x direction and v_i is the *ClusterHit* y direction, the




Figure 5.6: Schematic view of a cosmic muon (light blue) hitting the station (vertical orientation). The red lines are the *ClusterHits*, the dark blue dots are the *StrawHits* and the green dots are the *StereoHits*.

StereoHit coordinates will be:

$$\begin{aligned}
 x &= \frac{u_1 u_2 (y_{m,1} - y_{m,2} + \frac{v_2}{u_2} x_{m,2} - \frac{v_1}{u_1} x_{m,1})}{v_2 u_1 - v_1 u_2} \\
 y &= \frac{v_2}{u_2} \left(\frac{u_1 u_2 (y_{m,1} - y_{m,2} + \frac{v_2}{u_2} x_{m,2} - \frac{v_1}{u_1} x_{m,1})}{v_2 u_1 - v_1 u_2} - x_{m,2} \right) + y_{m,2}
 \end{aligned} \tag{5.8}$$

In two different planes, there are two *StereoHits* available, allowing for the reconstruction of the track on $x - y$ and $y - z$ planes. Referring to these *StereoHits* as (x_1, y_1, z_1) and (x_2, y_2, z_2) respectively, the final track slope on $x - y$ plane will be denoted as m_{xy} , while on the $y - z$ plane it will be denoted as m_{yz} . The corresponding y -intercepts will be q_{xy} and q_{yz} . The formulas for computing these parameters are as follows:

$$m_{xy} = \frac{y_2 - y_1}{x_2 - x_1} \tag{5.9}$$

$$q_{xy} = -m_{xy} \cdot x_1 + y_1$$

$$m_{yz} = \frac{y_2 - y_1}{z_2 - z_1} \tag{5.10}$$

$$q_{yz} = -m_{yz} \cdot z_1 + y_1$$

To determine the four hit positions of the reconstructed track on the panels, it is necessary to intersect the track with the panels' z_i coordinates. The reconstructed hits on the panels will be denoted as $R_{x,i}$ and $R_{y,i}$.

$$\begin{aligned}
 R_{y,i} &= m_{yz} \cdot z_i + q_{yz} \\
 R_{x,i} &= \frac{(m_{yz} \cdot z_i + q_{yz} - q_{xy})}{m_{xy}}
 \end{aligned} \tag{5.11}$$

The bias on the longitudinal position is the difference between the reconstructed coordinate and the simulated position on the panel frame.

5.5 Results

This study aims to verify the consistency of the calibration procedure with the specified requirements. In particular, since the procedure allows to perform an approximate reconstruction of the muon track and, thus, of the extrapolated track position on the panel, it is necessary to carefully determine what is the bias introduced by this procedure, to make sure it is well below the threshold of 4 cm.

Figure 5.7 shows the reconstructed longitudinal coordinate x_{track} , which is determined as the intersection of the reconstructed muon track with the mean z_i coordinate of the panel, for panel 0, taken as an example. Similar patterns were observed for all panels. In this distribution, the bumps are a consequence of the 4/4 requirement and overlap of the panels being located at the extreme edges. It is also important to observe that hits in different regions in x correspond to different straws, as can also be seen in Figure 5.4.

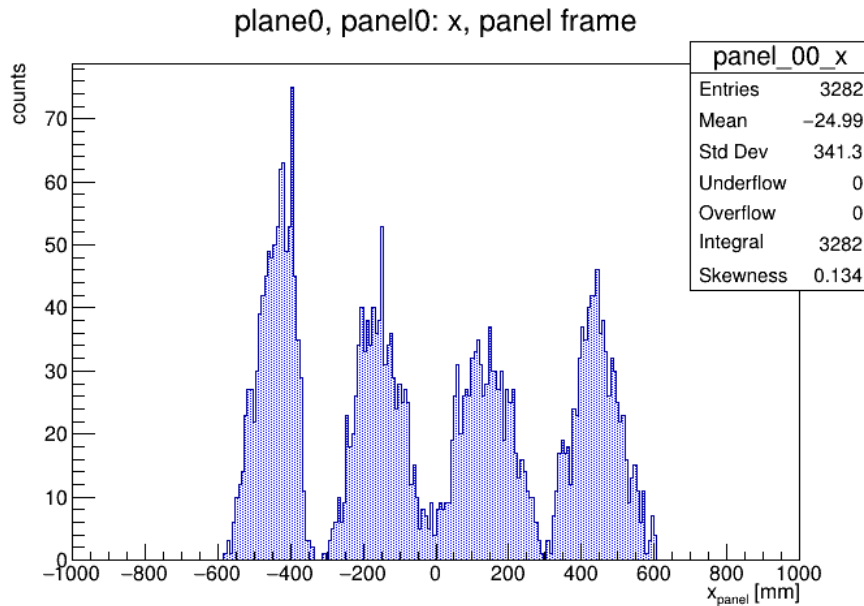


Figure 5.7: The reconstructed longitudinal coordinate in the 0th panel frame.

Figure 5.8 shows the distribution of the longitudinal bias Δx determined as the difference between the projection of the reconstructed track on the panel and its simulated coordinates. The true coordinate is computed as the mean coordinate of the Monte Carlo hits within the panel. The bias range is approximately $[-6,6]$ cm, indicating a significant systematic factor affecting the hit reconstruction. The distribution is similar for all panels. One problem is that this distribution is referred to multiple straws, so for different straws the bias could be very different.




Figure 5.8: The longitudinal bias Δx between the reconstructed and the true hit coordinate on panel 0. The true coordinate is the mean coordinates of Monte Carlo hits.

The systematics brought by this type of reconstruction, both for the horizontal and vertical orientation, arise from the fact that $m_{yz} = \frac{\Delta y}{\Delta z}$ is not accurately reconstructed and the reconstructed coordinates depend on this value, as demonstrated in Section 5.4. This discrepancy arises when the true hit position, far from the straws midpoint, results in incorrectly reconstructed tracks' direction on the $y - z$ plane, leading to the erroneous reconstruction of cosmic muons as horizontal tracks. Figure 5.9 reports the reconstructed m_{yz} distribution, which shows a significant peak at the value of 0, corresponding to cosmic muons erroneously reconstructed as horizontal. No horizontal muons are in the sample, as can be seen from Figure 5.5.

The main difference between the horizontal and vertical orientation can be better understood looking at Figure 5.10 and 5.11. The first one shows the 2D distribution of Δx versus the true position and the second one shows the profile histogram of Δx versus the true coordinates. Figure 5.10, shows four different spots on the x axis corresponding to the overlap regions. Each spot refers to different straws. The two spots on the y axis correspond to cosmics with different orientations.

To assess the systematic impact on the reconstructed coordinates, it is essential to look also at the profile histogram of the bias versus the true coordinates. Profile histograms are used to represent the mean value of y and its error for each bin in x . The default displayed error is the standard error on the mean. Figure 5.11 reveals a systematic effect in determining the longitudinal position within a range exceeding $[-4,4]$ cm. This plot suggests that the mean may not be a reliable estimator of the bias. The main issue with this plot is that, in the vertical orientation, the stereo reconstruction biases for tracks with positive and negative values of m_{yz} do




Figure 5.9: The reconstructed $y - z$ direction distribution of cosmic rays ($m_{yz} = \Delta y / \Delta z$).




Figure 5.10: The 2D histogram of the longitudinal bias versus the true longitudinal coordinate.

not cancel each other out. Considering the vertical orientation and imagining two cosmic muons with opposite $y - z$ orientations, since cosmic rays come from the top and travel downwards, the first particle would hit plane 0 and then plane 1, while the second would hit plane 1 first and then plane 0, with all hits occurring on different panels. For this reason, opposite $y - z$ orientated cosmic rays do not appear on the same




Figure 5.11: The profile of the longitudinal bias versus the true longitudinal coordinate.

profile histogram for a single panel. In the horizontal case, the same systematics apply as in the vertical case. However, in this scenario, both cosmic rays would first hit plane 0 and then plane 1, and since most muons have a vertical trajectory, they would hit most likely the same panels. Therefore, the mean could serve as a reasonable estimator of the bias in this case.

The narrower spot on the left (CAL side) in Figure 5.10 and 5.11 corresponds to the overlap of one panel with another located on a different face, oriented at 90° relative to the first one. Considering face 0 of the first plane and face 0 of the second plane, rotated by 180° about the y axis (the same applies for the other faces), there are three pairs of panels whose relative angle is 90° (Figure 5.3). This type of overlap minimizes the systematic effect and consequently the bias that was explained earlier.

Under these conditions, achieving the required accuracy of the calibration is expected to become challenging.

5.6 Conclusions

This chapter discusses the initial steps towards the tracker calibration. The ultimate goal is to perform a timing calibration of the first assembled station of the tracker using cosmic muons, aiming for a longitudinal hit position resolution better than 4 cm. This requires determining the signal propagation velocity and channel-to-channel delays. The TDCs will measure the signal arrival times, t_1 and t_2 . The calibration will be performed reconstructing the track position along the straw and comparing it with the arrival times. Due to operational constraints of the station, such as gas distribution, fragility, and space, a vertical station calibration was con-

sidered to be the best option for our goals. The focus of this chapter was on the reconstruction of the trajectories of simulated cosmic muons within a vertically oriented station, as well as the analysis of potential biases and systematic errors that could arise from this particular orientation. I reconstructed x_{track} using only the information of whether a straw was crossed or not by a cosmic muon, using only the geometry of the tracker station.

This analysis demonstrated that the vertical orientation for calibrating a station is not optimal. This is due to several reasons:

- the selection criteria in Section 5.3 select cosmic muons with orientations that affect panel illumination. The illumination is non-uniform, and moreover, there are almost no hits in the central region of the panel, which could result in waveform shape variations;
- the longitudinal bias Δx range is approximately $[-6,6]$ cm, indicating a significant systematic factor affecting the hit reconstruction. The systematics arise from the fact that m_{yz} is not accurately reconstructed. In fact, if the true hit position is far from the midpoint of the straws, the reconstructed tracks' direction on the $y - z$ plane is completely different from the true one;
- the 2D distribution of the Δx versus the true position shows four distinct spots on the x axis corresponding to the overlap regions. Each spot refers to different straws. The profile histogram of this distribution reveals a systematic effect in determining the longitudinal position within a range greater than $[-4,4]$ cm. The main issue is that, in the vertical orientation, the stereo reconstruction biases for tracks with positive and negative values of m_{yz} do not cancel each other out.

In addition to these results, there is one last factor we need to consider: the rate and the consequent data taking time. For the vertical orientation, the rate must be scaled by two additional factors compared to the horizontal one. Considering only muons at 45° (Figure 5.5), the first factor, approximately $1/\cos^2\theta \sim 1/2$, accounts for the angular dependence of the flux, and the second factor, $1/\sqrt{2}$, accounts for cosmic rays striking the station at a 45° angle. However, estimating the rate is quite complicated. In fact, these are not the main contributors to the difference between the two rates. With the horizontal orientation, it is possible to reconstruct a particle track without the 4/4 requirement, instead only requiring that the particle crosses two panels on different faces, as most muons are vertical. On the other hand, if the vertical orientation is the only possible configuration, the analysis must account for all the biases in every straw. Consequently, much more data is required, as the number of parameters in the analysis increases.

Therefore, an alternative orientation or method should be considered to optimize the calibration process and achieve more accurate results. For this reason, a new mechanical solution to address the problems outlined in Section 5.1.1 is currently under development.

Chapter 6

Pre-pattern recognition studies

The estimated data volume for Mu2e data-taking is at least 7 PByte per year: it will thus be crucial to exploit all the possible handles to optimize the CPU and memory usage. For example, the simulation shows that the primary source of hits in the Mu2e tracker will be low energy electrons and positrons, called δ -electrons. Therefore, the Mu2e Collaboration has made a huge effort to develop solutions to identify and flag these hits as soon as possible in the data-taking and not use them in pattern recognition and track reconstruction. The most important constraint is making sure that hits generated by CE are not erroneously flagged as δ -electrons, since this would compromise CE track reconstruction. There are currently two main algorithms being developed for this purpose. This Chapter reports on the first systematic study performed to compare the performance of the two algorithms and determine the best solution for data-taking.

6.1 δ -electrons as source of background

The performance of the detectors and the Mu2e physics reach have been thoroughly studied with the Monte Carlo simulation. In terms of occupancy, we know that the dominant source of hits in the tracker are low energy electrons and positrons, in the following referred to as δ -electrons. To be more precise, in decreasing order of importance, the primary sources of hits are electrons from Compton scattering, electron-positron pairs, and delta rays. Compton-scattered electrons are produced when photons, generated by various processes, interact with the detector material. These photons primarily originate from neutron captures, which excite the nuclei and lead to subsequent photon emission. Typically, these photons have energies of a few MeV. Neutrons are produced in the process of nuclear muon capture. Pair-production electrons and positrons are generated during nuclear processes, where pairs of electrons and positrons are created. Delta rays, or secondary ionization electrons, are generated when high-energy charged particles interact with the detector material.

6.1.1 Compton scattering

The Compton effect (Figure 6.1) is the scattering of a photon by a free or quasi-free electron. An electron is considered "quasi-free" when the energy of the incoming photon is significantly higher than the electron's binding energy ($E_\gamma \gg E_B$). The process is termed Compton scattering if the electron is ejected from the atom, carrying away the recoil momentum. This effect is most prominent in an extended energy region around 1 MeV, with the region being much larger for low Z materials compared to high Z materials.




Figure 6.1: The Compton effect [79].

Since the photon scatters quasi-elastically off the electron, the energy and angle of the scattered photon are correlated. To describe this relationship, we use the 4-momenta defined as follows: $k = (E_\gamma, \mathbf{k}c)$ and $p_e = (m_e c^2, 0)$ represent the 4-momenta of the photon and the electron (at rest) before scattering, and $k' = (E'_\gamma, \mathbf{k}'c)$ and $p'_e = (E'_e, \mathbf{p}'_e c)$ represent the 4-momenta after scattering. The angle between the scattered photon and the incident photon is denoted as θ_γ , while the angle of the electron is denoted as θ_e . By applying energy-momentum conservation:

$$k + p_e = k' + p'_e \quad (6.1)$$

$$(k - k')^2 = (p'_e - p_e)^2 \Rightarrow -k \cdot k' = m_e^2 c^4 - p'_e \cdot p_e \quad (6.2)$$

$$\Rightarrow E_\gamma E'_\gamma (1 - \cos \theta_\gamma) = m_e c^2 (E'_e - m_e c^2) = m_e c^2 (E_\gamma - E'_\gamma) \quad (6.3)$$

The right-hand side of the last equation uses the kinetic energy of the electron:

$$T = E'_e - m_e c^2 = E_\gamma - E'_\gamma \quad (6.4)$$

which follows from the energy part of equation 6.1. The energy of the scattered electron as a function of the photon scattering angle is derived from equation 6.2:

$$E'_e = \frac{E_\gamma \cdot \epsilon \cdot (1 - \cos \theta_\gamma)}{1 + \epsilon(1 - \cos \theta_\gamma)} + m_e \quad (6.5)$$

where $\epsilon = \frac{E_\gamma}{m_e c^2}$.

