

Dissertation

submitted to the
Combined Faculty of Mathematics, Engineering
and Natural Sciences
of Heidelberg University, Germany
for the degree of
Doctor of Natural Sciences

Put forward by
MIHAIL-BOGDAN BLIDARU

Born in: Buzău, România
Oral examination: 14.11.2024

Bent monolithic active pixel sensors and material budget imaging studies towards the ALICE ITS3 upgrade

Referees: Prof. Dr. Silvia Masciocchi
Prof. Dr. André Schöning

Abstract

The challenging environment of high-energy heavy-ion collisions demands continuous advancements in precision tracking and vertex detection. To meet these requirements, the ALICE ITS3 upgrade at CERN is designed to push the boundaries of CMOS monolithic detector technology by integrating wafer-scale, ultra-thin, and bent silicon sensors into the ALICE experiment.

Within the scope of this thesis two critical aspects of this innovative detector design are addressed: the feasibility of bending CMOS monolithic pixel sensors, specifically the ALPIDE sensors used in the ITS2 detector, and the material budget associated with carbon foam support structures for future wafer-scale sensors in the ITS3.

Through a series of extensive test beam campaigns, it was demonstrated that ALPIDE sensors retain their full functionality and performance metrics after bending. The sensors showed unchanged efficiency and spatial resolution compared to their flat counterparts, regardless of the bending radius, confirming the viability of bent CMOS MAPS technology for future detector designs, especially in applications requiring low material budgets and high tracking precision, like the ITS3.

Beyond the performance evaluation of the sensors, the thesis explores the characterization of carbon foam support structures intended for the ITS3 detector. As the primary structural component in the active area, the carbon foam plays a crucial role in minimizing the material budget, thereby enhancing tracking precision and efficiency. Using material budget imaging techniques, this study evaluates the material contribution of carbon foam samples, validating their potential in the final detector assembly.

Overall, this thesis provides much needed insights into the electrical, mechanical and material properties of key components in the ITS3 upgrade, reinforcing their potential for application in particle detectors that prioritize low mass and proximity to the interaction vertex.

Zusammenfassung

Hochenergetische Schwerionenkollisionen bieten eine anspruchsvolle, experimentelle Umgebung, die kontinuierliche Fortschritte in der Spur- und Vertexrekonstruktion elektrisch geladener Teilchen erfordert. Durch die geplante Integration von hauchdünnen und gebogenen Siliziumsensoren für das ALICE ITS3-Upgrade am CERN, werden deshalb die Grenzen der monolithischen CMOS-Detektortechnologie neu ausgelotet.

Im Rahmen dieser Doktorarbeit werden zwei kritische Aspekte dieses innovativen Detektordesigns untersucht. Zum einen wird die Machbarkeit von gebogenen, monolithischen CMOS-Pixelsensoren, insbesondere der ALPIDE-Sensoren, die im ITS2-Detektor verwendet werden, gezeigt. Zum anderen wird eine Analyse von Kohlenstoffschaum-Trägerstrukturen für die gebogenen Sensoren im Hinblick auf die Materialeigenschaften und Verteilung durchgeführt.

Durch umfangreiche Messkampagnen an einem Teststrahl konnte gezeigt werden, dass die ALPIDE-Sensoren nach dem Biegen ihre volle Funktionalität und Leistungskennzahlen beibehalten. Im Vergleich zu flachen Pendants, wurde für gebogene Sensoren eine unveränderte Effizienz und Ortsauflösung, unabhängig vom Biegeradius nachgewiesen. Diese Ergebnisse bestätigen die Realisierbarkeit von gebogenen Sensoren der CMOS-MAPS-Technologie und deren Anwendung in künftigen Detektordesigns, insbesondere wenn ein geringes Materialbudget und eine präzise Spurrekonstruktion erforderlich ist.

Zusätzlich zur Leistungsmessung der Sensoren wird in dieser Arbeit auch die Charakterisierung der, für den ITS3-Detektor vorgesehenen, Trägerstrukturen aus Kohlenstoffschaum untersucht. Als primäre mechanische Halterung im aktiven Bereich spielt Kohlenstoffschaum eine entscheidende Rolle bei der Minimierung des Materialbudgets, wodurch die Präzision und Effizienz der Teilchendetektion verbessert wird. In dieser Studie wird der Beitrag von Kohlenstoffschaumproben zum Streuverhalten passierender Teilchen mit Hilfe von Bildgebungsverfahren bewertet und somit deren Eignung für die endgültige Montage im Detektor validiert.

Insgesamt liefert diese Arbeit dringend benötigte Erkenntnisse in die elektrischen, mechanischen und materiellen Eigenschaften von Schlüsselkomponenten des ITS3-Upgrades und unterstreicht deren Potenzial für die Anwendung in Teilchendetektoren, bei denen eine geringe Masse und die Nähe zum Wechselwirkungspunkt im Vordergrund stehen.

Table of contents

Abstract	iii
Acknowledgments	ix
1 Introduction	1
1.1 Particle physics and the Standard Model	2
1.1.1 The Standard Model of particle physics	2
1.1.2 The Large Hadron Collider	5
1.1.3 The ALICE experiment	7
1.2 Particle detection	9
1.2.1 Interaction of particles with matter	9
1.2.2 Silicon detectors	17
1.3 Pixel sensors	28
1.3.1 Hybrid vs monolithic design	29
1.3.2 High-Voltage (HV) and High-Resistivity (HR) CMOS MAPS	30
1.3.3 Modified process for HR CMOS MAPS	33
1.4 ALICE upgrades	35
1.4.1 ITS2	36
1.4.2 ALPIDE	39
1.4.3 ITS3	43
1.4.4 ALICE 3	52
1.5 Sensor characterization	54
1.5.1 Testbeams and particle telescopes	54
1.5.2 Track models	58
1.5.3 Alignment models	60
1.5.4 Parameters of interest	62
2 Bent sensors	65
2.1 Bending ALPIDE sensors	66
2.1.1 Bending CMOS sensors	66
2.1.2 Bending along the short axis	66

2.1.3	Bending along the long axis	69
2.2	Alignment of bent sensors in software	74
2.3	Data analysis of bent sensors at different radii	80
2.3.1	Sensors bent along the short edge	80
2.3.2	Sensors bent along the long edge	86
2.4	Inelastic interactions in the μ ITS3 with a target	106
2.4.1	Experimental setup	106
2.5	Discussion	120
3	Material budget imaging	125
3.1	The road to nearly massless detectors	126
3.1.1	Material budget changes from ITS2 to ITS3	126
3.1.2	Carbon-foam design	128
3.2	Measurement idea	132
3.2.1	Carbon foam samples	132
3.2.2	Procedure	139
3.2.3	Experimental setup	143
3.3	Image reconstruction	148
3.3.1	Kink angle reconstruction	148
3.3.2	Estimators of the width of the distribution	151
3.3.3	Quantiles	158
3.3.4	Position resolved scattering angle distribution	160
3.3.5	Momentum divergence considerations	162
3.3.6	Corrections and calibrations	165
3.3.7	Material budget calibration	175
3.3.8	Material budget measurement for the foam samples	178
3.4	Discussion	181
4	Short summary and outlook	185
	Appendix A Image cell size	189
	Appendix B Telescope setup configurations	195
B.1	Setup configuration for 1 GeV simulated data	196
B.2	Setup configuration for 2.4 GeV simulated data	198
B.3	Setup configuration for 5.4 GeV simulated data	200

Table of contents	vii
Appendix C Quadratic subtraction	203
Appendix D Quadratic subtraction for the foam samples	207
Appendix E Material budget calibration	213
Appendix F 3D-printed jigs cross-section	217
Appendix G Effects of the pixel pitch asymmetry on the kink angle distributions	219
List of publications	223

Acknowledgments

As I reach the completion of this thesis, it marks a key moment in my journey as a scientist. This milestone symbolizes the culmination of years of hard work, curiosity, and determination. Yet, this achievement is not merely the end of a long academic path but the beginning of a new chapter, filled with opportunities for further exploration and discovery. As I reflect on this accomplishment, I am fully aware that it was made possible not solely by my efforts but by the unwavering support, mentorship, and inspiration from so many others. This journey has taught me that while individual dedication is essential, true success is often built on the collective efforts and encouragement of others. The adventure, however, is far from complete. I recognize just how much there is still to learn, and I humbly look forward to the challenges and opportunities that lie ahead. The completion of this thesis is not an endpoint, but rather a stepping stone toward further inquiry and understanding. So, as I take a momentary pause, I reflect on the many people I wish to thank.

I have made every effort during these last few stressful days to recall and include all the names here. However, I deeply apologize to anyone who has been unintentionally overlooked due to my incomplete memory or oversight. Your contributions are no less valued, and you truly deserve to be mentioned and appreciated here. If this was the case, please don't hesitate to let me know... the next round of beers is on me!