The differential cross section per (free) electron, known as the Klein-Nishina formula, is calculated using methods from quantum electrodynamics:

$$\frac{d\sigma}{d\Omega} = \frac{r_e^2}{2} \frac{1 + \epsilon(1 - \cos \theta_\gamma)}{[1 + \epsilon(1 - \cos \theta_\gamma)]^2} \left(1 + \cos^2 \theta_\gamma + \frac{\epsilon^2(1 - \cos \theta_\gamma)^2}{1 + \epsilon(1 - \cos \theta_\gamma)} \right) \quad (6.6)$$

An electron bound in an atom can only be considered quasi-free if the photon's energy is significantly higher than the electron's binding energy. As the photon energy increases, more shell electrons become quasi-free, leading to the Compton cross section per atom approaching proportionality to Z , with individual electrons contributing incoherently:

$$\sigma_C^{\text{atom}} = Z\sigma_C \quad (6.7)$$

where σ_C is the Klein-Nishina cross section for a single free electron. The Compton cross section decreases at lower energies, where coherent scattering (Rayleigh scattering) off the entire atom (without ionizing the electron shell) becomes dominant.

By reformulating the Klein-Nishina formula, one can obtain the differential dependence of the Compton cross section on the kinetic energy of the recoil electron $T = E_\gamma - E'_\gamma$:

$$\frac{d\sigma}{dT} = \frac{\pi r_e^2}{m_e c^2 \epsilon^2} \left[2 + \frac{t^2}{\epsilon^2(1-t)^2} + \frac{t}{1-t} \left(t - \frac{2}{\epsilon} \right) \right] \quad (6.8)$$

where $t = T/E_\gamma$. Because the scattering process is elastic, there is a one-to-one relationship between the energy and angle θ_e of the electron:

$$\cos \theta_e = \frac{T(E_\gamma + m_e c^2)}{E_\gamma \sqrt{T^2 + 2m_e c^2 T}} = \frac{1 + \epsilon}{\sqrt{\epsilon^2 + 2\epsilon/t}} \quad (6.9)$$

The maximum energy transfer to the electron is obtained from equation 6.5 for backward scattering of the photon ($\theta_\gamma = 180^\circ$), corresponding to forward scattering of the electron ($\theta_e = 0^\circ$). The electron's kinetic energy reaches its maximum value in this case, $T \rightarrow T_{\text{max}}$. In the measured energy spectrum, this leads to the so-called "Compton edge" at:

$$T_{\text{max}} = \frac{E_\gamma \cdot 2\epsilon}{1 + 2\epsilon} \quad (6.10)$$

which lies slightly below the photopeak. The energy difference between the photopeak and the Compton edge $E'_\gamma(\theta = \pi)$ decreases with increasing E_γ and approaches:

$$E'_\gamma(\theta = \pi) \approx \frac{m_e c^2}{2} \text{ for } E_\gamma \gg m_e c^2 \quad (6.11)$$

6.1.2 Pair production

In the Coulomb field of a charge, a photon can convert into an electron-positron pair (Figure 6.2)¹.




Figure 6.2: The pair production [79].

The energy of the photon must exceed twice the electron mass plus the recoil energy transferred to the field-producing charge. For most elements, pair production predominantly occurs in the Coulomb field of the nucleus. For nuclei, the recoil energy is usually negligible, leading to a threshold energy for pair production of:

$$E_{\gamma} \geq 2m_e c^2 + 2 \frac{m_e^2}{m_{\text{nucleus}}} c^2 \quad (6.12)$$

If the nuclear charge is not screened by atomic electrons (for low energies, the photon must come relatively close to the nucleus to make pair production probable, meaning it interacts with the "bare" nucleus),

$$1 \ll \epsilon \ll \frac{1}{\alpha Z^{1/3}} \quad (6.13)$$

the pair-production cross section is given by:

$$\sigma_{\text{pair}} = 4\alpha r_e^2 Z^2 \left(\frac{7}{9} \ln 2\epsilon - \frac{109}{54} \right) \text{ cm}^2/\text{atom} \quad (6.14)$$

However, for complete screening of the nuclear charge ($\epsilon \gg 1/\alpha Z^{1/3}$):

$$\sigma_{\text{pair}} = 4\alpha r_e^2 Z^2 \left(\frac{7}{9} \ln \frac{183}{Z^{1/3}} - \frac{1}{54} \right) \text{ cm}^2/\text{atom} \quad (6.15)$$

At high energies, pair production can occur even at relatively large impact parameters between the photon and the nucleus. In this case, the screening effect of atomic electrons must be considered. For large photon energies, the pair-production

¹Photon emission by an electron (bremsstrahlung) and pair production are closely related processes. By modifying the bremsstrahlung diagram—changing the outgoing photon to an incoming one and the incoming electron to an outgoing positron—one obtains the pair production diagram. The matrix elements of these processes are related, at least in the lowest order. Consequently, both processes are treated together in the foundational work by Bethe and Heitler, often referred to as the 'Bethe-Heitler processes'.

cross section approaches an energy-independent value as given by Equation 6.15. Ignoring the small term in the equation, the asymptotic value of $1/54$ is expressed as:

$$\sigma_{\text{pair}} \approx \frac{7}{9} \cdot 4\alpha r_e^2 Z^2 \ln \left(\frac{183}{Z^{1/3}} \right) \approx \frac{7}{9} \cdot \frac{1}{X_0} \cdot \frac{A}{N_A \cdot \rho} \quad (6.16)$$

The energy is uniformly distributed between the produced electrons and positrons at low and medium energies, but becomes slightly asymmetric at high energies.

The field of the nucleus is formed by the coherent sum of Z nucleon charges, leading to the Z^2 dependence of the pair production cross section.

Even with large momentum transfers Δp to the nucleus, the energy transfer $(\Delta p)^2/2M$ remains small due to the large nuclear mass M . After pair creation, the remaining energy is equally divided between the e^+ and the e^- .

6.1.3 Delta rays

δ -rays, or *knock-on* electrons, are produced when a projectile particle collides centrally with shell electrons, resulting in significant energy transfers. These electrons gain high kinetic energy and can be described through elastic collisions with quasi-free electrons. By considering the energy-momentum conservation relation and using the Lorentz factors γ and β , the relationship between the kinetic energy T of the δ -ray and the emission angle θ can be derived as:

$$\cos \theta = \frac{T(\gamma + m_e/M)}{\gamma\beta\sqrt{T^2 + 2Tm_e c^2}} \quad (6.17)$$

$$T(\theta) = \frac{2m_e c^2 \beta^2 \gamma^2 \cos^2 \theta}{\gamma^2(1 - \beta^2 \cos^2 \theta) + 2\gamma m_e/M + m_e^2/M^2} \quad (6.18)$$

The maximum energy transfer T_{max} occurs at $\theta = 0^\circ$, while the minimum energy, T_{min} , occurs at $\theta = 90^\circ$. At highly relativistic energies ($\gamma \gg 1$ and $\theta \gg 1/\gamma$), the energy-angle relationship becomes independent of the incoming particle's properties.

The rate of δ -rays per energy interval dT and path length dx is given by:

$$\frac{d^2 N}{dx dT} = n_e \frac{d\sigma}{dT} \quad (6.19)$$

which, when combined with the electron density and the differential cross section, becomes:

$$\frac{d^2 N}{dx dT} = \frac{1}{2} z^2 \frac{Z}{A} K \rho \frac{1}{\beta^2} \frac{F(T)}{T^2} \quad (6.20)$$

Here, K is the constant from the Bethe-Bloch formula, and $F(T)$ is a function accounting for spin dependence. Integration over T and x provides the number of δ -rays in a medium of thickness Δx :

$$N = \frac{1}{2} z^2 \frac{Z}{A} K \rho \Delta x \frac{1}{\beta^2} \left(\frac{1}{T_{\min}} - \frac{1}{T_{\max}} \right) \approx 0.077 \frac{\text{MeV cm}^2}{\text{g}} z^2 \rho \Delta x \frac{1}{T_{\min}} \quad (6.21)$$

The emission angle dependence is given by:

$$\frac{dT}{d \cos \theta} = 4 m_e c^2 \frac{\cos \theta}{\sin^4 \theta} \quad (6.22)$$

Substituting this into the rate equation yields:

$$\frac{d^2 N}{dx d \cos \theta} = \frac{1}{2} z^2 \frac{Z}{A} K \rho \frac{1}{\cos^3 \theta} \frac{1}{m_e c^2} \approx 0.15 \frac{\text{cm}^2}{\text{g}} z^2 \rho \frac{1}{\cos^3 \theta} \quad (6.23)$$

This expression diverges as θ approaches 90° , where T approaches zero, indicating a limitation in the assumption of a free electron. The resulting distributions suggest that δ -rays emitted at small angles can significantly affect the spatial resolution in detectors, particularly through ionization clusters that broaden the track of the mother particle.

6.2 Monte Carlo samples

For our studies, we used three different Monte Carlo samples: two samples of CE signal generated at two different proton pulse intensities, and one sample of antiproton annihilation events. The production of particles within the PT and their tracking from the PS to the DS is handled by the Mu2e Offline software (Appendix B). The simulation of particle interactions and the event processing are based on GEANT4, and they are handled by the art framework and data management is governed by the SAM system (Appendix B).

Mu2e uses a multi-stage simulation to generate and simulate events efficiently. The method involves generating events, partially simulating them, and saving the intermediate results. Later stages resume the simulation from the saved state. This approach optimizes both the time required for the event generation and the disk space usage.

At the first stage, the interactions of protons at the PT are simulated. Produced secondary particles are traced up to the DS and the information about particles which reached the DS is stored. At the second stage, the surviving particles are propagated through the upstream portion of the DS and muons stopped in the ST are recorded. The third and the following stages deal with the simulation of the physics processes in the Mu2e detector and the hit generation and digitization. For the CE dataset, a fraction of muons are assumed to decay into CEs. For the \bar{p} dataset, \bar{p} annihilation at rest in the ST is simulated based on the positions and times of the stopped antiprotons. Background electrons from annihilation result from decays such as $\pi^0 \rightarrow \gamma\gamma$, followed by photon conversions, and $\pi^- \rightarrow \mu^- \nu$,

followed by μ^- decays. During this stage, the raw simulated data are digitized into simple C++ classes or structs, using the detector's raw data.

A typical Mu2e event includes multiple pileup hits from particles produced by muon captures in the ST, as well as particles entering the DS from the TS. The pileup hits make the majority of the detector's hits. The pileup level depends on the proton pulse intensity. The Mu2e pileup simulation assumes that the pulse intensity varies on a time scale much longer than $2 \mu\text{s}$, meaning all proton pulses around the simulated one have the same intensity. Under this assumption, a transformation is applied to hits with time $T_i > 1695 \text{ ns}$ outside the microbunch limits, assigning them a residual time of $t_i = T_i \div 1695 \text{ ns}$, effectively accounting for hits from previous proton pulses that would otherwise be attributed to later microbunches.

For the low-intensity mode, with a mean intensity of 1.6×10^7 protons/pulse, which in Mu2e jargon is named "1BB", approximately 25,000 muons stop in the ST per pulse. In the high-intensity mode, named "2BB", this number is about 2.5 times higher (Section 2.4.1).

The datasets used for CE signal plus pileup analysis will be referred to as $CE - 1BB$ and $CE - 2BB$ for 1BB and 2BB pileup, respectively. The dataset for antiproton analysis without pileup will be referred to as $PBAR - 0BB$. For datasets with pileup, the pileup hits are explicitly added to the hits from the signal process. The antiproton sample, however, has been simulated w/o the pileup.

6.3 δ -electrons in the Mu2e tracker

In Mu2e, low-energy electrons and positrons, with a momentum below 20 MeV/c and referred to as δ -electrons, generate the majority of hits in the tracker. Identifying the δ -electrons hits as early as possible is important for improving the track reconstruction efficiency and the optimization of memory and the CPU usage. There are also several physics reasons why it would be important to identify those hits correctly:

- the main Mu2e goal is the search for the CE signal: flagging erroneously even a small fraction of the hits generated by the CE as δ -electrons reduces the CE track reconstruction efficiency. Figure 6.3 shows that in the Monte Carlo sample $CE - 1BB$, the simulated CE hits are just 1% of the total hits in the tracker;
- counting the number of protons will be a complementary procedure to the STM to determine the muon stopping rate: the simulation shows that is possible to estimate the number of muons captured in the stopping target by counting the number of protons produced in the nuclear muon capture;
- misidentification of muons and pions as δ -electrons may result in erroneous background estimate. This is particularly significant for the antiproton background. $p\bar{p}$ annihilation at rest in the ST can produce signal-like electrons,

which constitute a background to the CE search. An estimate of this background is presently affected by a large systematic uncertainty. Mu2e has developed a data-driven procedure to improve the estimate. $p\bar{p}$ annihilation at rest in the ST can produce events with two or more tracks, each with a momentum around 100-200 MeV/c. For $p\bar{p}$ annihilation, the rate of multi-track events is about 500 times higher than the rate of events with a single signal-like electron. For 10^4 $p\bar{p}$ annihilation events generated, about 3.7% of the events contained two reconstructable particle tracks. Therefore, the identification and reconstruction of multi-track events could be used to constrain the $p\bar{p}$ background. Thus, it is crucial not to flag hits generated by muons or pions as δ -electron hits, since this would compromise the reconstruction efficiency.

Figure 6.3 shows the momentum distribution (Monte Carlo truth) of particles that make at least one hit in the tracker for a simulated CE sample ($CE - 1BB$, Section 6.2). For each Monte Carlo particle, the histogram is filled a number of times equal to the number of hits generated by that particle. The distribution shows that the majority of hits originate from low-energy electrons and positrons (orange), which constitute approximately 75% of the total number of hits. There is an asymmetry between the number of hits below 20 MeV/c produced by electrons and positrons: electrons account for 71% of all hits in the tracker, while positrons contribute only 4%. The difference arises because electrons are produced also by Compton scattering, which is the primary source of hits in the energy range of 1 MeV. This difference will be crucial in the following sections when discussing the δ flagging efficiency.

The distribution also shows that 14% of the total hits are due to protons, which are produced by nuclear processes. According to the simulation, the bump at low energy in the Monte Carlo momentum distribution of protons and deuterons producing hits in the tracker (green) results from inelastic neutron scattering. Larger momentum values correspond to protons produced in muon captures at rest. Their kinetic energy ranges from about 5 to 20 MeV, resulting in low $\beta\gamma$ values, which makes them heavily ionizing particles. The deposited energy will be one of the variables used to discriminate protons. The distribution has a maximum and then decreases at 250 MeV/c, which is primarily due to the finite length of the tracker and the increase in longitudinal momentum of protons within the tracker acceptance.

It is important to note that the bump around 50 MeV/c in the positron distribution should not be present. The source is so far unexplained. We expect $N(\mu^+ \rightarrow e^+)/N(\mu^- \rightarrow e^-)$ to be about 10^{-3} for muons entering the DS. The DIO on the IPA (Section 2.5.3) should also be around 10^{-3} compared to the DIF of negative muons. The simulation of μ^+ may contain some errors and we are still investigating this discrepancy. However, this issue is not problematic for the analysis of low-momentum electrons and positrons, as the momentum ranges are different.

Figure 6.4 shows the momentum distribution of particles produced in the $p\bar{p}$ annihilation at the ST. The data sample ($PBAR - 0BB$) is described in Section 6.2. For each Monte Carlo particle, the histogram has an entry per hit generated by that particle. Particles produced in the $p\bar{p}$ annihilation are mostly pions and muons.




Figure 6.3: Momentum distribution (Monte Carlo truth) of particles producing at least one hit in the tracker ($CE - 1BB$ data sample). The momentum distribution of all particles making hits is depicted in dark blue, with electrons shown in pink, positrons in light blue, δ -electrons in orange, protons and deuterons in light green, and CEs in dark green.

The momentum distribution has a maximum in the 100-200 MeV/c range. Photons are also produced, and they can undergo Compton scattering and pair production, which explains the presence of a δ -electron peak that is about ~ 100 times lower than the one in Figure 6.3.

Figure 6.3 shows that the majority of the hits originate from δ -electrons, with protons and deuterons being the second most common source. Figure 6.5 presents an example of comparison of an event display before (Left) and after (Right) the background hits have been flagged and removed. Flagging those hits is crucial for several reasons: it prevents unnecessary data from being sent to the pattern recognition algorithms, improving their efficiency and conserving the CPU resources.

6.4 δ -electrons flagging algorithms

A brief description of the Mu2e Offline software tools is reported in Appendix B, and the reconstruction process is described in Appendix C. Flagging δ -electrons is done before the time clustering and pattern recognition.

Mu2e has developed two types of hit flagging algorithms:

- *FlagBkgHits*, described in Section 6.4.1;




Figure 6.4: Momentum distribution (Monte Carlo truth) of particles producing at least one hit in the tracker ($PBAR - 0BB$ data sample). The momentum distribution of all particles making hits is depicted in dark blue, with electrons shown in pink, positrons in light blue, δ -electrons in orange, protons and deuterons in green, pions in light brown and muons in black.

- *DeltaFinder*, described in Section 6.4.2.

6.4.1 The *FlagBkgHits* algorithm

The detailed description of multivariate analysis (MVA) and the process of MVA training is beyond the scope of this work. Nevertheless, since this technique is one of the options developed for the background hit flagging, I will briefly outline the fundamental principles involved.

When searching for patterns in a multivariable space, a common procedure involves defining a set of statistical models that analyze the measured variables and estimate the probability that these are consistent with the sought pattern. Once the variables are selected, the MVA is trained to recognize patterns by evaluating examples known to the trainer, allowing for the feedback to refine the pattern identification procedure.

The first selection concerns the mean energy deposited in the *ComboHits* used to create a *StereoHit* and excludes those with a deposited energy above 5 keV. The algorithm then classifies as the proton hits all hits with a deposited energy above

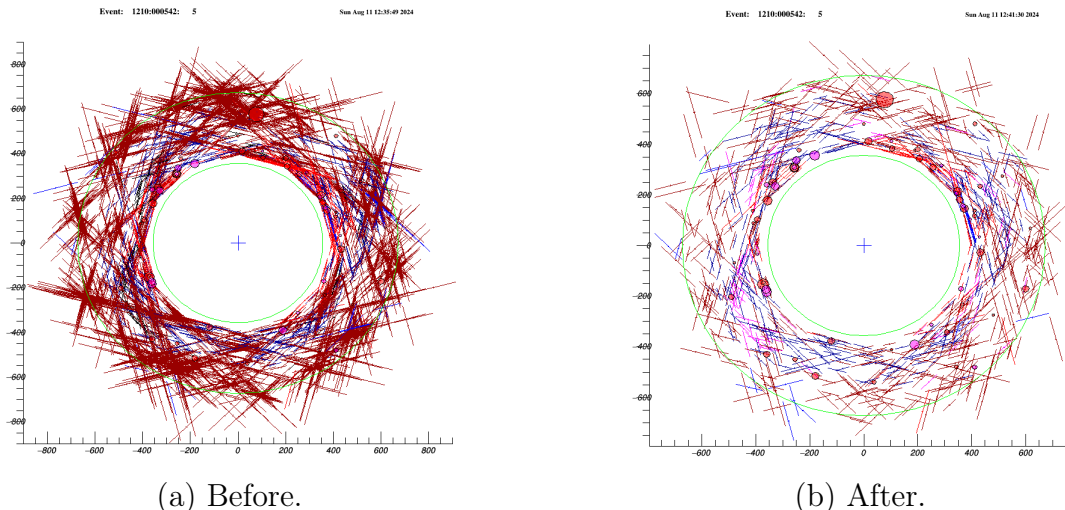


Figure 6.5: Before and After background hits flagging. The $x - y$ plane views of a CE event with 2BB pile-up in the tracker (Section 2.4.1). The segments are the *hit* tracker straws. The hits marked in red are from electrons and the ones in blue are from positrons. Left (Right): before (after) δ -electron hit flagging.

4.5 keV. Since δ -electrons tend to form small, dense clusters of hits on $x-y$ plane, an algorithm based on a clustering approach was developed: concentrated clusters in the $x-y$ plane are sought using a clustering algorithm, as δ -electrons are highly likely to create such hit clusters. The *FlagBkgHits* algorithm uses an artificial neural network (ANN)-based technique to classify those clusters.

The ANN was trained in a supervised mode, using CE and δ -electrons hits from Monte Carlo data sample.

The resolution on the measurement of the hit position in the wire direction is limited to a few cm. However, an improved position measurement can be achieved by using various layers of straws that have large overlaps in the transverse plane. By obtaining two hit measurements from a pair of intersecting straws and ensuring they fall within a time window on the order of the maximum drift time, one can deduce that these hits were produced by the same particle and occurred at the intersection of the two straws in the $x - y$ plane (*StereoHit*).

The clustering process uses the x and y coordinates of the selected hits. A random hit is chosen to define the initial cluster, after which the following iterative steps are applied. The centroid of each cluster is computed, and hits whose distances from a cluster centroid fall within a specified inner threshold and time window are added to that cluster. Hits with distances from all existing clusters greater than an outer threshold are used to seed new clusters, while hits that fall between the two thresholds remain unassigned. This process generates a new set of clusters, preparing input for the next iteration. During each iteration, every hit is reconsidered as a potential new point in a cluster, including those already assigned. The clustering process continues until convergence is achieved, i.e., when an iteration no longer

results in any changes.

6.4.2 The *DeltaFinder* algorithm

The *DeltaFinder* algorithm has been designed to identify δ -electron hit patterns. This algorithm relies on the fact that δ -electrons typically form a straight line in the r - z plane (cyan lines in Figure 6.6) and appear as a spot with a diameter of less than ~ 3 cm in the x - y plane within the 1 T magnetic field. On the other hand, CEs create entirely different patterns of hits, which appear as oblique segments in the r - z plane due to their helical trajectories (red lines in Figure 6.6).

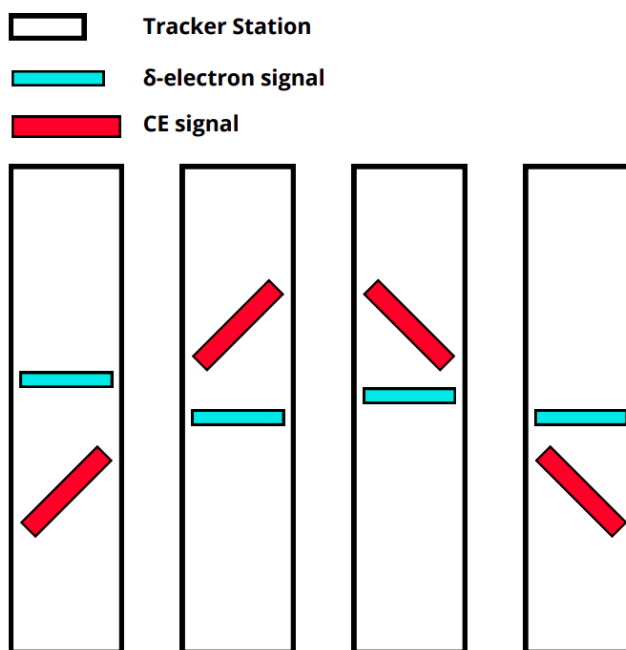


Figure 6.6: δ -electrons and CE patterns in the $r - z$ plane. The white rectangles represent four of the eighteen tracker stations. The cyan straight lines represent a possible δ -electron patterns, while the red lines a possible CE pattern.

6.4.2.1 Step 1: identifying δ -electron segments

DeltaFinder first tries to identify δ -electron track segments within each station individually. These segments, parallel to the beam axis, are called *seeds*. Since hits from the same electron should be close in both time and space, and δ -electrons may hit multiple straws within the same panel, the algorithm clusters these straw hits in space and time, applying various cleanup cuts to ensure the selected patterns are consistent with the δ -electron hit patterns. The maximum allowed time difference between two hits within a station to form a *seed* is set to 40 ns.

This cleanup based on the x - y coordinates is performed by computing a χ^2 . The *seed* is reconstructed using three, or four *ComboHits*. In each case, the intersections

between pair of straws are determined, and the x - y coordinates of the *seed* are given by the center of gravity of these intersections.

The algorithm extends the *seeds* by requiring hits to be sufficiently close to the intersection point in time and space, performing multiple checks to avoid over-efficiency in hit flagging. For each *seed*, the mean deposited energy is calculated from the energy deposited in all *ComboHits*. This selection optimizes data processing by reducing the total number of hits that need to be analyzed.




Figure 6.7: A δ candidate *seed*. The four coloured segments are the tracker straws that were hit in the same station.

It is necessary to clarify one detail. Figure 6.8 shows the distribution of the simulated deposited energy in the tracker for δ electrons, CEs, and protons in the case of a 1BB pileup.

To optimize data processing, a hit energy threshold could be applied to the *DeltaFinder* to reduce the total number of hits that need to be analyzed, thereby speeding up the process. Moreover, only about 4% of CE hits have energies above 3.5 keV (and 1% above 5 keV), so the loss of CE hits would be minimal, resulting in a faster overall algorithm.

However, such an energy cutoff would have a significant impact for the algorithm's performance. Starting with fewer hits, especially in stereo intersections, reduces the efficiency of identifying the correct *seed*.

6.4.2.2 Step 2: connecting the *seeds*

After the selection based on mean deposited energy, *DeltaFinder* attempts to connect segments that are close in both the x - y plane and time across different stations to form δ -electron candidates. A valid candidate must have at least two segments and a minimum of five *ComboHits*. Accidentally reconstructed segments of 100




Figure 6.8: The Monte Carlo deposited energy distribution in the tracker ($CE-1BB$ data sample). The red distribution refers to CEs, while the green and the blue one to δ -electrons and protons respectively. The peaks and tails correspond to particles with such high energy deposition that they result in a saturated waveform.

MeV electrons remain unconnected due to their separation in the x - y plane.

DeltaFinder links δ -electron *seeds* across stations, attempting to associate new *seeds* with already found δ -electron candidates. If no match is found, a new candidate is created. Good δ -electron candidates are marked, and their hits are flagged to prevent their inclusion in proton candidate searches.

6.4.2.3 Step 3: identifying proton candidates

Finally, *DeltaFinder* identifies proton candidates using the *seeds* with the mean deposited energy above 3 keV. First, it checks if a *seed* is consistent in time with any existing proton candidate. If no consistency is found, the hits of the *seed* are added to a new proton candidate.

6.5 Performance analysis and comparison

The comparison between *FlagBkgHits* and *DeltaFinder* is performed at two levels, with each testing different aspects of the algorithms:

- **Hit-level comparison:** this phase focuses on estimating how accurately individual hits are flagged, providing the most direct method for assessing and comparing the performance of the two algorithms. This stage allows for an un-

biased comparison without the influence of subsequent reconstruction stages. It also includes a direct evaluation of proton counting;

- **High-level comparison:** this phase focuses on the comparison of the tracks reconstruction efficiencies. The main figure of merit at this stage is the reconstruction of CE tracks.

6.5.1 Hit-level comparison

Before starting with the hit-level comparison of the two algorithms, during the analysis, we first observed an over-efficiency in the proton hit flagging by *DeltaFinder*. This primarily impacted the flagging of protons, muons, and pions, as shown in Table 6.1 and Table 6.2. In Tables 6.1, 6.2, 6.3, 6.4 and 6.5, f_p and f_e represent the number of *ComboHits* flagged as protons and electrons, respectively, divided by the total number of *ComboHits*. DF and FBH denote the *DeltaFinder* and *FlagBkgHits* algorithms, respectively. Each row corresponds to the particle under study.

	f_p	f_e
p	96.0%	1.0%

Table 6.1: *DeltaFinder* proton flagging results before the adjustment -selecting a proton candidate with more than four hits and a deposited energy above 3 keV- (*CE – 1BB* data sample). f_p and f_e represent the number of proton *ComboHits* flagged as proton and electron hits, respectively, divided by the total number of the proton/electron *ComboHits*.

	f_p	f_e
μ	5.8%	5.0%
π	2.5%	11.2%

Table 6.2: *DeltaFinder* muon and pion flagging results before the adjustment -selecting a proton candidate with more than four hits and a deposited energy above 3 keV- (*PBAR – 0BB* data sample). f_p and f_e represent the number of particle *ComboHits* flagged as proton and electron hits, respectively, divided by the total number of *ComboHits*.

The number of muon (5.8%) and pion (2.5%) hits flagged as protons is quite high, since as already mentioned in Section 6.3 the fraction of protons in this type

of events is really low. Figure 6.9 shows the distribution of the total number of muon *ComboHits* (red) and those flagged as δ -electrons (blue) and protons (green) as a function of the particle momentum. The distribution shows that a significant number of muons hits are misidentified as protons hits at low momentum. According to the Bethe-Bloch formula, such hits have higher energy loss and can thus be most likely flagged as protons. We thus implemented a number of corrections to the algorithm. Since *seeds* may have accidentally attached hits, we imposed a condition requiring a *good* proton candidate to have more than four hits with a deposited energy greater than 3 keV.

This adjustment reduces the proton hit flagging efficiency by a $\sim 10\%$ (Tables 6.4 and 6.3). However, it significantly reduces the number of muons and pions flagged as protons, by approximately a factor of 2 and a factor of 6 (Table 6.5), respectively.




Figure 6.9: The distribution of the total number of muon *ComboHits* (red) and of the muon *ComboHits* flagged as δ -electrons (blue) and protons (green) as a function of the particle momentum (*PBAR* – *0BB* data sample).

Moving to the performance and comparison between the two algorithms, Tables 6.3 and 6.4 show the hit-level comparison between the two algorithms for the *CE* – *1BB* and *CE* – *2BB* data samples.

The fraction of electron and positron hits flagged as δ -electron hits differs due to the different momentum distributions of these particles at low energies. Figure 6.10 shows a zoomed-in view of the δ -electron energy distribution (Monte Carlo truth) for particles with at least one hit in the tracker for energies below 2 MeV. Electrons are shown in pink and positrons in black. Each bin represents the number of hits corresponding to a specific Monte Carlo particle. The pink distribution peaks near

	f_p DF	f_e FBH	f_e DF
e^- ($p < 20$ MeV/c)	3.7%	75.1%	72.7%
e^- ($20 < p < 80$ MeV/c)	2.2%	49.0%	29.9%
e^- ($80 < p < 110$ MeV/c)	0.8%	7.4%	4.3%
p	83.6%		2.2%
e^+ ($p < 20$ MeV/c)	0.4%	84.2%	87.9%

Table 6.3: Electrons, positrons and protons hit-level comparison ($CE - 2BB$ data sample). FBH and DF denote the *FlagBkgHits* and the *DeltaFinder* algorithms, respectively. f_p and f_e represent the number of particle *ComboHits* flagged as proton and electron hits, respectively, divided by the total number of *ComboHits*.

	f_p DF	f_e FBH	f_e DF
e^- ($p < 20$ MeV/c)	2.5%	75.9%	72.5%
e^- ($20 < p < 80$ MeV/c)	1.0%	50.0%	27.4%
e^- ($80 < p < 110$ MeV/c)	0.3%	5.7%	3.4%
p	83.7%		1.0%
e^+ ($p < 20$ MeV/c)	0.2%	85.5%	88.5%

Table 6.4: Electrons, positrons and protons hit-level comparison ($CE - 1BB$ data sample). FBH and DF denote the *FlagBkgHits* and the *DeltaFinder* algorithms, respectively. f_p and f_e represent the number of particle *ComboHits* flagged as proton and electron hits, respectively, divided by the total number of *ComboHits*.

the electron mass, that corresponds to a value of $\theta_\gamma \sim 0$ (Section 6.1.1), while the black distribution falls down to zero.