First and foremost, I would like to express my deepest gratitude to my supervisor, mentor and friend, Prof. Dr. Silvia Masciocchi. I am so happy I met you and that you accepted me into your group! Since the beginning of my M.Sc., through the Ph.D., and most likely in future endeavors, your mentorship has been invaluable. You have shaped me into the scientist I am today, and I know that your influence will continue to guide me indirectly long after. Your unwavering support and insightful guidance have been the cornerstone of my academic journey, and for that, I am profoundly thankful. Your critical questions constantly pushed me to sharpen my thinking and brought greater depth to my work. It is because of your influence that I chose to continue down this path, and it is with you that I wish to pursue it further.

Secondly, I would like to thank Prof. Dr. André Schöning for his willingness to serve as the second referee for this thesis. I am deeply grateful to him for being my first point of contact when I joined Heidelberg as a fresh master's student and for introducing me to the world of pixel sensors within his group.

I would also like to extend my gratitude to Prof. Dr. Jan M. Pawlowski and Prof. Dr. Joao Seco for their willingness to be part of the examination committee. Additionally, I would like to thank Prof. Dr. Peter Fischer for serving as a co-adviser throughout my PhD project and for his numerous valuable suggestions during our regular meetings.

I also want to extend my heartfelt thanks to my dear friend and colleague, Pascal, who has shared in both the challenges and triumphs that hardware life brought along this journey. Without having somebody like you to rely on, ask questions, and seek guidance, I would not have managed to get here. We've both had great pleasure and an immense amount of fun engaging in all our common pixel activities, coordinating lab work, and pushing ourselves constantly. I truly miss our test beam adventures, but who knows what the future holds. Only you know the tough times we've faced, most of which we've shared, and it is together that we've managed to become the experts we are today. It always amazes me just how much physics knowledge you possess; you are a brilliant scientist and a true friend.

To Michael, my good friend and colleague, your companionship throughout this journey has been a constant source of encouragement. Working alongside you has made the tough moments more bearable, and your friendship is something I deeply cherish. God bless your coffee grinder and maker, as well as your marvelous coffee-making skills — they've been lifesavers! It's a real pity you didn't get to attend a test beam with Pascal and me. I'm already missing you from the office life and hope you'll drop by now and then, especially on Christmas Eve. Thank you for your tips on managing the final months of thesis writing; they were invaluable ;)

Some special thanks also go to Carolina for enduring my endless jokes and cries for help during office hours. Your patience and good humor have brightened many long days, and I truly appreciate the great times we've shared. I hope the office atmosphere wasn't too overwhelming, even though it could be challenging at times when all three of us were there together :). Your presence made it all the more enjoyable.

I would also like to thank the two Fedes, as well as Annalena for their support and collaboration. We share many great memories from Europa Park, the Neckarwiese, and all those fantastic group retreats — moments that made our time together even more enjoyable and memorable. Many thanks to Andrea for making the physics part somewhat bearable and for patiently listening to my constant complaints while always helping me push forward. Your support has been invaluable!

A special thanks to the entire hardware team (Bent, Maurice, Simon, Sarah, Johannes, Anouk, Fabian, David, Marius). Your technical expertise and problem-solving skills have been indispensable, and I am deeply appreciative of the time and effort you've invested in our

collective endeavors. I hope I was able to offer the right advice when needed and that you enjoyed working together as much as I did. We've all had our share of fun, and I'm so proud to see how each of you has become an expert in your own field. You are all truly amazing!

I would like to extend my heartfelt thanks to my colleagues at CERN and around the world in ALICE with whom I have collaborated (Lukas, Magnus, Miko, Felix, Matthew, Giacomo, Alexander Ferk, Alexander Kalweit, Jian, Paolo, Alice, Roberto, Jory, Serhiy, Shyam, Francesco, Gianfranco, Francesca, Manuel, Nicolo, Alperen) and to the people from the local ALICE group (Ralf, Ana, Dariusz, Kai, Denise, Maria, and many more). Your insights, feedback, and great spirit have greatly enriched my work and contributed to my growth as a researcher. The shared commitment to pushing the boundaries of knowledge in our field has been truly inspiring, and I am grateful to have been part of such a dedicated and innovative community.

Moreover, I have to thank all the other people I have collaborated with and who have helped me at DESY (Lennart, Simon, Finn, Håkan, Stephan, Daniil, Sara, Paul, Adriana, Marcel, Gianpiero). Moreover, a big debt of gratitude to my colleagues at GSI and FAIR (Michael Deveau, Thomas Morhardt, Oleg Kiselev, Christian Schmidt).

To everyone who has played a part in this journey, whether mentioned here or not, please know that your support has been invaluable. This work is as much a reflection of your efforts as it is of mine, and for that, I am eternally grateful.

On a lighter note, a special thanks goes out to those who have contributed to my new-found mastery of creative profanity, my early onset of academic dementia, and the sleepless nights that have made coffee an essential part of my bloodstream. You have truly reinforced the notion that academia is a relentless maze of madness — yet, ironically, it's a maze I seem determined to keep wandering through.

Special credits go to Bauhaus, the unsung hero of our experimental setups, whose strong adhesive tapes have saved our outrageously expensive equipment built in a hurry from collapsing more times than I care to admit — a reminder that an engineering background can work wonders in the most unexpected ways. Many thanks to Nescafé, Dallmayr, and Melitta — only divine insight could explain how a person can guzzle SO MUCH coffee daily like it's a cold beer on a hot summer's day.

Then, a special shoutout to Flensburger and my dear Erdinger, and almost none to Welde (though I had to endure that... concoction... since the Bothanik café seems blissfully unaware of other beer brands — may they learn from this grievous error, so I can one day enjoy a proper brew). God bless Pascal's vast knowledge of Scottish whiskey and those blessed Christmas Eve tastings... Glühwein has nothing on that!

A huge bow to Kai's unmatched poker prowess and Michael's strategic genius—your skills have left both me and Christian impressed and utterly bankrupt. Endless apologies to Caro for having to endure my presence in the office, where I seemed to have a talent for amplifying the mood in all the wrong ways. Statistics (or the ever-full swear jar) clearly indicate that the frequency of profanities more than doubled whenever I was around. But hey, at least now you can curse fluently, nonstop, from Heidelberg to Hamburg in at least five languages — German, English, Romanian, Russian, and Polish!

I'm still convinced that this is why Annalena made her grand exit. Sorry about that! On the bright side, you're now earning triple what the rest of us do combined, so maybe it all worked out in the end!

And last but not least, a shoutout to myself and the two surviving neuron cells that somehow managed to keep firing during the final months of writing this manuscript. You've earned a well-deserved break!

To all of you, my deepest gratitude for the unwavering support in making this research possible. Without you, I probably would have graduated two years earlier!

Finally and **above all**, on the most serious tone, there are three people without whom this journey would not be possible and to whom this thesis is dedicated.

I owe an immeasurable debt of gratitude to my mother, Marilena, whose unwavering support and encouragement have been a constant source of strength throughout my life, including during this PhD journey. From a young age, you instilled in me the values of perseverance, curiosity, and dedication, which have guided me through every challenge and triumph. Your belief in my abilities, even when I doubted myself, has been the foundation upon which I've built my academic and personal achievements. Your love and immense sacrifices have made it possible for me to pursue my dreams, and for that, I am forever grateful.

To my late father, Nicolae, I owe so much of who I am today. Your wisdom, patience, and quiet strength have shaped my character and instilled in me the determination to pursue excellence in everything I do. Even though you are no longer here, your guidance and the values you imparted continue to influence my life and my work. I know you would have been proud of this achievement, and I carry your memory with me as I move forward. The lessons you taught me about resilience, integrity, and the importance of family remain with me every day.

Finally, to my wonderful wife, Maria, words cannot express the depth of my gratitude. Your unwavering support, understanding, and love have been my anchor throughout this

demanding journey, as well as many, many years before. You have been there for me during the late nights, the stressful deadlines, and the moments of doubt, always ready with a kind word or a reassuring hug. Your belief in me has been a driving force, and your presence has made this journey not just bearable, but meaningful. Thank you for your endless patience, for always helping me find my way, and for being my partner in every sense of the word. This accomplishment is as much yours as it is mine, and I look forward to the future we will continue to build together. I love you!

Introduction

“Nature uses only the longest threads to weave her patterns, so each small piece of her fabric reveals the organization of the entire tapestry.”

– Richard Feynman, *The Character of Physical Law*

In the grand tapestry of the universe, every thread, weave, and knot tells a story – a narrative of cosmic proportions that dates back to the very inception of time itself. The poetic insight of Richard Feynman into the fabric of reality serves as a profound prelude to our exploration into the realm of particle physics, a domain where the seemingly insignificant is anything but. It is here, amidst the subatomic ballet of quarks, leptons, and bosons, that we begin to discern the patterns of the grand design of the universe.

Particle physics stands as a testament to humanity’s relentless quest to understand the fundamental constituents of matter and the forces that govern their interactions. At the forefront of this exploration is collider physics, a field that involves the study of high-energy particle collisions within advanced accelerator systems. These collisions allow physicists to probe the microscopic world, uncovering the fundamental components of nature and the forces that shape them.