Figures 6.11a and 6.11b show the efficiency (i.e., the number of δ -electron *ComboHits* flagged as δ -electrons over the total number of *ComboHits* for each particle type) as a function of particle momentum for positrons (red) and electrons (blue) using the *FlagBkgHits* (left) or *DeltaFinder* (right) algorithms.

This plot shows a dependence on the particle momentum. The efficiency plot is, in fact, convoluted with the momentum distribution.

At low momentum values (below 1-2 MeV/c), the efficiency tends to be higher for positrons. This is because the total number of positron hits with such low momentum is extremely small, as positrons are not produced by Compton scattering. A common case of failure occurs when particles produce only a single hit in the tracker, making both stereo reconstruction and *seed* reconstruction impossible. Compton electrons typically produce one or two hits per station, which does not allow to properly flag them. There are lots of electron *ComboHits* in this energy range which remain unflagged. This explains the lower electron hit flagging efficiency.

However, *FlagBkgHits* performs better with electrons having $E < 2$ MeV. This is because the minimum number of hits required to generate a seed with *DeltaFinder* is three, and in this energy range, many particles produce only one or two hits in

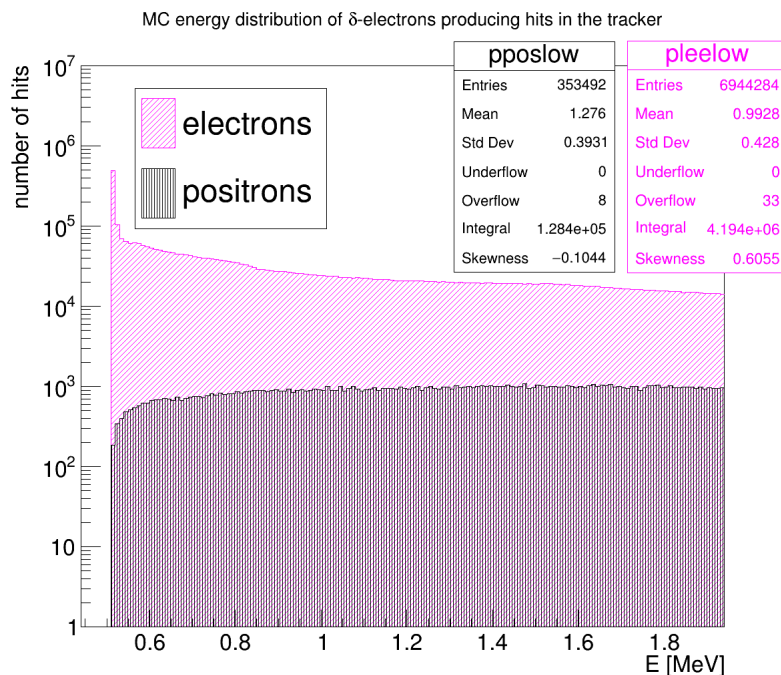


Figure 6.10: The electron (pink) and positron (black) energy distribution for $E < 2$ MeV.

the tracker. Since positrons are not present in this energy range (as they are not produced by Compton scattering), the efficiency of positron hit flagging is higher for *DeltaFinder* (Tables 6.3 and 6.4).

At higher momentum, the two algorithms have a similar performance and the differences between positrons and electrons become negligible. Figure 6.11 shows that as the momentum increases, the efficiency decreases. This occurs because higher momenta correspond to larger radii and a greater spread of hits on the x - y plane. This trend is also observed for particles in the next selected momentum range ($20 \text{ MeV}/c < p < 80 \text{ MeV}/c$).

A common case of δ flagging failure is when hits occur on straws that are parallel to each other, especially those near the center of the tracker, where stereo and *seed* reconstruction is not possible.

FlagBkgHits flags 70% more CE hits as δ -electrons compared to *DeltaFinder* (Tables 6.3 and 6.4). The main difference between the two algorithms lies in the primitives: *StereoHits* and *seeds*. If the same δ -electron hits three different straws, the *StereoHit* information results in multiple intersection points being created. On the other end, intersecting hit wires and reconstructing a *seed* helps to determine the delta electron segment position in 3D.

Concerning proton flagging, we could not compare the two algorithms since the *FlagBkgHits*, before creating *StereoHits*, applies a preliminary selection cut on the energy deposited in the *StrawHits* of 5 keV. *FlagBkgHits* then identifies protons as all particles with a deposited energy greater than 4.5 keV. It was not possible to

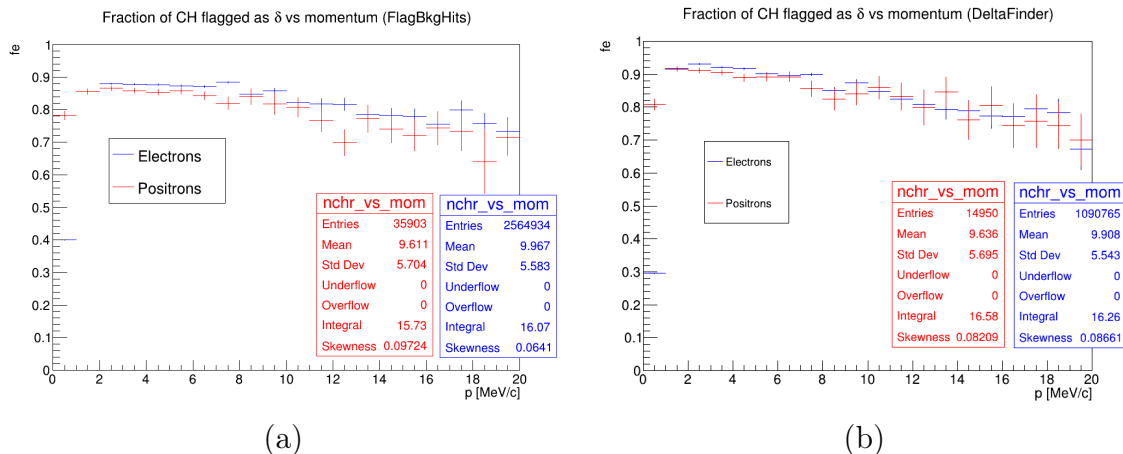


Figure 6.11: The electrons (blue) and positrons (red) hit flagging efficiency versus particle momentum in the low-momentum range. The left plot shows the results for *FlagBkgHits*, while the right plot shows those for *DeltaFinder*.

perform a direct comparison, as *DeltaFinder* does not apply any cuts on the energy deposited in the *StrawHits*, making the comparison unreliable and biased. Therefore, we reported the fraction of protons flagged as protons only by *DeltaFinder*.

The results for the 1BB and 2BB data samples are comparable.

Concerning μ and π (Table 6.5), we observe an approximate factor of 4 difference between the probabilities of hit mis-ID by the two algorithms for muons, and about a factor of 3.3 for pions. This occurs because muons and pions often produce more than one hit in a single station. However, the *DeltaFinder* algorithm can distinguish these hits as it is based on identifying δ -electron hit patterns. The main problem of *FlagBkgHits* is that it is based on a supervised training. It can distinguish one type of particle (CEs) from another (δ -electrons) but could get confused when the data include particles not used in the training, such as those from $p\bar{p}$ annihilation.

On average, pions produced from $p\bar{p}$ annihilations have higher momenta than muons, resulting in a higher false positive rate. This occurs because the curvature of the pion tracks is smaller than that of the muon tracks, making them more likely to form close in ϕ segments, thus appearing more similar to δ -electron.

	f_p DF	f_e FBH	f_e DF
μ	2.7%	13.0%	3.2%
π	0.4%	23.8%	7.3%

Table 6.5: Muon and Pion hit-level comparison (*PBAR* – *0BB* data sample). FBH and DF denote the *FlagBkgHits* and *DeltaFinder* algorithms respectively. f_p and f_e represent the number of particle *ComboHits* flagged as proton and electron hits, respectively, divided by the total number of *ComboHits*.

6.5.2 High-level comparison

The unflagged hits are then passed to the time clustering algorithm followed by the pattern recognition algorithm, after which the fit of the found track candidates is performed. Figure 6.12a shows the momentum distribution of reconstructed tracks for the $CE - 1BB$ data sample using both algorithms. In the case of *FlagBkgHits* (blue), less proton hits are correctly flagged, and non-flagged hits are used in the pattern recognition. On the other end, *DeltaFinder* (red) successfully flags proton hits, preventing them from being used by the pattern recognition. Figure 6.12b provides a zoomed-in view of the reconstructed track momentum distribution within the CE momentum range, where the results of track reconstruction for the two flagging algorithms are almost identical.




Figure 6.12: The momentum distribution of reconstructed tracks. Tracks reconstructed using *FlagBkgHits* (*DeltaFinder*) as the δ -flagger are shown in blue (red). The distribution in (a) covers a higher momentum range, while (b) provides a zoomed-in view of the momentum range [100, 107] MeV/c.

There is an excess of reconstructed tracks in the region of momentum above 120 MeV/c for *FlagBkgHits*. The excess is due to unflagged proton and deuteron hits used by the pattern recognition.

Table 6.6 reports the fraction of the simulated CE events that have at least one reconstructed track (for both $CE - 2BB$ and $CE - 1BB$ data samples). This fraction of CEs is nearly identical between the two algorithms, as expected from Figure 6.12. Table 6.4 shows that the difference in the number of hits between the two algorithms is less than one hit per track, which justifies the observed similarities in Table 6.6. For both 1BB and 2BB datasets, the relative difference between the CE reconstruction efficiencies corresponding to two algorithms is less than 1%.

To better understand the performance of the two algorithms, we analyzed events where at least one track was reconstructed using one hit flagger while no track was

	FBH 2BB	DF 2BB	FBH 1BB	DF 1BB
fraction of CE events with $N_{tracks} > 0$	36.7%	36.5%	37.9%	37.9%

Table 6.6: Fraction of reconstructed CE events ($CE - 2BB$ and $CE - 1BB$ data sample). FBH and DF denote *FlagBkgHits* and *DeltaFinder* algorithms, respectively.

reconstructed with the other hit flagger. A well defined class of events was identified in which the effects of hit flaggers are mitigated by the time clustering algorithm during the reconstruction.

Figures 6.13 (*FlagBkgHits*) and 6.14 (*DeltaFinder*) show the *time* versus *z* coordinate (the one aligned with the tracker axis) for the same event processed by both algorithms. CE hits are represented by large red dots, δ -electrons by small brown dots, and protons by large blue dots. Violet rectangles denote particles associated with the same *TimeCluster*, which is formed by grouping hits that align along a linear path in time versus *z* space. The time clustering process begins by




Figure 6.13: *TimeCluster* (green) on *time* – *z* plane. The *TimeCluster* is reconstructed using the *FlagBkgHits* output. CE hits are represented by large red dots, δ -electrons by small brown dots, and protons by large blue dots. Violet squares denote particles associated with the same *TimeCluster*.

combining hits within a specific *time* – *z* window to create *chunks* (with a 20 ns time window and a 5-plane *z*-window). A window qualifies as a *chunk* if it contains more than three *ComboHits*. Every potential pair of chunks within a certain time proximity is tested together, and the pair, that minimizes the $\chi^2/ndof$ when the hits are fit to a linear line, is combined. This procedure is repeated until no further combinations yield a $\chi^2/ndof$ below a set threshold. If a chunk exceeds a minimum number of straw hits, it is saved as a cluster.




Figure 6.14: *TimeCluster* (green and orange) on $time - z$ plane. The *TimeCluster* is reconstructed using the *DeltaFinder* output. CE hits are the large red dots, δ -electrons the small brown ones, and protons the large blue ones. Violet squares denote particles associated with the same *TimeCluster*.

Since particles grouped together may be at different z values, *TimeClusters* can exhibit varying z coordinates. In Figure 6.13, all hits are grouped into a single cluster (green), while in Figure 6.14 (green and orange), CE hits are divided into two different *TimeClusters*. Specifically, when the *DeltaFinder* algorithm is used, hits from the same particle are split into separate time clusters and no track are reconstructed. In contrast, with *FlagBkgHits*, unflagged hits are utilized by the time clusterer to *connect* particle hits, leading to successful track reconstruction.

This example highlights the importance of refining the time cluster finder to improve the track reconstruction.

Table 6.7 shows the fraction of reconstructed $p\bar{p}$ events, with at least two reconstructed tracks with the reconstructed momenta above 80 MeV/c or 90 MeV/c. The 80 MeV/c cut approximately corresponds to the minimum particle momentum for which a track of a particle coming from the ST can be reconstructed in the Mu2e tracker. The 90 MeV/c cut is applied to ensure the suppression of DIO events to a negligible level. For Run I, requiring a DIO electron momentum above 90 MeV/c yields an estimate of approximately 10^{-2} events with two DIO electrons. With a track reconstruction efficiency of approximately 0.1, this corresponds to roughly 10^{-4} events with two reconstructed DIO electron tracks. Assuming a uniform distribution in time, the number of such events within a 100 ns window is approximately 10^{-5} .

The *DeltaFinder* algorithm shows an advantage of approximately 22% in reconstructing two tracks, attributed to its ability to distinguish between muon and pion hits compared to δ -electrons (Table 6.5). The difference in hit-level efficiencies

results in more than 12 hits per track difference between the two algorithms. After applying momentum cuts, the difference between the algorithms appears to be similar to that observed without a momentum cut.

	FBH	DF
fraction of events with $N_{tracks} \geq 2$	1.8%	2.2%
fraction of events with $N_{tracks} \geq 2$ & $p > 80$ MeV/c	1.7%	2.1%
fraction of events with $N_{tracks} \geq 2$ & $p > 90$ MeV/c	1.6%	2.0%

Table 6.7: Fraction of reconstructed $p\bar{p}$ events ($PBAR - 0BB$ data sample). FBH and DF denote the *FlagBkgHits* and the *DeltaFinder* algorithm, respectively.

6.6 Conclusions

The simulation shows that the majority of hits in the tracker is generated by δ -electrons, i.e. electrons and positrons with momenta below 20 MeV/c. Two distinct algorithms have been developed in the Mu2e Offline to identify those hits and exclude them from the pattern recognition: *FlagBkgHits* and *DeltaFinder*. The *FlagBkgHits* algorithm initially clusters the hits and then uses an ANN to classify them, while *DeltaFinder* reconstructs only clusters of hits consistent with those produced by low-momentum particles or protons and deuterons.

I performed a systematic two-level comparison of the performance of the two algorithms:

- hit-level comparison: the focus is on evaluating the accuracy with which individual hits are flagged, providing a direct method for comparing the algorithms' performance. In terms of δ -electron flagging, both algorithms perform similarly in the considered energy range. Electrons (below 1-2 MeV/c) often produce only one or two hits, making them harder to flag, which reduces the electron hit flagging efficiency. The *FlagBkgHits* algorithm works better for electrons below 2 MeV, since the minimum number of hits to form a proper *seed* for *DeltaFinder* is three. Since positrons are not present in this energy range (as they are not produced by Compton scattering), the efficiency of positron hit flagging is higher for *DeltaFinder*. At higher momentum, both positron and electron flagging efficiency become similar, but overall efficiency decreases as momentum increases due to a wider spread of hits. *FlagBkgHits* flags approximately 70% more CE hits than *DeltaFinder*. The main difference between the two algorithms results from the primitives: *StereoHits* and *seeds*. When a δ -electron hits three straws, the *StereoHit* creates multiple intersection points, while the intersection of hit wires and the consequent reconstruction of a *seed* allows to better determine the delta electron segment

position in 3D. Furthermore, *DeltaFinder* can flag proton hits (approximately 84%), while *FlagBkgHits* can only flag hits based on large energy deposits, but it does not specifically identify proton candidates. To correctly estimate the background, it is crucial to not misidentify product hits of $p\bar{p}$ annihilation as δ -electrons and protons. This necessitates examining muon and pion hit flagging. *FlagBkgHits* has been trained on datasets lacking muons and pions and flags these particles at rates roughly 4 and 3.3 times higher, respectively, than *DeltaFinder*;

- high-level comparison: this is a study of the algorithms' impact on subsequent stages of event reconstruction and the track reconstruction efficiency. The main observed difference between the two algorithms is an excess of tracks in the momentum range above 120 MeV/c for *FlagBkgHits*. This increase is due to *FlagBkgHits* failing to properly flag proton and deuteron hits, allowing these particles to be sent to the reconstruction stage. For both 1BB and 2BB datasets, the relative difference between the CE reconstruction efficiencies corresponding to two algorithms is less than 1%. Moreover, the *DeltaFinder* algorithm demonstrates a significant advantage, showing approximately a 22% higher fraction of \bar{p} events with at least two reconstructed tracks compared to the alternative method.

A comparison of the timing performance of the two algorithms has also been performed. For *FlagBkgHits*, it was necessary to account for the time spent creating the *StereoHits*, while *DeltaFinder* independently reconstructed the *seeds*. Processing $CE - 1BB$ sample required about 0.14 ms/event for *FlagBkgHits* plus an additional 0.15 ms/event, compared to 0.42 ms/event for *DeltaFinder*. For $CE - 2BB$ events, *FlagBkgHits* took approximately 0.37 ms/event plus 0.34 ms/event, while *DeltaFinder* required 1.1 ms/event. The Mu2e specification states that the processing time per event should not exceed 5 ms. The difference in processing time is approximately 0.13 ms/event with 1BB pileup and 0.39 ms/event with 2BB pileup. While this difference does not appear critical at the moment, *DeltaFinder* needs to improve its timing performance.

The primary drawback of *FlagBkgHits* is its reliance on the supervised ANN training using CE and δ -electron samples. The problem with this method is when other particles, such as cosmic muons and those from $p\bar{p}$ annihilation, are introduced into the algorithm. In principle, additional samples could be incorporated into the training process, but these would increase the technical complexity. Additionally, the training was performed using Monte Carlo data rather than real data, posing a potential risk when transitioning to actual data taking.

Chapter 7

Conclusions

The success of Mu2e depends on many factors, one of which is the performance of the tracker. The tracker must provide an excellent momentum resolution, approximately 1 MeV/c FWHM, in order to distinguish the monochromatic CE signal from the background. The Mu2e collaboration has chosen a detector based on the straw tube technology. Its distinctive annular geometry is highly effective in reducing the background. The tracker is placed inside the DS, downstream from the ST, in a uniform 1 T magnetic field. The tracker has a modular design and consists of 18 tracking stations.

This Thesis provides a comprehensive contribution to the Mu2e tracker development, covering from its initial stages of commissioning to the optimization of DAQ and FEE and preliminary calibration steps. My work at Fermilab focused on thorough testing of the DAQ system from both hardware and software perspectives, and included involvement in the Vertical Slice Test (VST) of the tracker. The VST encompasses the entire testing chain, from the straws to the readout, to processed data on disk. Throughout the chapters, we explored critical aspects of the tracker performance, focusing on the robustness and reliability of data acquisition processes and the crucial role of calibration in ensuring precise measurements. Part of my work was dedicated to the offline analysis, particularly to pre-pattern recognition studies, investigating techniques of flagging δ -electrons during data taking. I performed the very first systematic study comparing the performance of two algorithms to determine the best approach for data-taking.

7.1 Commissioning of the tracker DAQ and FEE

The commissioning of the tracker DAQ and FEE systems, described in Chapter 4, has been a crucial step in validating the functionality and performance of the tracker readout chain.

The first test I performed was to verify the correct functioning of the ROC buffering and to understand its logic. During these tests, a single ROC was connected to one DTC. Data were collected with digi-FPGAs pulsed by their internal pulsers,

with the ROC set in external mode. The ROC's readout logic can be emulated with a bit-level C++ simulation, and I contributed to the simulation code development. The Monte Carlo simulation and the experimental data were compared in two different modes: *underflow* and *overflow* readout mode.

The ROC has an internal buffer which can store up to 255 hits, and, depending on T_{gen} and T_{EW} , the total number of hits within the event window can be respectively less than or equal to 255. These configurations depend on whether or not the buffer is getting filled up. By studying the timing distributions of the signal delays between different channels, we were able to incorporate relative channel-to-channel delays into the simulation. The comparison between the data and the bit-level simulation demonstrated good agreement, validating the simulation accuracy in reproducing the expected ROC behavior in two modes. The agreement between the simulation and the experimental data was highly satisfactory, with the deviations at a level of 10^{-3} .

The second test focused on evaluating the performance of preamplifiers, particularly looking for dead channels, cross-talk between the channels, and unexpected pulse patterns. The first part of the test examined the channel occupancy, revealing dead channels and channels exhibiting a higher number of hits than expected. Cross-talk was observed in only the first channels, with an asymmetry in interference between odd and even channels. This phenomenon was attributed to the physical proximity of preamplifiers on the motherboard and the closeness of the first channels to each other. Although the cross-talk was not fully addressed at the time of writing, further investigations are ongoing. Subsequent waveform analysis helped explain abnormalities, particularly the inverted waveforms observed in some channels, which indicated the erroneous triggering of waveform on both the leading and trailing edges. The waveform analysis provided further insights into the charge and pulse height distributions, showing two or three peaks in their distribution. This was due to the ADCs and the test pulser having correlated clocks.

Overall, these tests provided valuable results that will inform future improvements to the preamplifier design and its integration into the readout chain. They also laid a foundation for the development of real-time diagnostic tools, along with the identification of key performance issues.

In summary, this commissioning phase successfully validated the core functionalities of the tracker DAQ and FEE systems, ensuring they are ready for the next stages of testing and integration. The insights gained, particularly in handling ROC overflow and timing synchronization, will be essential for optimizing the system's performance during full detector operation. Further work will focus on expanding the scope of testing with more complex configurations and refining the DAQ software for future runs.

7.2 First steps towards the station calibration

Chapter 5 describes the initial steps towards the tracker calibration. Our primary objective is to perform a time calibration of the first assembled station using cosmic muons, with the aim of achieving a resolution on the longitudinal hit position better than 4 cm.

This calibration requires determining the signal propagation velocity and channel-to-channel delays.

When a charged particle passes through the straw tube gas volume, the resulting ionization charge generates an electric signal along the anode wire, which propagates to both ends of the straw and reaches the front-end electronics, where TDCs measure the signal arrival times, t_1 and t_2 . The calibration will be performed by reconstructing the track position along the straw and correlating it with the arrival times.

Due to operational constraints, such as the gas distribution, fragility, and space, a calibration with a station oriented vertically was considered to be the best option. I focused on the reconstruction of x_{track} using only the information about whether a straw was crossed or not by a cosmic muon. The only way to deduce x_{track} from this "yes-no" information is to use the geometry of the station, which consists of a set of rotated panels arranged in four faces orthogonal to the detector axis. I examined potential biases and systematic errors that could arise from using the vertical orientation.

The first step in my analysis involved studying the selection criteria for muons crossing a station and then various hit distributions resulting from the selection. The requirement of having at least one hit per face resulted in a non-uniform illumination, with almost no hits in the central region of the panel.

To reconstruct the 3D trajectory of a cosmic muon moving along the straight line, at least two 3D points along its track are needed. The x and y coordinates of each point can be determined by the intersection of one pair of non-parallel straws. This analysis indicates that the longitudinal bias along the straw range spans approximately $[-6,6]$ cm. This indicates a significant systematic factor affecting the hit reconstruction. The 2D distribution of the longitudinal bias versus the true position shows four distinct spots along the longitudinal axis, each corresponding to overlap regions of different straws. The systematic shift in the reconstructed longitudinal position could be as large as ± 4 cm, with the magnitude varying along the wire. This suggests that the mean is not a reliable estimator of the bias under these conditions, making the calibration process more challenging. The main issue with this configuration is that the stereo reconstruction biases for tracks with positive and negative values of m_{yz} do not cancel out.

In addition to the results, we must consider the rate and the data-taking time. For the vertical orientation, the rate must be scaled by two additional factors compared to the horizontal one: one that accounts for the angular dependence of the flux, and the second one that considers the angle between the station and cosmics. However,

estimating the rate is complex, as these factors are not the main contributors to the difference. With the horizontal orientation, it's possible to reconstruct a particle track with fewer requirements, relying on most muons being vertical. In contrast, with the vertical orientation, the analysis must include all biases for every straw, increasing the data volume needed as more parameters are involved.

Therefore, the horizontal orientation of the station is worth considering despite the difficulties with connecting the gas supply lines. A new mechanical solution to use the horizontal orientation is currently under development to address these issues.

7.3 Pre-pattern recognition studies

Given the high data volume expected during Mu2e data taking, which could be as large as 7 PBytes per year, optimizing memory usage and minimizing the CPU time consumption are critical. One significant challenge is flagging δ -electron hits (i.e. electrons and positrons with momenta below 20 MeV/c), which are the primary source of hits in the tracker, without compromising the efficiency of reconstructing the CE tracks. These particles come from Compton scattering, electron-positron pairs, and delta rays. Compton scattered electrons are produced when photons from neutron capture interact with the detector material, and typically have energies of a few MeV. Pair production electrons and positrons are generated during nuclear processes, and delta rays are produced when high-energy charged particles collide with the detector material.

A detailed study of pre-pattern recognition and a comparison of the two δ -electron flagging algorithms is presented in Chapter 6. This study has strong motivations. First, the primary goal of Mu2e is to detect the CE signal: flagging even a small fraction of hits generated by the CE as δ -electrons reduces the CE track reconstruction efficiency. Second, misidentification of muons and pions as δ -electrons may result in erroneous background estimate. Another reason is that counting the number of protons will be a complementary procedure to the STM to determine the muon stopping rate.

Two algorithms have been developed in Mu2e collaboration to identify δ -electron hits: *FlagBkgHits* and *DeltaFinder*. The *FlagBkgHits* algorithm initially clusters the hits and then uses an ANN to classify them, while *DeltaFinder* reconstructs only clusters of hits consistent with those produced by low-momentum particles or protons and deuterons.

I conducted a systematic, *two – level* comparison of the performance of the two algorithms:

1. **hit-level comparison:** this study evaluates the accuracy of individual hit flagging, providing a direct method to compare the algorithms' performance. In terms of δ -electron flagging, the two algorithms have a similar performance. At low momentum values (below 1-2 MeV/c), positron hit flagging is more ef-

ficient because positrons are not produced by Compton scattering. Electrons often produce only one or two hits, making them harder to flag, which reduces electron flagging efficiency. The *FlagBkgHits* algorithm works better for electrons below 2 MeV, since the minimum number of hits to form a *seed* for *DeltaFinder* is three. At higher momentum, both positron and electron flagging efficiency become similar, but overall efficiency decreases as momentum increases due to a wider spread of hits. *FlagBkgHits* flags approximately 70% more CE hits than *DeltaFinder*. The main difference between the two algorithms results from the primitives: *StereoHits* and *seeds*. When a δ -electron hits three straws, the *StereoHit* creates multiple intersection points, while the intersection of hit wires and the consequent reconstruction of a *seed* allows to better determine the delta electron segment position in 3D. Furthermore, *DeltaFinder* can flag proton hits (approximately 84%), while *FlagBkgHits* can only flag hits based on large energy deposits, but it does not specifically identify proton candidates. *FlagBkgHits* has been trained on datasets lacking muon and pion hits and thus flags these particles at rates roughly 4 and 3.3 times higher, respectively, than *DeltaFinder*;

2. **high-level comparison:** this study assesses the algorithms' impact on the track reconstruction efficiency. The main observed difference between the two algorithms is an excess of tracks in the right part of the *FlagBkgHits* momentum distribution of reconstructed tracks. This increase is due to *FlagBkgHits* leaving in enough proton hits for the pattern recognition to find the corresponding tracks. For both 1BB and 2BB datasets, the relative difference between the CE reconstruction efficiencies corresponding to two algorithms is less than 1%. Moreover, the *DeltaFinder* algorithm shows an advantage of approximately 22% in reconstructing \bar{p} two tracks.

An important aspect of the study was the timing performance of the two algorithms. The Mu2e specification states that the processing time should not exceed 5 ms/event. The difference in processing time is approximately 0.13 ms/event with 1BB pileup and 0.39 ms/event with 2BB pileup, with an advantage for *FlagBkgHits*. However, this difference does not appear to be critical.

The primary drawback of *FlagBkgHits* is its reliance on supervised training using CE and δ -electron samples. This approach fails when other particles, such as cosmic muons and those from $p\bar{p}$ annihilation, are introduced into the algorithm. Additional samples could be added to the training process, but these would increase the technical complexity.

Furthermore, the training was performed using the Monte Carlo data rather than the real data, posing a potential risk when transitioning to the data taking. The focus should be on how to handle training during early data taking and determining the strategy for incorporating real data into the training process.

Appendix A

Charged particle in a magnetic field

The motion of a charged particle in a magnetic field can be described by the Lorentz force:

$$\mathbf{F} = q \mathbf{v} \times \mathbf{B} \quad (\text{A.1})$$

where q is the particle charge, \mathbf{v} is the particle speed and \mathbf{B} is the magnetic field. A charged particle moving in a uniform solenoidal field describes the combination of a free trajectory and a circular motion, namely a helix, with the property:

$$|\mathbf{B}|\rho = \frac{p_{\perp}}{|q|} \quad (\text{A.2})$$

where ρ is the helix radius, p_{\perp} is the transverse momentum component. This is the simplest example, although more complex magnetic fields can produce a more complex trajectories. In this appendix, I will go over how to use a gradient to accelerate particles and also how particles perpendicularly drift in a curved magnetic field. The intensity of a magnetic field with a non-null gradient changes with the position. The force derived from the gradient is proportional to the particle magnetic momentum μ , where E represents energy. It can be written as:

$$\mathbf{F} = -\mu \nabla \mathbf{B} \quad \mu = \frac{c^2 p_{\perp}^2}{2EB} \quad (\text{A.3})$$

The force alters the direction of momentum, not the particle's energy. If the gradient is strong enough, it can reverse the direction of motion and reflect the particle like a magnetic mirror. Complex gradients can trap particles in certain regions or prevent them from staying there for too long. The other important feature is the use of curved magnetic fields. In a curved solenoid, the particle orbit drifts perpendicular to the bending plane. The drift velocity v_D and total displacement D can be determined from the path along the curved solenoid S :

$$v_D = \frac{m\gamma c}{eBR} (v_{\parallel}^2 + \frac{1}{2}v_{\perp}^2) \quad (\text{A.4})$$

$$D \propto pS\left(\frac{1}{\cos\theta} + \cos\theta\right) \quad (\text{A.5})$$

In the equations above, parallel and perpendicular refer to the magnetic field direction and R represents the solenoid bending radius. θ is the angle of the helix axis and the magnetic field and the sign of the drift relies on the sign of the charge. These characteristics show that a curved solenoid can be used to separate particle beams of opposite charge.