The tools of this exploration are the particle detectors — marvels of engineering and physics that allow us to “see” the unseeable. Each particle collision observed in these detectors adds more threads of knowledge, weaving them into the larger narrative of the universe.

A key innovation in this field are the silicon pixel sensors. Akin to the sensors found in ordinary mobile phone cameras, these devices have revolutionized particle detection. Composed of an array of pixels, each smaller than a human hair, they convert the energy of incident particles into electrical signals. This enables physicists to determine, with remarkable accuracy, the exact location where a subatomic particle has passed through.

The precision and efficiency of pixel sensors has been instrumental in achieving many of the landmark breakthroughs in physics and have made them increasingly central to the design of modern particle detectors. In fact, most future high-energy physics experiments are expected to incorporate these sensors as vital tracking layers, underscoring the pivotal role this technology will continue to play in advanced tracking systems.

This journey into the heart of particle physics is rooted in the belief, inspired by Feynman, that by understanding the smallest parts of nature, we can unlock the secrets of the vast universe. In understanding these fundamental particles and forces, we do more than just satisfy our curiosity; we unlock new technologies, enhance our comprehension of the cosmos, and, perhaps most importantly, learn our place within it.

1.1 Particle physics and the Standard Model

1.1.1 The Standard Model of particle physics

The Standard Model of particle physics is arguably the most successful theoretical framework in the history of science, achieving a level of numerical precision in its predictions that remains unmatched by any other scientific theory. It consistently and accurately predicts the outcomes of countless experiments, demonstrating an extraordinary level of reliability.

It categorizes elementary particles by their corresponding charges and explains their interactions through fundamental forces. Three of the four fundamental forces that govern the universe are described: the electromagnetic force, the strong force and the weak force.

The Standard Model includes twelve elementary particles with half-integer spin, known as fermions, and classifies them into two groups: leptons, which are influenced only by the electromagnetic and weak interactions, and quarks, which participate in all three types of interactions. Additionally, there are four gauge bosons (the gluon, photon, W, and Z bosons) mediating the interactions and the scalar Higgs boson. Particles may possess multiple charges, allowing them to be influenced by more than one fundamental force. This is shown in Fig. 1.1.

This entire theoretical framework is encapsulated in the Standard Model Lagrangian, a single equation that represents our comprehensive understanding of the subatomic world.

Despite its elegance, the Standard Model is not a complete theory of fundamental interactions. While it successfully describes three of the four known fundamental forces — electromagnetism, the strong force, and the weak force — it does not account for gravity. Additionally, it fails to explain the observed asymmetry between baryonic and anti-baryonic matter in the universe, lacks a candidate for dark matter that aligns with cosmological observations, and does not incorporate a mass for neutrinos, even though neutrino oscillations were observed.

Among the fundamental interactions, the strong force stands out due to its unique properties and its critical role in the structure of matter.




Figure 1.1 – The Standard Model of particle physics showing 12 particles (with corresponding antiparticles), the force carrying bosons and the Higgs boson. Adapted from [1].

The strong force is mediated by eight massless gluons and is behind the existence of atomic nuclei, as it governs the interactions of particles carrying color charge.

The strength of the strong interaction is determined by the strong coupling constant ¹, α_s , which depends on the energy scale of the interaction being considered (momentum transfer Q^2). Unlike QED, the gluons in QCD carry a color charge themselves, leading to self-interaction. This in turn leads to a stronger coupling constant at low Q^2 values (see Fig. 1.2) with restricted interaction range, confining quarks and gluons within composite hadrons. As such, bare quarks and gluons are never observed directly. For this reason low-energy QCD calculations are not possible using perturbation theory. At high Q^2 (> 100 GeV), typical scale of modern collider experiments, α_s becomes sufficiently small, resulting in asymptotic freedom, where quarks and gluons behave as nearly free particles.

Due to the significant variation in the strong coupling constant, nuclear matter can exhibit drastically different properties, particularly under extreme conditions such as high

¹The coupling strength or constant is a number that determines the strength of the force exerted in an interaction




Figure 1.2 – Summary of the experimental values of the strong coupling constant (α_s) as function of the energy scale Q , compared to predictions from QCD. Taken from [2].

temperatures and/or densities. These extreme states of matter are believed to have existed in the first few microseconds after the Big Bang, as well as in the dense cores of neutron stars.

To study such effects, we recreated conditions akin to those of the early universe or at the core of neutron stars in laboratory, by colliding heavy nuclei at multi-TeV energy levels.

As a result of heavy-ion collisions, an exotic state of matter called the quark-gluon plasma (QGP) is formed where the quarks and gluons composing the baryonic matter are no longer confined within hadrons, but exist in a free, strongly interacting medium. The QGP is a short-lived, non-spherical hot fireball which rapidly expands and cools. It behaves as an almost ideal liquid, with its dynamics governed by Quantum Chromodynamics (QCD). Understanding the dynamics of the fireball and its evolution is crucial for probing the properties of the QGP and advancing our knowledge of this exotic state of matter.

The investigation of ultra-relativistic nucleus-nucleus collisions at very large energies offers a way to study the QGP. These collisions aim to recreate the extreme conditions that existed in the early universe, approximately 10 ps after the Big Bang [3]. During a brief period, lasting about 10 μ s, the universe was thought to be in a QGP phase before transitioning to the hadronic matter we observe today.

Heavy-ion experiments are specifically designed to probe the QGP by analyzing various observables, despite the challenge that the QGP exists for only an incredibly short time (10^{-22} s [3]) before it transitions back into hadronic matter. As a result, direct experimental

observables of the QGP are not available, and most insights are derived indirectly from the properties of the particles generated in the heavy-ion collision. These particles may originate from different stages of the collision, such as the initial hard parton scatterings, occurring before the QGP forms, during the expansion of the QGP or at the phase transition.

For instance, heavy quarks (charm, and bottom) are particularly revealing because they experience the entire evolution of the fireball, carrying crucial information about their interactions with the QGP. In addition to heavy-ion collisions, comparisons with proton-proton collisions (as well as other, mixed systems), where a deconfined state of matter is not expected, serve as a valuable reference for understanding the differences between these two opposing collision systems at the extremes of energy density.

Heavy-flavor hadrons containing charm or beauty quarks are of particular interest since they represent a sensitive probe of the hot and dense QCD matter. However, they are produced in relatively small quantities during heavy-ion collisions, making them rare signals that are difficult to detect. The challenge is compounded by the fact that they are often obscured by a large background of soft particle production, which occurs much more frequently.

Isolating these rare heavy-flavor signals from the overwhelming background and therefore obtaining statistical significance requires advanced detection techniques, including high-resolution tracking and precise vertexing capabilities and the possibility to position highly-granular detectors as close as possible to the interaction point.

For this, it is crucial to employ detectors with exceptional tracking resolution and position them as close as possible to the interaction point. Moreover, having highly granular detectors with as little material as possible allows for more detailed and accurate reconstruction of particle trajectories, facilitating the identification and separation of heavy-flavor particles from the background. This precision is essential for enhancing our understanding of the QGP and the role heavy-flavor quarks play in its dynamics.

1.1.2 The Large Hadron Collider

The most extreme of conditions are found in collisions at the Large Hadron Collider (LHC) at CERN owing to its very high collision energies.

The LHC, the most powerful quasi-circular ² accelerator ever constructed, lies in the former LEP tunnel approximately 100 meters underground along the French-Swiss border near Geneva. The 27 km long accelerator machine houses two separate rings where particle beams are accelerated by radiofrequency (RF) cavities and bent around the ring by superconducting

²The LHC consists of eight straight and eight curved sections, enabling consecutive acceleration and bending of particle beams.

magnets. The counter-rotating bunches of particles can be collided at four distinct interaction points along the circumference of the tunnel.

The LHC dedicates the majority of its operational time to proton-proton (pp) collisions. However, for about one month each year, it shifts its focus to accelerating and colliding ions (and sometimes even asymmetric systems, such as proton-ion collisions), with lead (Pb) being the most commonly used ion. The collision products from these experiments are measured at four primary experimental sites: ATLAS, CMS, LHCb, and ALICE. ATLAS and CMS serve as general-purpose detectors, LHCb is specialized in flavor physics, and ALICE is dedicated to the study of heavy-ion collisions.

Before collisions and measurements can occur, several stages of acceleration and processing are required, as illustrated by the CERN accelerator complex in Fig. 1.3.