Appendix B

Simulation and Analysis tools

The Mu2e software framework has been designed to evaluate in advance the performance of the Mu2e detectors, assess signal reconstruction efficiency, and analyze the characteristics of background signals. This framework is built upon GEANT4, a widely recognized simulation toolkit that uses Monte Carlo methods for precise modeling.

Implemented in the C++ programming language, the framework encompasses a comprehensive suite of tools, including those for tracking, geometry configuration, and physics modeling. It provides detailed models of numerous physical processes, such as particle scattering, energy loss, and the decay of long-lived particles across a broad energy spectrum. The simulation environment faithfully reproduces the Mu2e geometry, and it integrates advanced pattern recognition and track reconstruction algorithms to enhance the accuracy and reliability of the simulations.

B.1 *art* and FHiCL

The Mu2e Offline software is structured around the *art* framework (Figure B.1), a versatile system developed and maintained by the Fermilab Scientific Computing Division (SCD). This framework is utilized by several Intensity Frontier experiments at Fermilab due to its robustness and adaptability.

art is an event-processing framework that is driven by command-line inputs and written in C++. It functions in a non-interactive mode, sequencing events as specified by the user. The framework is designed to address the extensive needs of high-energy physics experiments, including tasks such as high-level software triggering, online data monitoring, calibration, event reconstruction, simulation, and data analysis.

The configuration of *art* at runtime is managed using the Fermilab Hierarchical Configuration Language (FHiCL), a specialized data definition language created at Fermilab. FHiCL files define the C++ classes that implement *art* services. Within these classes, various algorithms - ranging from simulation and reconstruction to analysis codes - are developed and compiled into dynamic libraries known as mod-

ules. The FHiCL files specify which modules will be loaded, the order in which they will be executed, and the input and output files involved in the process.

The simulation process is initiated using a FHiCL file, which configures the entire process by specifying the required modules and referencing additional files containing essential geometry and physics data.

B.2 STNTUPLE and ROOT

STNTUPLE is a data format for n-tuples and a lightweight n-tuple analysis framework written primarily in C++. Originally developed for the CDF experiment at Fermilab, it has been adapted for use in the Mu2e experiment. One of the *art* plug-in modules is specifically designed for working with STNTUPLE [90]. Each STNTUPLE file is a ROOT file, which includes multiple branches corresponding to distinct data blocks. These blocks serve as optimized containers for I/O operations and analysis, storing Mu2e raw and/or reconstructed data.

Using the appropriate module within the *art-job* configuration file (with the *.fcl* extension), a STNTUPLE file can be generated from the data stored in an *art* file. The data type stored in this format is highly customizable. Once the *.stn* file is produced, analysis can proceed directly from this data format, thereby eliminating the need to rerun the reconstruction process. STNTUPLE is built on the ROOT package, a data analysis and graphics software developed at CERN, enabling users to leverage all of ROOT's interactive features for their analysis.

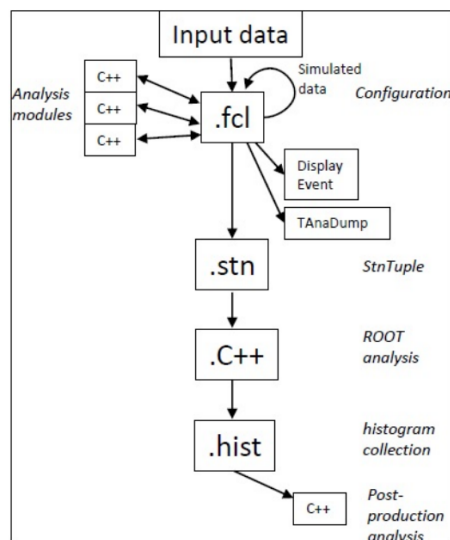


Figure B.1: Both the generation and reconstruction of events are performed through *art*-jobs configured using FHiCLs, which import the necessary C++ modules. Additionally, modules such as the Event Display and TAnaDump are utilized for debugging purposes. The analysis workflow relies heavily on the use of STNTUPLE.

Appendix C

Mu2e Offline event reconstruction

The primary goal of the Mu2e reconstruction algorithms is to efficiently reconstruct electrons within the specific range of conversion electrons. To meet this goal, the algorithms, along with certain user-defined parameters, are preconfigured with default values that are carefully optimized for this purpose. This appendix details the various stages involved in the Mu2e event reconstruction process.

C.1 Hits reconstruction and pre-filtering

C.1.1 Tracker hit reconstruction and pre-filtering

Charged particles traversing the tracker volume generate ionization charges within the gas enclosed by the straws. These charges are collected and produce electrical signals. Since the straws are read out by the front-end electronics from both ends, the initial step in the hit reconstruction process involves combining the two resulting electrical signals to estimate both the hit time and position along the wire. In the reconstruction code, this information is encapsulated in an object named *StrawHit*.

One of the most significant challenges in track reconstruction within the Mu2e experiment is the presence of numerous *StrawHit* objects within the $1.7\mu\text{s}$ time window that corresponds to a single event. The first critical task is, therefore, to identify *StrawHit* instances that are close in time and likely to have been generated by the same particle traversing the tracker. To enhance spatial resolution and minimize combinatorial complexities during track searching, adjacent *StrawHit* objects within a panel, which are most likely due to the same particle, are combined into a more complex object termed *ComboHit*. While retaining the information from individual *StrawHits*, the *ComboHit* provides the average time and position of the cluster.

This process is complicated by the presence of numerous hits produced by low-energy $p < 20 \text{ MeV}/c$ electrons, which are knocked out by Compton scattering, γ

conversion, or δ -rays. These electrons, commonly referred to as δ -electrons, follow trajectories with small radii and can generate a large number of hits, all confined within a limited volume, leading to regions of high occupancy. Fortunately, the hit patterns generated by δ -electrons are markedly different from those generated by particles within the energy range of interest for Mu2e, as further discussed in Chapter 6. These can be identified and filtered out using either the *FlagBkgHits* or *DeltaFinder* algorithms.

At this stage of the reconstruction process, the tracker's information has been translated into a collection of *ComboHit* objects. Each *ComboHit* is the result of combining a set of associated *StrawHits* and is characterized by its specific time and position. The clustering algorithm employed not only aids in this combination but also facilitates a partial reduction and filtering of background hits due to secondary electrons.

C.1.2 Calorimeter Hit Reconstruction

The logical equivalent of a *ComboHit* in the tracker is termed a *Cluster* in the calorimeter. A *Cluster* is formed by combining the signals generated by a group of crystals when a particle impacts the detector. The reconstruction of a *Cluster* begins by identifying the crystal with the highest energy deposition and subsequently adding all adjacent crystals that produce signals within a 2 ns window and have an energy above 50 MeV. The reconstruction process involves grouping adjacent crystals within a 2 ns time window, with an adjustable energy threshold, currently set at 50 MeV. This procedure is iteratively applied, starting from the added crystals, until no further crystals meet the inclusion criteria. Given the precision of the time measurements provided by the calorimeter, the timing of a *Cluster* can be utilized to define a temporal window within which all *ComboHits* generated by a particle passing through the tracker should be located. Thus, the calorimeter plays a crucial role in seeding the pattern recognition process, significantly reducing the combinatorial background in the tracker.

C.2 Helix search

The primary objective of the Mu2e tracking software is to reconstruct the trajectories of charged particles moving within the magnetic field of the DS. Ideally, these trajectories would form helices if no other effects were involved.

The helical paths can be described by the following set of parameters:

$$\vec{\eta} \equiv (d_0, \phi_0, \omega, z_0, \tan \lambda)$$

where d_0 represents the distance of the point of closest approach to the solenoid axis in the $x - y$ plane, with its sign determined by the particle's angular momentum relative to the origin. ϕ_0 denotes the direction of the momentum at the point of closest approach, while $\omega = 1/R$ is the curvature in the transverse plane. z_0 is the

z -coordinate at the point of closest approach, and $90^\circ - \lambda$ defines the pitch angle between the momentum vector \vec{p} and the xy plane.

The number of full rotations completed by a particle within the tracker volume is a function of its pitch and momentum. Most particles of interest in the Mu2e experiment complete more than one full rotation, implying that the actual trajectory for most tracks will be a long helix rather than just a segment of one. Due to the detector's design, particles may develop a portion of their trajectories within the bore, resulting in sequences of hits in the tracker that form multiple arcs.

The combination of *ComboHits* in the tracker, alongside the potential simultaneous presence of *Clusters* in the calorimeter, forms the foundational data required to reconstruct the helical trajectories. This search is executed in two successive phases, known respectively as *TimeClustering* and *Pattern – Recognition*.

C.2.1 Time Clustering

Given that the duration of a Mu2e event is several orders of magnitude longer than the time it takes for a particle to traverse the tracker, the first essential task is to identify which *ComboHits* could have originated from the same particle. This can be achieved by analyzing the time distribution of the *ComboHits*.

The procedure can be logically divided into two main steps:

- **Analyzing the *ComboHits* Time Distribution:** *ComboHits* produced by the same particle typically cluster into peaks within the time distribution. At each peak, a new entity known as a *TimeCluster* is created, and the collection of *ComboHits* associated with that peak is assigned to this *TimeCluster*. To enhance the accuracy of associating *ComboHits* with *TimeClusters*, the time distribution is adjusted by propagating all hit times to the central plane of the tracker ($z = 0$). This propagation is performed by assuming values for β and the angular velocity λ , which depend on the hypothesized particle identity. Similar to how *ComboHits* are formed from *StrawHits*, the *TimeCluster* consists of a list of *ComboHits*, but it is also associated with a time and position, both estimated from the *ComboHits*;
- **Refining the *ComboHits* Collection:** The time and position of the *TimeCluster* are subsequently used to further refine the collection of associated *ComboHits*. Several criteria are applied to the *ComboHits* linked with each *TimeCluster*, such as a maximum angular distance in the transverse $x - y$ plane. Following this refinement, the list of *ComboHits* associated with the *TimeCluster* may undergo minor adjustments. At this stage, the time and position of the *TimeCluster* are re-evaluated, and an additional iteration may include *ComboHits* that now meet the selection criteria. This iterative process continues until the list of *ComboHits* associated with the *TimeCluster* stabilizes, meaning no more *ComboHits* are added or removed.

It is important to consider that the avalanche processes occurring within the straws have a finite velocity, and these avalanches are initiated at random distances from the wires. Given that the radius of a straw is 2.5 mm, the standard deviation of the uniform distribution for the distance between the avalanche's starting point and the wire can be approximately estimated through a back-of-the-envelope calculation as $2.5 \text{ mm}/\sqrt{12} \sim 700 \text{ }\mu\text{m}$. Assuming a drift velocity of $50 \text{ }\mu\text{m}/\text{ns}$, this leads to a width estimate of approximately 14 ns. On the other hand, the hit times are propagated based on an assumed particle identity (specified by β and pitch), meaning that *TimeClusters* generated by different particles (with varying β) will have small differences and a similar spread in time.

The entire procedure is slightly modified if an energy cluster is detected in the calorimeter. In this case, the cluster can define the time window and provide a rough estimate of the *TimeCluster*'s $x - y$ position. Once this procedure is complete, all *TimeClusters* containing more than a programmable number of hits are stored.

C.2.2 Pattern Recognition

The next step involves searching for patterns within the list of *TimeClusters*. The current Mu2e code implements two primary pattern recognition algorithms. Pattern recognition in Mu2e is categorized into two main approaches: Tracker-only Pattern Recognition and Calorimeter + Tracker Pattern Recognition.

C.2.2.1 TrkPatRec: Tracker-only Pattern Recognition

The Tracker-only Pattern Recognition algorithm, referred to as *TrkPatRec* in Mu2e terminology, follows a two-step process. Initially, *TrkPatRec* analyzes the $x-y$ plane to find the track projection on the transverse plane. This analysis determines the track's radius (which is correlated with the transverse momentum) and the impact parameter with respect to the Stopping Target. Subsequently, the reconstruction in the $\phi-z$ plane is carried out, where the 2π ambiguity is resolved, and the track pitch is determined.

C.2.2.1.1 Reconstruction in the $x-y$ plane To identify the optimal circle that fits the hit distribution, a loop is executed over all possible triplets of *ComboHits* that belong to the same *TimeCluster*. For each triplet, if it spans a sufficient area, the (x, y) coordinates of the intersection of the two perpendicular bisectors are recorded. The median operator is then employed to combine the results from all triplets, yielding a point that represents a stable approximation of the helix center. Once the circle's center is established, a second loop determines the radial distance of the *ComboHits* from the helix axis, providing information about the track's radius. A pictorial representation of this procedure is shown in Figure C.1.

C.2.2.1.2 Reconstruction on $\phi - z$ plane To estimate the pitch of the track, the first step is to resolve the 2π ambiguity associated with the angular position of




Figure C.1: A pictorial representation of the procedure used to search for the center of the helix's x - y projection using triplets of *ComboHits* is shown. If a triplet covers a sufficient area, the position of the intersection of the bisectors is recorded. The median of these points provides an estimate of the helix axis.

the hits. Specifically, the angular positions ϕ of hits generated in the n -th loop of the track must be shifted by $2\pi n$. To make this correction, the angular velocity $\frac{d\phi}{dz} = \frac{1}{\lambda}$ of the particle is required, and therefore the first necessary step is to estimate $\frac{1}{\lambda}$.

A histogram is created using the variable $\lambda_{i,j;k}$, defined as:

$$\frac{1}{\lambda_{i,j;k}} = \frac{(\phi_j + 2\pi k) - \phi_i}{z_j - z_i} \quad (\text{C.1})$$

where i and j denote two different hits and range from 0 to $NCH - 1$, while k accounts for the number of full rotations, ranging from 0 to 10. The peaks in the resulting distribution are used to assign hits to the corresponding k -th loop, thereby resolving the ambiguity.

Figure C.2 illustrates how resolving the ambiguity affects the position of the hits in the $\phi - z$ plane. Once the ambiguity is resolved, it is possible to generate the histogram for $\frac{1}{\lambda_{i,j}} = \frac{\phi_j - \phi_i}{z_j - z_i}$, where the peak provides the best estimate of the helix's $\frac{d\phi}{dz}$.




Figure C.2: Sketch illustrating the resolution of the 2π ambiguity. Assigning the hits to the correct loop enables the determination of the track's angular velocity.

C.2.2.2 CalPatRec: Calorimeter+Tracker Pattern Recognition

The *CalPatRec* algorithm utilizes *CaloClusters* as initial *seeds* for pattern recognition. If a *CaloCluster* with reconstructed energy exceeding 50 MeV is present, its time and position are used to filter the *ComboHit* collection. The selected hits must fall within a ± 40 ns window relative to the calorimeter cluster's time and must lie within the same semi-plane (Figure C.3). This region is defined by first determining the angular position of the *Cluster* with respect to the beam axis. The tracker is then divided into two halves using a plane perpendicular to the *Cluster*'s position vector that passes through the beam axis. The half that contains the *Cluster* is retained.

Instead of using triplets of hits, the *CalPatRec* algorithm begins with the calorimeter cluster position, one *ComboHit*, and the solenoid center as initial points. A loop over the *ComboHits* allows the algorithm to flag hits that are sufficiently close to the helix projection. At this stage, the solenoid center can be dropped as a fixed position, and different *ComboHits* are iteratively used to adjust the helix parameters. The parameters are updated using two separate reduced- χ^2 fits for the x - y and ϕ - z planes. A critical step in this process is the accurate projection of hit uncertainties, accounting for the orientation of the straws relative to the helix.

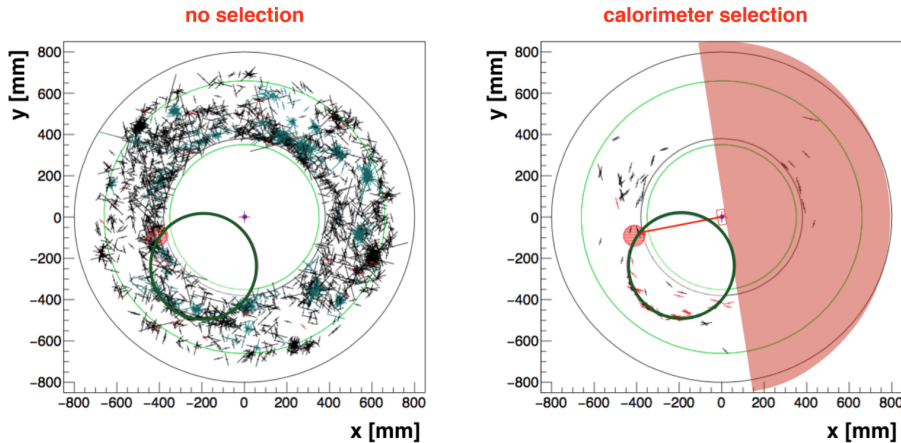


Figure C.3: Combinatorial background reduction achieved by exploiting the calorimeter clusters seeding. (Left plot) Typical Mu2e event with a conversion electron projected on the $x - y$ plane. The green circle represents the transverse projection of the conversion electron trajectory and the black crosses are *StrawHits* (the long arm indicates the direction of the straw); (Right plot) Same event after applying the calorimeter seeding.

C.3 Kalman Fitting

After the pattern recognition algorithms have been executed, an initial estimate of the track parameters $\vec{\eta}$ becomes available. At this stage, there are still numerous

effects that must be accounted for to optimize track reconstruction. Some of these effects are straightforward, such as the non-uniformity of the magnetic field, while others are subtler. For instance, there is an intrinsic symmetry in a straw hit regarding which side the particle traversed, leading to what is commonly referred to as left-right ambiguity, as illustrated in Figure C.4.




Figure C.4: The symmetry of the straw generates an ambiguity for the hits.

To address these effects that influence particle trajectories, the well-established Kalman fitting algorithm is employed. The Mu2e experiment utilizes a fitter developed for the BaBar experiment. The standard fitting process begins with a simplified Kalman fit, known as KSF in the Mu2e offline framework. Although this initial fit does not account for all effects, such as particle interactions with detector material, it improves the accuracy of track parameter reconstruction.

For more comprehensive effect corrections, a second Kalman fit (referred to as KK in Mu2e Offline) can be applied to account for residual effects. In this approach, the parameter vector $\vec{\eta}$ and the position along the beam axis (z) are used for track optimization, denoted as $F(\vec{\eta}; z)$. This fitting process provides the optimal estimate of $\vec{\eta}$, along with the corresponding 5×5 covariance matrix V . The complexity of this process increases when the parameter vector depends on the running variable, $\vec{\eta}(z)$, which is the case in the Mu2e experiment.

In the momentum range of interest, around $p \sim 100$ MeV/ c , different charged particle species, including electrons, muons, and protons, exhibit distinct behaviors within the Mu2e tracker. Electrons are highly relativistic with $\beta_e = v_e/c \sim 1$, while muons are significantly slower with $\beta_\mu \sim 0.7$. Protons at 100 MeV/ c are deeply non-relativistic. The energy loss characteristics also differ among these particles. Electrons and muons experience energy losses on the order of 1-2 MeV in the tracker, which is small compared to their total energy. Conversely, protons primarily lose all their energy due to ionization in the tracker. To account for these diverse cases, the offline track reconstruction involves multiple passes, each assuming specific hypotheses about particle mass and propagation direction.

Appendix D

Panels illumination

The illumination of cosmic hits on all twelve panels in the first station, as described in Section 5.3. The patterns are similar across all panels, with differences arising from the varying rotations of the panels relative to the station's center.

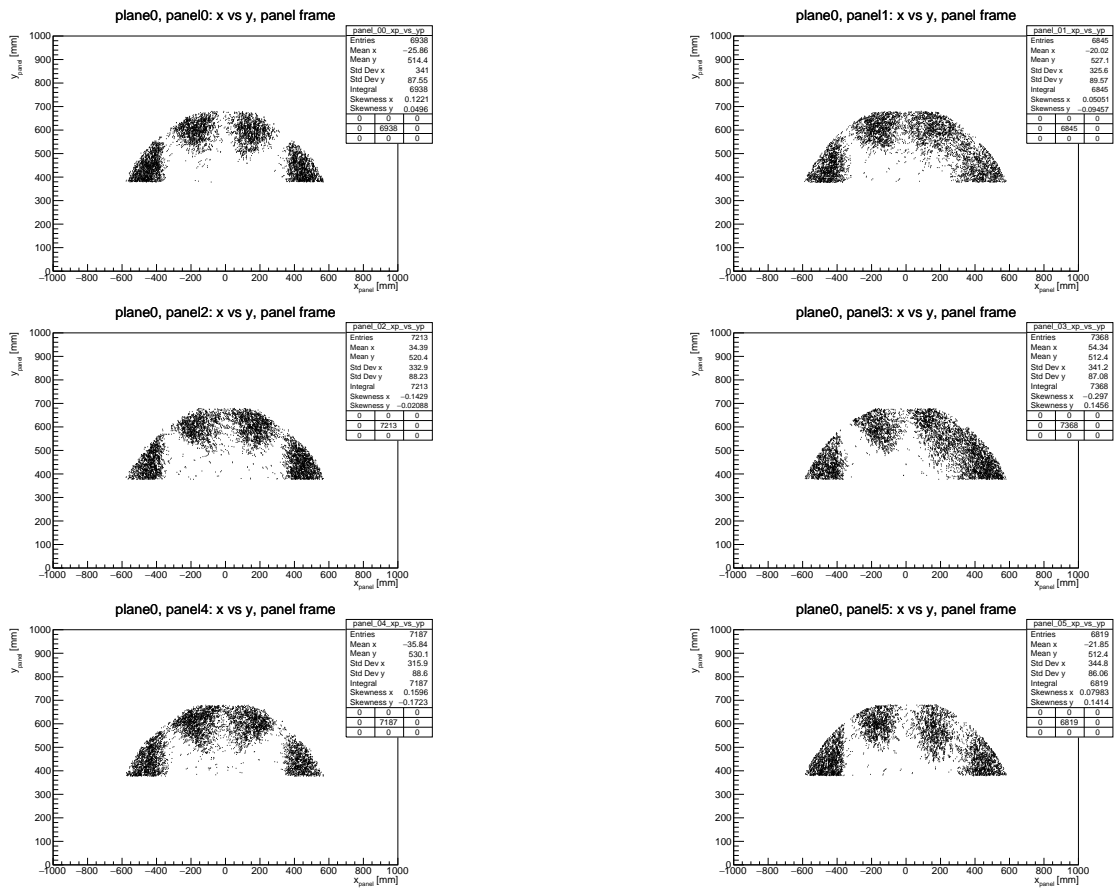


Figure D.1: Illumination pattern of cosmic hits on the panels of plane 0.

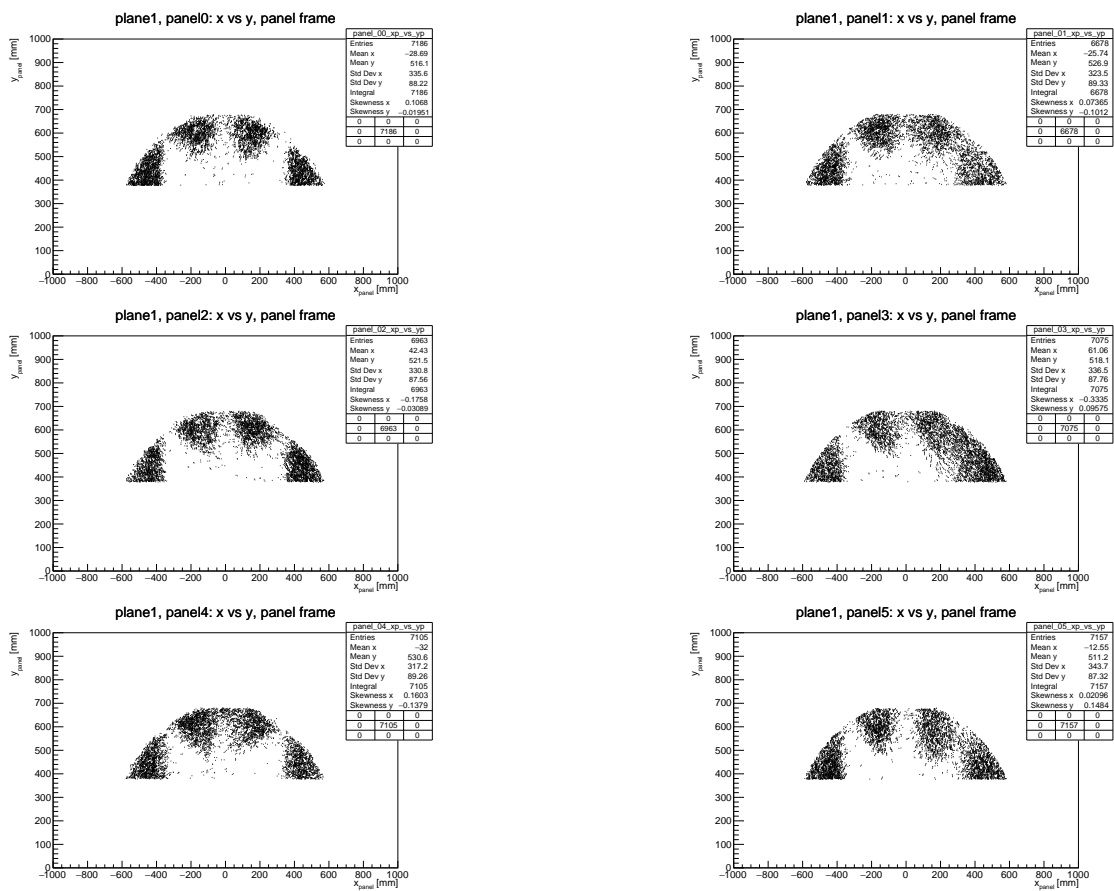


Figure D.2: Illumination pattern of cosmic hits on the panels of plane 1.

Bibliography

- [1] Mu2e Collaboration. Mu2e Technical Design Report, 2015.
- [2] Wilhelm H. Bertl et al. A Search for muon to electron conversion in muonic gold. *Eur. Phys. J. C*, 47:337–346, 2006.
- [3] Bernstein, Robert H. The Mu2e Experiment. *Frontiers in Physics*, 7, 2019. ISSN 2296-424X. URL <https://www.frontiersin.org/articles/10.3389/fphy.2019.00001>.
- [4] Bernstein, R.H. and Cooper, Peter S. Charged lepton flavor violation: An experimenter’s guide. *Physics Reports*, 532(2):27–64, November 2013. ISSN 0370-1573. URL <http://dx.doi.org/10.1016/j.physrep.2013.07.002>.
- [5] L. Cabibbi, G. Signorelli. Charged Lepton Flavour Violation: An Experimental and Theoretical Introduction. *La Rivista del Nuovo Cimento*, 41(2):71–174, Jan 2018. ISSN 0393697X, 0393697X. URL <https://doi.org/10.1393/ncr/i2018-10144-0>.
- [6] Cabibbo, Nicola. Unitary Symmetry and Leptonic Decays. *Phys. Rev. Lett.*, 10:531–533, Jun 1963. URL <https://link.aps.org/doi/10.1103/PhysRevLett.10.531>.
- [7] Ardu, Marco and Pezzullo, Gianantonio. Introduction to Charged Lepton Flavor Violation. *Universe*, 8(6), 2022. ISSN 2218-1997. URL <https://www.mdpi.com/2218-1997/8/6/299>.
- [8] W.J. Marciano and A.I. Sanda. Exotic decays of the muon and heavy leptons in gauge theories. *Physics Letters B*, 67(3):303–305, 1977. ISSN 0370-2693.
- [9] Davis, Raymond and Harmer, Don S. and Hoffman, Kenneth C. Search for Neutrinos from the Sun. *Phys. Rev. Lett.*, 20:1205–1209, May 1968. URL <https://link.aps.org/doi/10.1103/PhysRevLett.20.1205>.
- [10] Weinberg, Steven. Baryon- and Lepton-Nonconserving Processes. *Phys. Rev. Lett.*, 43:1566–1570, Nov 1979. URL <https://link.aps.org/doi/10.1103/PhysRevLett.43.1566>.

- [11] Harnik, Roni and Kopp, Joachim and Zupan, Jure. Flavor violating Higgs decays. *Journal of High Energy Physics*, 2013(3), mar 2013. ISSN 1029-8479. URL [http://dx.doi.org/10.1007/JHEP03\(2013\)026](http://dx.doi.org/10.1007/JHEP03(2013)026).
- [12] Pati, Jogesh C. and Salam, Abdus. Lepton number as the fourth "color". *Phys. Rev. D*, 10:275–289, Jul 1974. URL <https://link.aps.org/doi/10.1103/PhysRevD.10.275>.
- [13] I. Doršner and S. Fajfer and A. Greljo and J.F. Kamenik and N. Košnik. Physics of leptoquarks in precision experiments and at particle colliders. *Physics Reports*, 641:1–68, 2016. ISSN 0370-1573. URL <https://www.sciencedirect.com/science/article/pii/S0370157316301314>. Physics of leptoquarks in precision experiments and at particle colliders.
- [14] Leike, A. The phenomenology of extra neutral gauge bosons. *Physics Reports*, 317(3-4):143–250, aug 1999. ISSN 0370-1573. URL [http://dx.doi.org/10.1016/S0370-1573\(98\)00133-1](http://dx.doi.org/10.1016/S0370-1573(98)00133-1).
- [15] Particle Data Group, Zyla et al. Review of Particle Physics. *Progress of Theoretical and Experimental Physics*, 2020(8):083C01, 08 2020. ISSN 2050-3911. URL <https://doi.org/10.1093/ptep/ptaa104>.
- [16] MEG II collaboration, A.M. Baldini et al. A search for $\mu^+ \rightarrow e^+ \gamma$ with the first dataset of the MEG II experiment, 2024.
- [17] Bellgardt, U. and others. Search for the Decay $\mu^+ \rightarrow e^+ e^+ e^-$. *Nucl. Phys. B*, 299:1–6, 1988.
- [18] P. Wintz. Proceedings of the First International Symposium on Lepton and Baryon Number Violation. *IOP Publishing*, page 534, 1998.
- [19] L. Willmann et al. New bounds from searching for muonium to anti-muonium conversion. *Phys. Rev. Lett.*, 82:49–52, 1999.
- [20] Aubert et al. Searches for Lepton Flavor Violation in the Decays $\tau^\pm \rightarrow e^\pm \gamma$ and $\tau^\pm \rightarrow \mu^\pm \gamma$. *Physical Review Letters*, 104(2), jan 2010. ISSN 1079-7114. URL <http://dx.doi.org/10.1103/PhysRevLett.104.021802>.
- [21] Hayasaka et al. Search for lepton-flavor-violating τ decays into three leptons with 719 million produced $\tau^+ \tau^-$ pairs. *Physics Letters B*, 687(2-3):139–143, apr 2010. ISSN 0370-2693. URL <http://dx.doi.org/10.1016/j.physletb.2010.03.037>.
- [22] R. Aaij et al. Search for the Lepton-Flavor-Violating Decays $B_s^0 \rightarrow e^\pm \mu^\mp$ and $B^0 \rightarrow e^\pm \mu^\mp$. *Phys. Rev. Lett.*, 111:141801, Sep 2013. URL <https://link.aps.org/doi/10.1103/PhysRevLett.111.141801>.