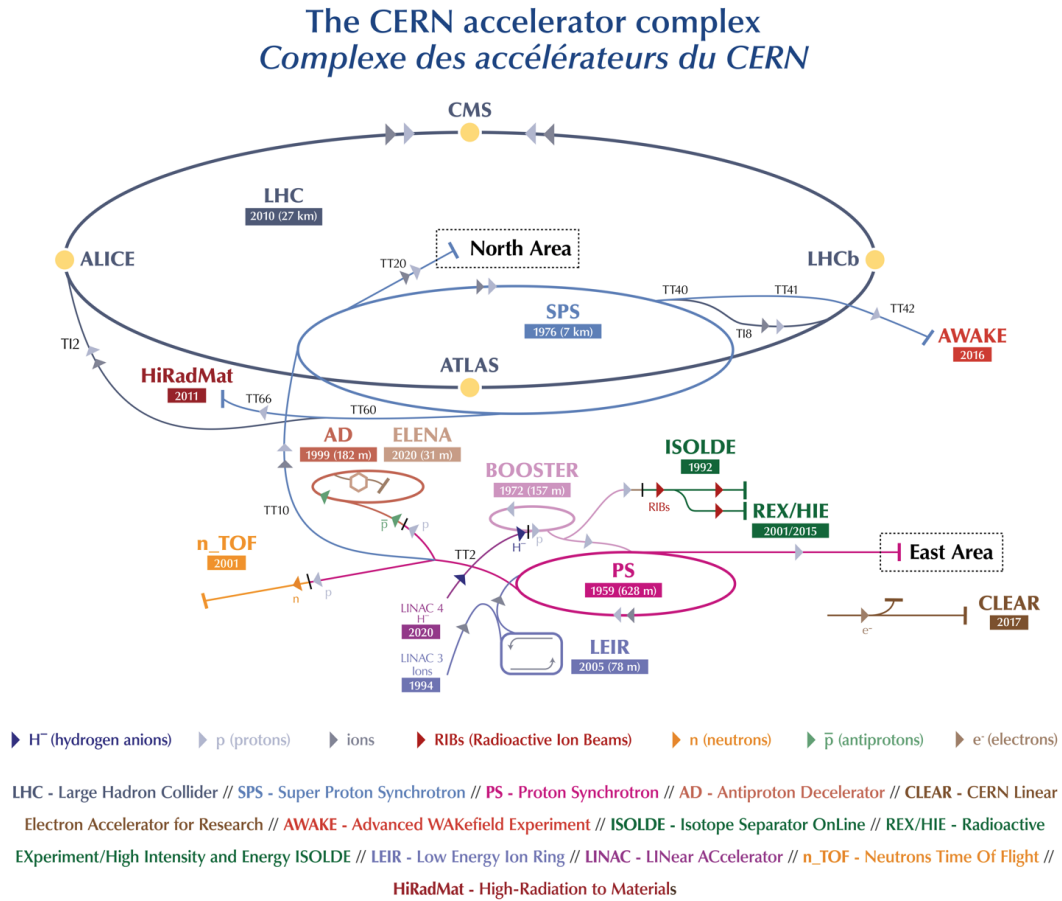


Figure 1.3 – The CERN accelerator complex. Taken from [4].

To begin, hydrogen atoms from a gas bottle are stripped of their electrons using an external electric field, producing protons. These protons are then injected into a linear accelerator

(LINAC 2), where they are initially accelerated to 50 MeV [5]. The Proton Synchrotron Booster (PSB) subsequently increases the beam energy to 1.4 GeV. Next, the Proton Synchrotron (PS) raises the energy further to 25 GeV, after which the Super Proton Synchrotron (SPS) boosts it to 450 GeV. At this point, they are injected into the LHC which accelerates each beam to a final energy of 7 TeV.

The heavy ions are accelerated in a similar manner. Lead nuclei are produced by heating a highly purified lead sample to approximately 500 °C, creating lead vapor that is ionized by an electron current. This process generates nuclei with various charge states, peaking around Pb^{29+} . These nuclei are selected and accelerated to 4.2 MeV per nucleon (MeV/u) before passing through a carbon foil, stripping most to Pb^{54+} . This beam is accumulated and accelerated to 72 MeV/u in the Low Energy Ion Ring (LEIR) before being transferred to the PS. The PS accelerates the beam to 5.9 GeV/u, after which it passes through another foil to be fully stripped to Pb^{82+} . The SPS then accelerates the nuclei to 177 GeV/u and transfers them to the LHC, where they are further accelerated to 5.36 TeV/u.

When lead nuclei collide centrally (most nucleons participate in the collision) in the LHC, they generate temperatures in excess of 10^{12} K, and on average, each collision event produces around 2000 tracks at mid-rapidity, with pions, protons, and kaons being the most commonly observed particles.

The LHC operates in phases called Runs. The first Run occurred from 2009 to 2013, followed by Run 2 from 2015 to 2018. Run 3 began in 2022 and is expected to continue until the end of 2025. A fourth Run is anticipated to start in 2029 and last until 2032. Between these Run periods, extensive technical shutdowns take place, during which the detectors are typically upgraded. The upgrade work and detector assembly starts a few years in advance of the upgrade operation.

1.1.3 The ALICE experiment

ALICE (A Large Ion Collider Experiment) is the dedicated heavy-ion experiment at the LHC, devoted to the study of QCD under extreme conditions of high temperatures and energy densities.

To achieve its physics objectives, the experimental program of ALICE requires good particle identification (PID) capabilities across a broad range of transverse momentum for hadrons, electrons, photons, and muons. This capability is crucial, particularly in the high-multiplicity environment of Pb-Pb collisions, where as many as 2000 particles can be produced at mid-rapidity. These requirements are met by a set of 18 subsystems, each with unique

detection, tracking and PID capabilities, which are illustrated in the schematic overview of the ALICE detector provided in Fig. 1.4.




Figure 1.4 – ALICE experiment – detector subsystems as present during Run 3 (after the LS2 upgrade). Taken from [6].

The ALICE detector is composed of two main sections: a central barrel, which houses subdetectors within the L3 solenoid magnet (depicted in red) that operates at a nominal magnetic field of 0.5 T, and a forward muon spectrometer located outside the central barrel. The individual detector subsystems and their contributions to physics goals of ALICE are discussed comprehensively in [7]. In this chapter, only two detector systems are described, with one being explored in greater detail in a subsequent section.

The first detector is the Inner Tracking System (ITS), situated closest to the interaction point. The ITS consists of seven concentric layers of monolithic active pixel sensors. This ITS2 was entirely changed before the start of Run 3, replacing the previous generation (used in Runs 1 and 2) that comprised a mix of silicon pixel, drift, and strip detectors.

Key design objectives for the ITS2 include highly efficient tracking — both standalone and in conjunction with the Time Projection Chamber — especially at low momenta, and the precise reconstruction of primary and secondary vertices, particularly from hadrons containing charm and beauty quarks. These objectives were met by moving the first layer closer to the interaction point, reducing the material budget of the inner layers, using more granular pixels, and adding a seventh detection layer.

Further out from the ITS2 is the Time Projection Chamber (TPC), the primary tracking detector of ALICE. The TPC is an 88 m² gas-filled cylinder equipped with readout chambers at its endcaps [8]. For Run 3, the TPC was upgraded with novel readout chambers composed of a

stack of four Gas Electron Multiplier (GEM) foils, enabling continuous readout and recording rates of up to 50 kHz in Pb-Pb collisions.

The TPC is designed to provide accurate measurements of charged particles, ensuring good separation of tracks and enabling particle identification by measuring the energy loss of charged particles as they pass through the detector volume.

1.2 Particle detection

1.2.1 Interaction of particles with matter

Particles that traverse matter interact with it at the most fundamental level, perceiving it as a collection of electrons and nuclei, and even subcomponents of these nuclei. The nature of these interactions depends on the type of radiation, its energy, and the material it encounters. Reactions can occur with the atoms or nuclei as a whole, or directly with their constituent particles, through various allowed channels.

The likelihood of these interactions is dictated by quantum mechanics and the relative strengths of the fundamental forces involved. For charged particles and photons, electromagnetic interactions — particularly inelastic collisions with atomic electrons — are by far the most common processes, given the relative strength and long range of the Coulomb force compared to other interactions.

Upon interacting electromagnetically, a particle can be deflected and/or lose energy through elastic scattering from nuclei and inelastic collisions with atomic electrons, respectively. The elastic scattering is occurring less frequent than the collisions with nuclei and usually involves only a small energy transfer, given that the nuclei of most materials are significantly more massive than the incident particle.

The following paragraphs will focus on the specific interactions pertinent to the work discussed in this thesis. The explanations provided here are based on the descriptions found in [9, 10]. For a more comprehensive review of particle interactions in matter, refer to those two works.

1.2.1.1 Ionization, excitation, energy loss

Most particle detectors operate by detecting the ionization or the emission of light by charged particles. Electrically neutral particles, such as neutrons, must interact with matter to produce charged particles, which then generate measurable signals. Neutrinos are an

exception, as their interactions with matter are extremely weak, making them particularly challenging to detect.

Energy loss by particles traversing matter is a stochastic process, meaning it cannot be precisely predicted on an individual particle basis. The number of interactions and the interaction strength fluctuate, but average energy losses are well understood and can be described in detail. These losses are often discussed separately for heavy charged particles on one hand, and electrons and positrons on the other, to accurately account for effects such as Bremsstrahlung and the indistinguishability of electronic and impact electrons at the quantum mechanical level.

As charged particles are passing a detector, they lose energy primarily through ionization and excitation of the atoms of the detector medium. This process is dominant up to the threshold of radiative effects at large energies.