- [23] B. Auberta et al. Searches for the decays $B^0 \rightarrow l^\pm \tau^\mp$ and $B^+ \rightarrow l^+ \nu$ ($l = e, \mu$) using hadronic tag reconstruction. *Phys. Rev. D*, 77:091104, May 2008. URL <https://link.aps.org/doi/10.1103/PhysRevD.77.091104>.
- [24] D. Ambrose et al. New limit on muon and electron lepton number violation from $K_0(L) \rightarrow \mu^\pm e^\mp$ decay. *Phys. Rev. Lett.*, 81:5734–5737, 1998.
- [25] A. Sher et al. Improved upper limit on the decay $K^+ \rightarrow \pi^+ \mu^+ e^-$. *Phys. Rev. D*, 72:012005, Jul 2005. URL <https://link.aps.org/doi/10.1103/PhysRevD.72.012005>.
- [26] Abouzaid, E. and others. Search for lepton flavor violating decays of the neutral kaon. *Phys. Rev. Lett.*, 100:131803, 2008.
- [27] W. Love et al. Search for Lepton Flavor Violation in Upsilon Decays. *Physical Review Letters*, 101(20), nov 2008. ISSN 1079-7114. URL <http://dx.doi.org/10.1103/PhysRevLett.101.201601>.
- [28] G. Aad et al., ATLAS Collaboration. Search for the lepton flavor violating decay $Z \rightarrow e\mu$ with the ATLAS detector. *Physical Review D*, 90(7), oct 2014. ISSN 1550-2368. URL <http://dx.doi.org/10.1103/PhysRevD.90.072010>.
- [29] R. Akers et al., OPAL Collaboration. A Search for lepton flavor violating Z^0 decays. *Z. Phys. C* 67, 1995.
- [30] P. Abreu et al., DELPHI Collaboration. Search for lepton flavor number violating Z^0 decays. *Z. Phys. C* 73, 1997.
- [31] Search for lepton flavour violating decays of the Higgs boson to e and $e\mu$ in proton-proton collisions at $\sqrt{s} = 8$ TeV. 2016. URL <https://hdl.handle.net/1721.1/134602>.
- [32] Sirunyan A. et al. Search for lepton flavour violating decays of the Higgs boson to $\mu\tau$ and $e\tau$ in proton-proton collisions at $\sqrt{s} = 13$ TeV. page , 03 2018.
- [33] De Gouvêa, André and Saoulidou, Niki. Fermilab’s Intensity Frontier. *Annual Review of Nuclear and Particle Science*, 60(1):513–538, 2010. URL <https://doi.org/10.1146/annurev-nucl-100809-131949>.
- [34] Hewett, J L et al. Fundamental physics at the intensity frontier. Report of the workshop held December 2011 in Rockville, MD. 6 2012. URL <https://www.osti.gov/biblio/1042577>.
- [35] M. Aoki et al. A New Charged Lepton Flavor Violation Program at Fermilab. In *Snowmass 2021*, 3 2022.
- [36] Final results of the MEG experiment 2016. *Il Nuovo Cimento C*, 39(4):1–8, Dec 2016. ISSN 03905551, 03905551. URL <https://doi.org/10.1393/ncc/i2016-16325-7>.

- [37] Mu2e Collaboration. Mu2e Run I Sensitivity Projections for the Neutrinoless $\mu^- \rightarrow e^-$ Conversion Search in Aluminum. *Universe*, 9(1), 2023. ISSN 2218-1997. URL <https://www.mdpi.com/2218-1997/9/1/54>.
- [38] F Cei. Status and perspectives of lepton flavour violation experiments with muons. *Journal of Physics: Conference Series*, 1526(1):012020, apr 2020. URL <https://dx.doi.org/10.1088/1742-6596/1526/1/012020>.
- [39] MEG II Collaboration. Operation and performance of MEG II detector, 2024.
- [40] Eichler, Ralph and Grab, Christophorus. The SINDRUM-I Experiment. *SciPost Physics Proceedings*, page , 09 2021.
- [41] Gavin Hesketh and Sean Hughes and Ann-Kathrin Perrevoort and Nikolaos Rompotis. The Mu3e Experiment, 2022.
- [42] Papa, Angela. Towards a new generation of Charged Lepton Flavour Violation searches at the Paul Scherrer Institut: The MEG upgrade and the Mu3e experiment. *EPJ Web of Conferences*, 234:01011, 01 2020.
- [43] E. C. Dukes. Mu2e-II: The Mu2e Experiment in the PIP-II Era. 2021.
- [44] TESHIMA, Natsuki. DeeMe experiment to search for muon to electron conversion at J-PARC MLF. page 109, 02 2018.
- [45] Abramishvili R. et al. COMET Phase-I technical design report. *Progress of Theoretical and Experimental Physics*, 2020(3), 2020. ISSN 2050-3911. URL <http://dx.doi.org/10.1093/ptep/ptz125>.
- [46] The Belle detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 479(1):117–232, 2002. ISSN 0168-9002. URL <https://www.sciencedirect.com/science/article/pii/S0168900201020137>. Detectors for Asymmetric B-factories.
- [47] The LHCb Collaboration. 3(08):S08005, aug 2008. URL <https://dx.doi.org/10.1088/1748-0221/3/08/S08005>.
- [48] Branchini, Paolo. The Belle II Experiment: Status and Prospects. *Universe*, 4(10), 2018. ISSN 2218-1997. URL <https://www.mdpi.com/2218-1997/4/10/101>.
- [49] J Bennett. Belle II Physics Prospects, Status and Schedule. *Journal of Physics: Conference Series*, 770(1):012044, nov 2016. URL <https://dx.doi.org/10.1088/1742-6596/770/1/012044>.
- [50] Kargiantoulakis, Manolis. A search for charged lepton flavor violation in the Mu2e experiment. *Modern Physics Letters A*, 35(19):2030007, apr 2020. ISSN 1793-6632. URL <http://dx.doi.org/10.1142/S0217732320300074>.

- [51] R.M. Dzhilkibaev and V.M. Lobashev. On the search for $\mu \rightarrow e$ conversion process in a nucleus. pages 622–624, 1989.
- [52] D.F. Measday. The nuclear physics of muon capture. *Physics Reports*, 354(4):243–409, 2001. ISSN 0370-1573. URL <https://www.sciencedirect.com/science/article/pii/S0370157301000126>.
- [53] Czarnecki, Andrzej and Garcia i Tormo, Xavier and Marciano, William J. Muon decay in orbit: Spectrum of high-energy electrons. *Phys. Rev. D*, 84:013006, Jul 2011. URL <https://link.aps.org/doi/10.1103/PhysRevD.84.013006>.
- [54] R. Ehrlich. Cosmic Ray Backgrounds in the Mu2e Experiment at Fermilab. *DPF 2017 Poster 2017-07-31*, 2017. URL https://indico.fnal.gov/event/11999/contributions/11447/attachments/7562/9672/poster_DPF2017_5.pdf.
- [55] L Michel. Interaction between Four Half-Spin Particles and the Decay of the μ -Meson. *Proceedings of the Physical Society. Section A*, 63(5):514, may 1950. URL <https://dx.doi.org/10.1088/0370-1298/63/5/311>.
- [56] Kinoshita, Toichiro and Sirlin, Alberto. Radiative Corrections to Fermi Interactions. *Phys. Rev.*, 113:1652–1660, Mar 1959. URL <https://link.aps.org/doi/10.1103/PhysRev.113.1652>.
- [57] A. Gaponenko. Summary of backgrounds for CD3. (*Mu2e Backgrounds 2017-06-08*), *mu2e-doc-7464-v13*, 2017.
- [58] Bistirlich, J. A. and Crowe, K. M. and Parsons, A. S. L. and Skarek, P. and Truöl, P. Photon Spectra from Radiative Absorption of Pions in Nuclei. *Phys. Rev. C*, 5:1867–1883, Jun 1972. URL <https://link.aps.org/doi/10.1103/PhysRevC.5.1867>.
- [59] Stratakis, Diktys and Convery, Mary E. and Johnstone, Carol and Johnstone, John and Morgan, James P. and Still, Dean and Crnkovic, Jason D. and Tishchenko, Vladimir and Morse, William M. and Syphers, Michael J. Accelerator performance analysis of the Fermilab Muon Campus. *Phys. Rev. Accel. Beams*, 20:111003, Nov 2017. URL <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.20.111003>.
- [60] S. J. Werkema. CD-3c Review: Mu2e Accelerator Systems. (*DOE CD-3c Review 2016-06-14*), *mu2e-doc-7530-v6*, 2016.
- [61] E. Prebys and S.J. Werkema. Out of Time Beam Extinction in the Mu2e Experiment. In *Proc. 6th International Particle Accelerator Conference (IPAC'15), Richmond, VA, USA, May 3-8, 2015*, number 6 in International Particle Accelerator Conference, pages 3996–3999, Geneva, Switzerland, June 2015. JACoW. ISBN 978-3-95450-168-7.

- [62] D. R. Pushka. Production target (Hayman 2)-Requirement and Design. *Mu2e internal doc. Database, doc-db-30814-v2*, 2020.
- [63] G. Ginther. CD-3c Review: Muon Beamline Overview. (*DOE CD-3c Review 2016-06-14*), mu2e-doc-7531-v5.
- [64] Kashikhin V. V. et al. Conceptual Design of the Mu2e Production Solenoid. *IEEE Transactions on Applied Superconductivity*, 23(3):4100604–4100604, 2013.
- [65] Mu2e Collaboration. MuonBeamlineIntro. URL <https://mu2ewiki.fnal.gov/wiki/MuonBeamlineIntro>.
- [66] Cirigliano, Vincenzo and Kitano, Ryuichiro and Okada, Yasuhiro and Tuzon, Paula. Model discriminating power of $\mu \rightarrow e$ conversion in nuclei. *Phys. Rev. D*, 80:013002, Jul 2009. URL <https://link.aps.org/doi/10.1103/PhysRevD.80.013002>.
- [67] Kitano, Ryuichiro and Koike, Masafumi and Okada, Yasuhiro. Erratum: Detailed calculation of lepton flavor violating muon-electron conversion rate for various nuclei [Phys. Rev. D 66, 096002 (2002)]. *Phys. Rev. D*, 76:059902, Sep 2007. URL <https://link.aps.org/doi/10.1103/PhysRevD.76.059902>.
- [68] F. Abusalma et al. Expression of Interest for Evolution of the Mu2e Experiment, 2018.
- [69] T. Yamazaki et al. Negative Moun Spin Rotation. *Physica Scripta*, 11(3-4): 133, mar 1975. URL <https://dx.doi.org/10.1088/0031-8949/11/3-4/004>.
- [70] Nikolay Atanov et al. The Mu2e Crystal Calorimeter: An Overview. *Instruments*, 6(4):60, 2022.
- [71] N. Atanov et al. Towards the construction of the Mu2e electromagnetic calorimeter at Fermilab. *J. Phys. Conf. Ser.*, 2374(1):012021, 2022.
- [72] N. Atanov et al. The Mu2e Calorimeter Final Technical Design Report. 2 2018.
- [73] Nikolay Atanov et al. The Mu2e Crystal and SiPM Calorimeter: Construction Status. *IEEE Trans. Nucl. Sci.*, 70(7):1281–1287, 2023.
- [74] Luca Morescalchi et al. The Readout Electronics of the Mu2e Electromagnetic Calorimeter. *PoS*, PANIC2021:113, 2022.
- [75] S. Ceravolo et al. Design and qualification of the Mu2e electromagnetic calorimeter electronic system. *Nucl. Instrum. Meth. A*, 1047:167875, 2023.
- [76] Nikolay Atanov et al. Mu2e Crystal Calorimeter Readout Electronics: Design and Characterisation. *Instruments*, 6(4):68, 2022.

- [77] N. Atanov et al. The mu2e undoped csi crystal calorimeter. *Journal of Instrumentation*, 13(02):C02037, feb 2018. URL <https://dx.doi.org/10.1088/1748-0221/13/02/C02037>.
- [78] S. Giovannella. The detectors of the Mu2e experiment. *Journal of Instrumentation*, 15(06):C06022, jun 2020. URL <https://dx.doi.org/10.1088/1748-0221/15/06/C06022>.
- [79] H. Kolanoski, N. Wermes. *Particle detectors*. Oxford University Press, 2020.
- [80] T. Zhao, Y. Chen, S. Han, and J. Hersch. A study of electron drift velocity in ar:co2 and ar:co2:cf4 gas mixtures. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 340(3):485–490, 1994. ISSN 0168-9002. URL <https://www.sciencedirect.com/science/article/pii/0168900294901295>.
- [81] A. Mukherjee. Mu2e Tracker. In (*DOE IPR 2017 2017-07-25*), *mu2e-doc-11567-v3*, 2017.
- [82] V. L. Rusu. Mu2e Tracker Front End Electronics. *DOE IPR 2019 2019-12-10*, *mu2e-doc-30505-v3*, 2019.
- [83] Richie Bonventre, Greg Deuerling, Vadim L Rusu. Mu2e-doc-7849-v1. Technical report, 2016.
- [84] Antonio Gioiosa et al. Status of the data acquisition, trigger, and slow control systems of the Mu2e experiment at Fermilab. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1046:167732, 2023. ISSN 0168-9002. URL <https://www.sciencedirect.com/science/article/pii/S0168900222010245>.
- [85] R. Bernstein and A. Mukherjee, mu2e-doc-732-v22. Tracker Requirements Document. 2016.
- [86] Olive, K. A. et al. Charged lepton flavor violation: An experimenter’s guide. *Chin. Phys. C38*, Review of Particle Physics, 2014.
- [87] Mengyun Guan and Ming-Chung Chu and Jun Cao and Kam-Biu Luk and Changgen Yang. A parametrization of the cosmic-ray muon flux at sea-level, 2015.
- [88] Chris Hagmann, David Lange, and Doug Wright. Cosmic-ray shower generator (cry) for monte carlo transport codes. *2007 IEEE Nuclear Science Symposium Conference Record*, 2:1143–1146, 2007. URL <https://api.semanticscholar.org/CorpusID:24570080>.
- [89] D. Heck, J. Knapp, J. N. Capdevielle, G. Schatz, and T. Thouw. CORSIKA: A Monte Carlo code to simulate extensive air showers. 2 1998.
- [90] G. Pezzullo (Mu2e) P. Murat. Stntuple package. *Mu2e-doc-13916-v1*, 2017.