The mean energy loss for moderately relativistic particles ($0.1 < \beta\gamma < 100$) via both excitation and ionization is well described by the Bethe-Bloch formula, which offers an accuracy within a few percent:

$$\left\langle -\frac{dE}{dx} \right\rangle = K\rho \frac{Z}{A} \frac{z^2}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{\max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} - \frac{C(\beta\gamma, I)}{Z} \right] \quad (1.1)$$

where $K = 0.307 \text{ MeV cm}^2 \text{ mol}^{-1}$, z is the charge and β the velocity of the projectile, Z and A the atomic number and mass number of the detector medium, I the mean excitation energy and W_{\max} is the maximum possible energy transfer, calculated from the kinematics of an elastic collision of the particle with a shell electron. The maximum energy transfer happens for a central collision and is given by:

$$W_{\max} = \frac{2m_e c^2 \beta^2 \gamma^2}{1 + 2\gamma m_e/M + (m_e/M)^2} \quad (1.2)$$

The values are usually expressed in $\beta\gamma$ which is nothing else but the ratio of momentum and mass (p/m) of the incident particle.

At $\beta\gamma \sim 0.1$ the projectile speed is comparable to the “speed” of atomic electrons. At $\beta\gamma \sim 1000$ radiative effects become significant. Two corrections are applied: δ , a density correction that becomes important at high energies, and C/Z , a shell correction relevant at low projectile velocities. The density effect arises because the electric field of the projectile polarizes the atoms along its path, reducing the full electric field intensity experienced by distant electrons. Consequently, these outer electrons contribute less to the total energy loss than the Bethe-Bloch formula would predict. The shell correction accounts for cases where

the velocity of the incident particle comes in the range of shell electron velocities, at which point the Bethe-Bloch formula no longer accurately describes the interaction.

The Bethe-Bloch formula describes the stopping power of the charged particle projectile and is depicted in Fig. 1.5.




Figure 1.5 – Mean energy loss $\langle dE/dx \rangle$ as a function of the energy of the incoming particle for π^\pm in silicon. Taken from [9].

This quantity is different than the measured ionization charges in the detector, which is smaller, since it does not account for atomic excitations with consequent de-excitation via low energy photons. This distinction is crucial in silicon detectors, where part of the energy deposited is used to generate free electron-hole pairs, which are detected, while the rest is dissipated thermally via phonon excitation.

At low energies, the energy loss of particles exhibits a strong dependence on the $1/\beta^2$ term, leading to a steep rise in energy loss as the velocity of the projectile decreases. Conversely, at high energies, the $(\ln \beta\gamma)$ term becomes significant, causing the energy loss to increase more gradually. Between these two regions lies a broad minimum around $\beta\gamma \sim 3-3.5$, where projectiles are known as minimum-ionizing particles (MIPs). In high-energy physics (HEP), due to the relatively moderate increase in energy loss at higher energies compared to those corresponding to MIPs, it is common occurrence to refer to all particles with momenta above

the minimum-ionizing threshold as MIPs. This designation is used because these particles exhibit similar energy loss characteristics, despite having higher momenta.

The Bethe-Bloch equation needs adjustments for electrons and positrons due to their much smaller mass compared to heavier particles, which makes energy loss due to bremsstrahlung significant even at lower energies. The maximum energy transfer for electrons can be calculated using Eq. (1.2), but since atomic and projectile electrons are identical particles, it is only half of this value: $W_{\max} = m_e c^2 (\gamma - 1)/2$. For positrons, the division by two is not needed, and a similar formula can be derived.

The energy loss for electrons and positrons is shown below and taken from [11]:

$$-\left(\frac{dE}{dx}\right)^{\pm} = 0.1535 \frac{\rho Z}{A\beta^2} \left\{ \ln \left[\frac{\tau^2(\tau+2)}{2 \left(\frac{I^2}{mc^2}\right)} \right] + F(\tau)^{\pm} - \delta \right\} \text{ [MeV/cm]} \quad (1.3)$$

where

$$\tau = \frac{E_{\text{kin}}}{mc^2} \quad \beta = \frac{v}{c} = \frac{\sqrt{\tau(\tau+2)}}{\tau+1} \quad \gamma = \tau+1$$

$$F(\tau)^+ = 2 \ln 2 - \frac{\beta^2}{12} \left[23 + \frac{14}{\tau+2} + \frac{10}{(\tau+2)^2} + \frac{4}{(\tau+2)^3} \right]$$

$$F(\tau)^- = 1 - \beta^2 + \frac{\tau^2/8 - (2\tau+1) \ln 2}{(\tau+1)^2}$$

The mean energy loss for electrons and positrons in silicon is shown in Fig. 1.6 for a range of momenta, calculated using Eq. (1.3).

For a given particle, the actual energy loss will typically differ from the mean value due to statistical fluctuations in the number of collisions and the energy transferred in each collision. The mean is skewed by rare events with large single-collision energy losses, shifting the mean to higher values. Therefore, the most probable energy loss, which is smaller than the mean for such skewed distributions, is often a better estimator of energy loss. To exemplify this, the energy loss of pions in silicon is shown in Fig. 1.7

In detectors of moderate thickness, the energy loss probability distribution is usually described by the Landau distribution, which is effective for thin silicon absorbers ($> 300 \mu\text{m}$). For very thick absorbers, the distribution becomes less skewed but never fully approaches a Gaussian, as might be expected from the Central Limit Theorem due to the large number of collisions.

However, the Landau model does not accurately describe energy loss in ultra-thin silicon absorbers ($< 160 \mu\text{m}$) [2]. In such thin layers, where the number of collisions is small, and the

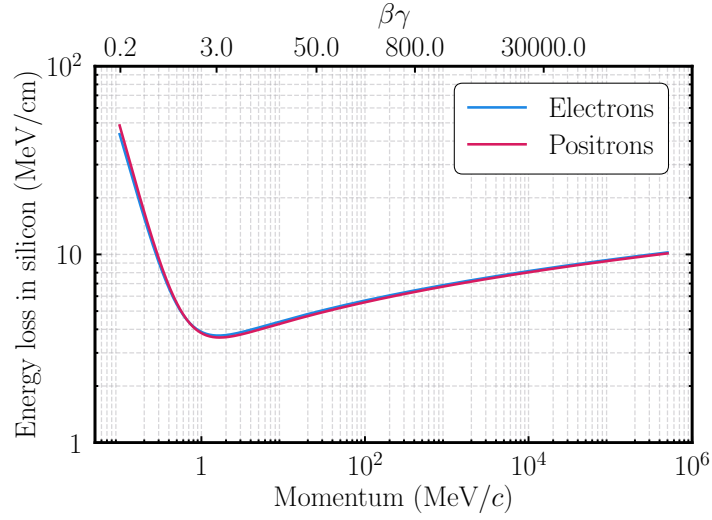


Figure 1.6 – Mean energy loss of electrons in silicon calculated using Eq. (1.3).

deposited energy still reflects the shell structure of the silicon atom, a more accurate model is the Bichsel model [12] (modeled as a convolution of a Landau and a Gaussian distribution), which has been shown to reproduce energy losses in ultra-thin silicon very accurately [13].

The most probable value of energy deposition in ultra-thin silicon sensors (13 μm to 100 μm) for projectiles with $\beta\gamma > 100$ can be calculated using the empirical formula from [12]:

$$\Delta_p[\text{eV}] = t [\mu\text{m}] \cdot (100.6 + 35.35 \cdot \ln t [\mu\text{m}])$$


This calculation carries an uncertainty of 1.2%. The most probable energy loss is on the order of 66 eV/ μm for sensitive layers about 25 μm thick.

In thin detector layers, highly energetic atomic electrons, known as δ electrons, that are ejected from the shell may escape the detector. When these electrons leave the detection volume, their energy is not fully measured, leading to a reduction in the mean observed energy.


Moreover, electrons and positrons lose energy not only through ionization, but also by emitting electromagnetic radiation in the Coulomb field of the nuclei, a process known as bremsstrahlung. The energy loss due to bremsstrahlung can be expressed as:

$$\left(\frac{dE}{dx} \right)_{\text{rad}} = -\frac{E}{X_0} \quad (1.4)$$

where X_0 is the characteristic length for energy loss through bremsstrahlung and is called radiation length. It is a material dependent quantity and represents the average distance



(a) Energy loss of 500 MeV pions, normalized, for different thicknesses of silicon



(b) Most probable energy loss, scaled to the mean energy loss of a MIP ($388 \text{ eV } \mu\text{m}^{-1}$, or $1.66 \text{ MeV cm}^2 \text{ g}^{-1}$)

Figure 1.7 – Energy loss fluctuations in silicon. Taken from [2].

an electron travels before its remaining energy is $1/e$ (approximately 63% of its energy is radiated off).

Radiation length values are typically tabulated (for example [2]) for various elements and compounds. An approximation can be obtained using:

$$\rho X_0 = \frac{716.41 A [\text{mol/g}]}{Z(Z+1) \ln \left(\frac{287}{\sqrt{Z}} \right)} \left[\frac{\text{g}}{\text{cm}^2} \right] \quad (1.5)$$

The dimensionless quantity x/X_0 , where x is the thickness of the material in units of length, is often of interest. x/X_0 represents the thickness of the scattering medium in units of radiation length.

1.2.1.2 Multiple Coulomb scattering

Electromagnetic interactions influence both the energy and the trajectory of a particle. In addition to inelastic collisions with atomic electrons, charged particles traversing matter undergo repeated elastic Coulomb scatterings from nuclei, although these occur with somewhat lower probability. These low probability scattering events are described by the Rutherford cross-section, which follows a $1/\sin^4(\theta/2)$ dependence, meaning most collisions result in only small angular deflections.

In the textbook scenario where the nuclei are much more massive than the incident particles, the energy transfer to the nucleus is minimal, leading to negligible deflection. As a result, the path of a particle through matter becomes a meandering journey filled with small

twists and turns. The cumulative effect of these small-angle scatterings is a net deflection and an offset from the original trajectory. Given the stochastic nature of scattering, identical particles entering the same medium with the same energy will not follow identical paths, ultimately producing a statistical distribution of scattering angles.

The mean number of scattering processes, N , in a material slab of thickness x , is given by [14] as $N = \sigma x \frac{N_A \rho}{A}$, where σ is the integrated elastic cross-section, N_A is Avogadro's constant, ρ is the density of the material and A the atomic mass number of the nucleus.

Then, according to [14] the number of scatterers can be expressed as:

$$N_{\text{scatt}} = \frac{x}{X_0} \frac{1}{\beta^2} \left(\frac{1.587 \cdot 10^7 Z^{\frac{1}{3}}}{(Z + 1) \ln(287 Z^{-\frac{1}{2}})} \right) \quad (1.6)$$

The number of scatterings as a function of the thickness of the scattering medium in units of radiation length is shown in Fig. 1.8, using Eq. (1.6).




Figure 1.8 – The number of scatterers as a function of the radiation length for three materials

As seen in the figure, the average number of scatterings depends on the radiation length x/X_0 . Since the radiation length has a density dependence, the number of scatterers in denser materials like Ni is larger, whereas for Al and Si, which have similar densities, the numbers are comparable.

When the number of scatterers exceeds about 20, the scattering process is referred to as multiple or Molière scattering. According to the central limit theorem, for an infinite number of scatterings, the distribution of scattering angles should approach a Gaussian distribution.




Figure 1.9 – (left) Quantities used to describe multiple Coulomb scattering. The particle is incident in the plane of the figure. Taken from [2]. **(right)** Projected angular distribution.

At small angles, the angular distribution closely resembles a Gaussian. However, as the angle increases, additional terms come into play, adding a broad tail. Large-angle deflections are typically due to single, large-angle Coulomb scatterings rather than the cumulative effect of many small-angle scatterings. Consequently, this broad tail should approximate the Rutherford $1/\sin^4(\theta/2)$ form for single scattering rather than a Gaussian distribution.

The distribution of scattering angles is accurately described by Molière theory [15].

The plane projection of the scattering angle distribution (see Fig. 1.9) is defined as:

$$\theta_0 = \sigma_\theta^{\text{plane}} = \frac{1}{\sqrt{2}} \sigma_\theta^{\text{space}}$$

The factor $\sqrt{2}$ accounts for the reduction in dimensionality when considering the scattering angle in one plane rather than in three-dimensional space, given the isotropic nature of scattering.

The RMS of the central 98% of the projected scattering distribution can be described by the Highland formula [16], with the corrections from Lynch and Dahl [17]:

$$\theta_0 = \frac{13.6 \text{ MeV}}{\beta c p} z \sqrt{\frac{x}{X_0}} \left[1 + 0.038 \ln \left(\frac{x z^2}{X_0 \beta^2} \right) \right] \quad (1.7)$$

Here p , β , and z are the momentum, velocity, and charge number of the incident particle, and x/X_0 represents the thickness of the scattering medium in radiation lengths.

For the inner 98% of the distribution core, this parametrization deviates by less than 11% [2, 17] from a Gaussian fit for all Z and for $10^{-3} < x/X_0 < 100$.

The formula from Highland simplifies multiple Coulomb scattering calculations by parameterizing Molière theory into an elegantly simple formula that includes a single scattering correction (in the parentheses).

Lynch and Dahl in [17], propose a more precise approximation of multiple Coulomb scattering that agrees with Molière scattering predictions to within 2% for all Z . However, this approach removes the dependence on radiation length X_0 . Since the focus of the entire Chapter 3 is on linking radiation length (and hence material budget) to the scattering angle, the Highland formula is preferred for this discussion.

1.2.2 Silicon detectors

Detectors capable of accurately tracking the passage of particles, allowing for the determination of key properties such as momentum, charge, and identity, are essential components of modern particle physics experiments, especially in the high-occupancy environments of hadron and heavy-ion colliders.

The performance of these tracking detectors is crucial to the success of all physics analyses. Heavy flavor physics, in particular, strongly depends on excellent secondary vertexing capabilities, which are crucial for identifying and studying hadrons containing b or c quarks. To achieve this, the geometry and materials of the detectors are meticulously chosen, with a preference for highly granular, low-mass designs that enhance precision while minimizing material-induced disturbances to the particle trajectories.

Silicon semiconductor detectors are the preferred technology in most current high-energy particle physics experiments due to their numerous advantages over alternative detector types [18]. When a charged particle traverses a silicon sensor, it creates a number of electron-hole pairs in the sensitive layer. These charge carriers then drift towards the electrodes under the influence of an electric field, which is the result of both the pn-junction built-in potential and (possibly) an externally applied bias. The movement of charges induces a current on the readout electrodes, which is subsequently amplified, shaped, possibly discriminated, and read out by the front-end electronics.

One of the key advantages of silicon detectors is their relatively low mean ionization energy, approximately 3.6 eV, which is significantly lower than that of gaseous tracking detectors, where the ionization energy is about ten times higher. This lower ionization energy in silicon and the higher material density allows for the generation of a greater number of charge carriers, enhancing the sensitivity of the detector.

The development of advanced detectors has been made possible by the widespread availability of high-purity silicon and significant progress in microfabrication techniques coming from industry. These sophisticated sensors are designed for precise measurements on very small scales. Position sensitivity in silicon detectors is accomplished by dividing the implants and electrodes into segments, which can be arranged as strips for one-dimensional position

detection or as pixels for two-dimensional detection. Today's silicon detectors are capable of achieving spatial resolutions as fine as $10\text{ }\mu\text{m}$. By combining data from multiple layers within the detector, the complete trajectory of a particle passing through can be accurately reconstructed.

The widespread availability of high-purity silicon, along with advancements in microfabrication techniques, has enabled the development of highly sophisticated detectors capable of performing precise measurements at very small scales. Position sensitivity in silicon sensors is achieved by segmenting the chips, which can be designed in either strip-like configurations for one-dimensional or pixel-like configurations for two-dimensional position measurements. Modern silicon detectors can achieve intrinsic spatial resolutions below $10\text{ }\mu\text{m}$. To reconstruct the full trajectory of a passing particle, measurements from multiple layers of the detector are combined. Like this, the positional uncertainty on the track position can reach even $1\text{ }\mu\text{m}$.

For over four decades, a wide array of experiments have leveraged the resolution and accuracy of silicon detectors [19] in order to explore the world of HEP. These detectors have continually evolved, expanding in their usable area, moving into three-dimensional configurations, and becoming faster and more complex. Silicon detectors are expected to continue evolving to meet the diverse needs of both science and technology.

This subchapter will specifically focus on monolithic silicon pixel sensors, particularly those used in the ALICE ITS upgrades. Only relevant aspects of these technologies will be discussed, and even those only briefly. For a more comprehensive exploration of other technologies and detailed explanations, readers are encouraged to consult [9, 20].

1.2.2.1 Semiconductor physics, doping

Semiconductors possess unique conductivity properties due to the small energy gap between bound electrons in the valence band and quasi-free electrons in the conduction band. Silicon is the most widely used material for semiconductor detectors, though other materials such as germanium (Ge), gallium arsenide (GaAs), cadmium telluride (CdTe), and diamond are also employed in specific applications.

Silicon crystallizes in a diamond cubic crystal lattice, where each silicon atom participates in a covalent bond forming a structure that is periodic in space. This tetrahedron structure is the primitive unit cell of the crystal whose translational repetition generates the crystal lattice.

In a solid-state lattice, the dense periodic arrangement of atoms causes the energy levels of individual atoms to split due to the influence of neighboring atoms. Some energy levels are so closely spaced that they form continuous energy bands to which electrons are confined.

Between the bands lies a band-gap, a region where electrons are forbidden because there are no allowable electronic states for them to occupy. However, when electrons in a lower band gain sufficient energy, they can jump to a higher band by absorbing a phonon (heat) or a photon (light). The two highest energy bands, known as the valence band (VB) and the conduction band (CB), are critical in determining the electrical conduction properties of semiconductors.

In semiconductors, the bonds between neighboring atoms are weaker than in insulators, resulting in a smaller energy band gap (1.12 eV in silicon). This gap can be overcome by thermal excitations or external electric fields, leading to the creation of a free electron in the conduction band and a corresponding hole in the valence band. These free charge carriers contribute to current flow under the influence of an electric field.

By introducing impurity atoms, thereby doping the material, the electrical conductivity properties of semiconductors can be selectively modified. For instance, doping silicon with pentavalent elements such as phosphorus, arsenic, or antimony, known as donors, introduces extra conduction electrons, resulting in n-type material (the majority of charge carriers in the crystal are negative electrons). Conversely, doping with trivalent elements such as boron, aluminum, or gallium, known as acceptors, introduces extra holes, creating p-type material (the majority of charge carriers in the crystal are positive holes). This is shown in Fig. 1.10. Depending on the doping concentration, semiconductors may be denoted as n^+ or n^{++} for higher doping levels, with analogous designations for p-type semiconductors.

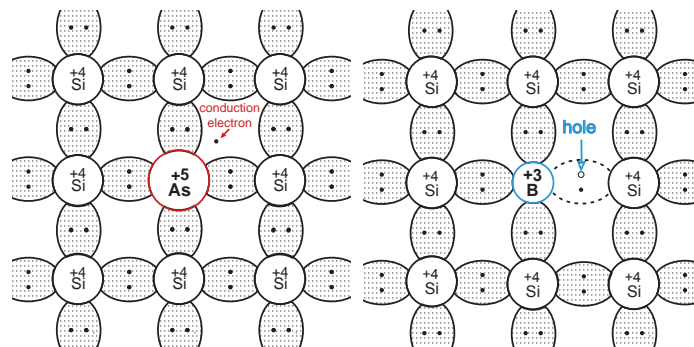


Figure 1.10 – (left) n-type and (left) p-type doping of a silicon lattice. Taken from [9]

Doped semiconductors are referred to as extrinsic semiconductors, in contrast to undoped intrinsic semiconductors. Despite the addition of impurities, the semiconductor remains electrically neutral, but the number of available charge carriers (electrons or holes) increases or decreases, thereby altering the conduction properties of the material.

In intrinsic semiconductors, the number of thermally generated electron-hole pairs at room temperature (which represents unwanted signal, the noise) is on the order of 10^8 for a $50\text{ }\mu\text{m}$ thick sensor with an area of 1 cm^2 [9]. Meanwhile, a MIP passing through the same sensor loses about 3.5 MeV/cm (see Fig. 1.6). With a mean electron-hole pair creation energy of 3.65 eV [9] and considering the most probable value, approximately 6000 electron-hole pairs are generated in the active area of the sensor. The number of electron hole pairs produced by the MIP is impossible to distinguish from the sea of free charge carriers.

Doping is a crucial step in the production of integrated circuits. Doping can be achieved either by thermal diffusion or by ion implantation. Ion implantation, the more prevalent method, involves generating an ion beam and directing it at the intrinsic silicon at precise energies (typically in the tens to hundreds of keV) so that the ions come to rest just beneath the surface (a few nanometers deep). This method is particularly important because it allows for localized doping, essential for the creation of integrated circuits and semiconductor devices.

1.2.2.2 PN-junction, reverse bias

When p- and n-doped semiconductor materials are brought into contact, they form a pn junction as shown in Fig. 1.11. At this junction, a concentration gradient of the two types of charge carriers (electrons and holes) leads to the formation of a diffusion current. Electrons from the n-type region diffuse into the p-type region, and holes from the p-type region diffuse into the n-type region. At the boundary, recombination of these carriers occurs, creating a region devoid of free charge carriers known as the depletion zone.

As charge carriers diffuse, an electric field builds up due to the immobile donor and acceptor ions, counteracting the diffusion process until equilibrium is reached. This results in a space charge region: the p-type region becomes negatively charged, and the n-type region becomes positively charged. The resulting built-in voltage, or diffusion potential, for silicon is approximately 0.5 V to 0.8 V . In this equilibrium state, drift and diffusion currents are balanced.

The behavior of the pn junction can be manipulated by applying an external bias voltage. When the p-doped side is connected to a positive voltage and the n-doped side to a negative voltage, the depletion zone is reduced, a condition known as forward bias. In contrast, applying a reverse bias, where the p-side is connected to a negative voltage and the n-side to a positive voltage, increases the depletion zone. This reverse-bias mode is typical for detector operation because it maximizes the depletion zone, creating a larger volume devoid of free charge carriers, which is ideal for detecting incoming particles. The generated electron-hole

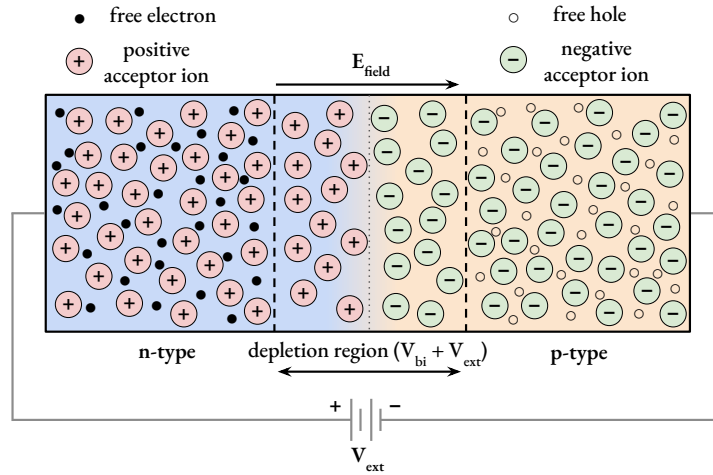


Figure 1.11 – A pn junction in thermal equilibrium with a reverse bias applied that enlarges the depletion region.

pairs within this zone drift under the influence of the electric field towards the electrodes, producing a current or charge signal that can be detected.

Boundaries between regions of the same doping type but with different doping concentrations (n^+ - n or p^+ - p) also create potential gradients similar to those found in pn junctions. Additionally, metal-semiconductor boundaries can form ohmic contacts, essential for creating functional electronic devices.

Pn junctions are the simplest semiconductor devices; a pn junction, when connected to a circuit, represents a diode. More complex devices, such as bipolar junction transistors (BJTs), are created by combining p-type and n-type semiconductors in configurations like n-p-n or p-n-p. Integrating multiple semiconductor devices on a single chip leads to the creation of integrated circuits.

A MOS (metal-oxide-semiconductor) transition is a double interface structure consisting of three layers: metal, oxide, and semiconductor. MOS structures are critical in microelectronics, particularly in field-effect transistors (FETs). The most widely used type is the MOSFET, which forms the backbone of most modern electronics, including the readout electronics for detectors. They can be made of either p-type or n-type semiconductors.

CMOS (complementary MOS) technology, which combines symmetrical NMOS and PMOS MOSFET transistors on the same substrate, is the most common approach for building complex circuits based on complementary logic. To accommodate both transistor types on the same substrate, one transistor type is embedded in a doped area called a “well” (see Fig. 1.12). For example, in a p-type substrate, PMOS transistors are embedded in n-wells. An essential advantage of CMOS technology is that it allows for circuit designs without resistors, reducing

power consumption, as transistors only draw current during switching, making power usage largely dependent on switching frequency.





Figure 1.12 – (left) Cross section of two MOSFET transistors in a CMOS gate. While the NMOS transistor is directly made on the substrate, the PMOS transistor is housed in a n-well. Taken from [9] **(right)** Metal lines cross the entire extent of the circuit and are connected vertically through vias. Adapted from [21].



Wells not only allow for the integration of both NMOS and PMOS transistors on any substrate type but also shield the transistor terminals from forming unwanted pn junctions with the substrate, which could otherwise interfere with the collection nodes and compete with charge collection. For example, in Fig. 1.12, the source and drain of the NMOS transistor, which are n-doped, form pn junctions with the p substrate.

On the top side of the sensor where the circuitry is located, multiple metallic layers are present. These layers are insulated from each other and are used for signal connections, power supply, or grounding. Thinner metal lines are used for local connections, while thicker metal lines on top connect globally. Stacked vias³ between layers enable vertical connections, as shown in Fig. 1.12.

In the context of imaging sensors, special processing steps are often introduced to optimize critical imaging parameters like leakage current and optical response. These processes are typically implemented on an epitaxial layer, which provides superior imaging performance, including better control of lateral cross-talk. An epitaxial layer is a thin, high-purity layer of silicon grown on top of the bulk silicon substrate, usually a few micrometers thick. This layer has superior electrical properties compared to the underlying bulk material and plays a crucial role in enhancing sensor performance by improving charge collection efficiency and reducing noise.

The epitaxial layer offers a more controlled environment for charge carriers (electrons and holes) generated by incident radiation or light within the pixel sensor. Without an

³Electrical connection between two or more metal layers; essentially a small drilled hole

epitaxial layer, the bulk silicon might contain higher levels of impurities and defects, which can trap charge carriers and reduce the sensitivity of a sensor. In contrast, the epitaxial layer is engineered to be as defect-free as possible, allowing charge carriers to move rapidly and efficiently toward the collection nodes, which is essential for high-resolution imaging.

1.2.2.3 CMOS nodes

The technology node denotes a specific semiconductor manufacturing process and its associated design rules. Different nodes often correspond to different generations of circuits and architectures. Generally, smaller technology nodes indicate smaller feature sizes, resulting in a denser net of transistors that are both faster and more power-efficient.

Recent technology nodes, such as those produced by Intel and TSMC, use advanced 2 nm to 3 nm extreme ultraviolet (EUV) lithography. These nodes offer exceptional performance but come with very high costs and complexity, making them unsuitable for many HEP applications. Moreover, at these small sizes, transistor gates are so tiny (about 20-30 atoms in length) that sustained radiation damage, common in HEP environments, can render them unusable.

The choice of fabrication process in microelectronics is not merely about selecting a feature size, but rather about opting for a specific process offered by a particular vendor. This decision is influenced by three key factors:

- It is crucial to use a fabrication process that has already been qualified for its intended purpose, such as ensuring radiation tolerance, to avoid the need for new qualification efforts. In the case of HEP sensors, a new technology shift is often accompanied by a long period of radiation hardness qualification, for both the sensors bulk properties, as well as for the integrated electronics. Using a process that is already qualified and understood reduces risk and simplified the design process.
- Sensor design teams are typically small, and a significant effort is made to retain expertise within the group while also recruiting additional talent to contribute to the work. Mastering the necessary tools takes time, and maintaining this specialized knowledge in-house is a top priority. This is why it is efficient to standardize the number of technologies that design teams need to master, which helps in reducing costs, adhering to schedules, and minimizing risks. Consequently, once a fabrication process is adopted, it often becomes the go-to option for future designs.
- Consistent collaboration that helps standardize the sensor design process. Typically, a collaboration is established between a microelectronics group and a specific foundry.

This partnership, often formalized through long-term contracts, is maintained over the years. This results in a uniform and efficient use of resources across multiple projects, serving the whole HEP community.

More advanced technologies offer several advantages, such as higher speeds and lower power consumption, both of which are critical for large-scale experiments with millions of channels or specialized applications requiring precise timing. Smaller feature sizes also enable finer segmentation of detectors, such as silicon pixel detectors, leading to better position resolution or lower occupancy rates. At the same time, deep sub-micron technologies are inherently highly tolerant to radiation, mainly due to the reduced oxide thickness [22].

However, these cutting-edge technologies come with significant downsides. They are more expensive and require more resources for training. Moreover, the designs are way more complex and extra verification steps are needed before they are approved for production. Although the cost per mm^2 may be somewhat offset by the reduced area needed for specific circuits, the initial costs, such as mask costs for prototyping, can be quite high. Nonetheless, even though smaller areas are required to integrate the same functionality in more advanced nodes, sensors in high-energy physics (HEP) experiments still need to cover large areas. This is because HEP experiments often involve detecting particles across wide acceptance regions, requiring sensor arrays that span significant physical dimensions.

Additionally, the reduction in rail voltages, a common trend in deep submicron technologies, presents significant challenges for both I/O and analog circuits, which are critical components of detector readout systems. As transistors in these technologies shrink, the operating voltages must be lowered to prevent excessive power dissipation and to ensure the reliability of the device. However, as supply voltages approach 1 V, the design of reliable analog circuits becomes increasingly complex. This is because lower voltages reduce the available headroom for signal swings and increase the susceptibility to noise, making it harder to achieve the necessary precision and performance. Moreover, the reduced voltage levels can also complicate the interfacing with other components that may operate at higher voltages, further challenging the design of robust and efficient detector systems.

While older technologies may be sufficient for many experiments, they have a finite lifetime as foundries gradually phase out older, slower processes due to dwindling customer demand. The constant evolution of the semiconductor industry, coupled with the benefits of newer technologies for certain experiments, drives the HEP community to continuously consider adopting new nodes.

The current generation of silicon pixel sensors in ALICE features sensors built in the 180nm CMOS node, while future upgrades will see a transition to a deeper submicron node, the 65nm.

In the ALICE experiment, the current generation of silicon pixel sensors is built using the 180 nm CMOS node (ALPIDE sensors in Run 3) of Tower Semiconductor. Future upgrades will involve transitioning to a deeper submicron node of the same parent company, specifically the 65 nm TPSCo process, to take advantage of its benefits in speed, power efficiency, and circuit density. The sensors installed in Run 4 and beyond will be produced in this node.

1.2.2.4 Signal formation

To detect traversing particles in silicon detectors, it is essential to produce, collect, and measure charge carriers. The signal formation in silicon detectors involves several well-understood mechanisms. The primary process that mediates energy loss is dependent on the energy of the incident particle and is described by the Bethe-Bloch formula Eq. (1.1). As previously mentioned, energy loss in thin sensors is subject to significant fluctuations, which can be modeled by a Landau or Bichsel distribution.

When a charged particle passes through a silicon detector, it ionizes the silicon atoms along its path, creating a trail of electron-hole pairs. On average, in a 50 μm thin silicon sensor about 66 electron-hole pairs per μm are created. At the pn junction, an electric field — established by the built-in potential and potentially enhanced by an externally applied voltage — causes these electron-hole pairs to drift toward the electrodes.

The electron-hole pairs can be categorized into two groups: those generated inside the depletion zone and those generated outside of it. Pairs generated outside the depletion zone either diffuse and recombine, or eventually diffuse and reach the depletion zone. Inside the depletion zone, the electric field separates the electrons and holes, driving them toward the contact points. As these charges move inside the sensor, they induce mirror charges on the collection electrodes positioned on top of the n-doped region. These induced charges are then amplified and used to detect the traversing particles.

According to the Shockley-Ramo theorem, a current is induced at the electrodes throughout the entire drift and diffusion process. The characteristics of this current strongly depend on the geometry of the implant and the sensor segmentation, as this influences the weighting potentials.

Once the charge is collected, a rapid potential drop occurs at the input node of the amplifying circuit. This input node is continuously reset to restore the potential to its baseline

level. Some sensors include a pulse injection capacitance, which is used to inject test charges into the analog front-end for calibration purposes (see Fig. 1.13).

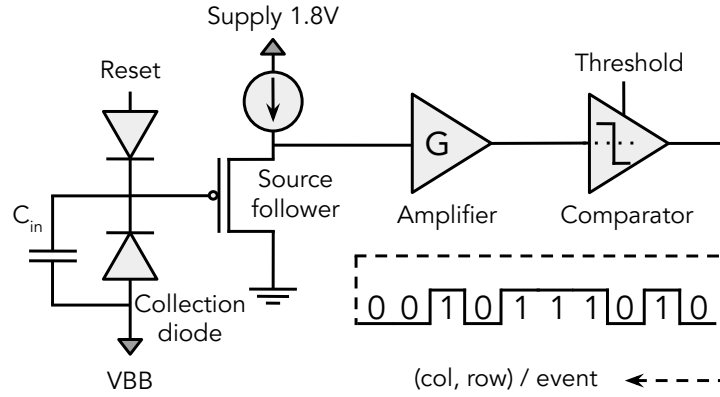


Figure 1.13 – Typical front-end electronics found inside a pixel. The incident particle induces a signal in the sensor. The pn-junction formed diode collects the signal, which is processed by the preamplifier. A source follower, is a type of transistor amplifier circuit used in some of the pixel sensors. Its primary role is to buffer and transfer signals between stages without altering the amplitude of the signal significantly. The comparator discriminates the signal before it is digitized by the ADC and read out.

The induced signal is fed into an amplifier to boost the signal to a level suitable for further processing. After amplification, the signal is sent to a comparator. If the signal exceeds a user-configurable threshold set by the comparator, the traversing particle is registered as a hit. In digital sensors, the digitized information (hit/no-hit) is then sent out, while some analogue designs retain the full or part of the analogue signal.

In many modern silicon pixel detectors, not only the two-dimensional pixel address (indicating the column and row number of the hit inside the matrix) is recorded, but additional information such as a hit timestamp or time-of-arrival (ToA) or time-over-threshold (ToT, a measure of the amount of charge) are also detected. However, the sensors discussed in this thesis record only the column and row information upon receiving an external trigger. The specialized logic in the readout electronics converts signals into a digital data stream. This stream is then captured by the data acquisition (DAQ) system and stored for later analysis offline.

1.2.2.5 Radiation damage

Radiation damage in silicon pixel sensors arises primarily from the interaction of high-energy particles with the sensor material, leading to structural changes that can significantly affect the performance of these devices. This damage, which is quantified by fluence — the integrated particle flux per unit area — is generally classified into two main types: ionizing

(surface) and non-ionizing (bulk) radiation damage. Bulk damage impacts the ability of the silicon sensor to generate signals, while surface damage predominantly affects the performance of the CMOS electronics situated on the top side of the sensor.

Ionizing radiation damage occurs when charged particles ionize atoms within the sensor material, particularly at the silicon-silicon dioxide (Si-SiO₂) interface⁴, creating electron-hole pairs. While electrons may escape the oxide layer, the less mobile holes often become trapped, resulting in a positive charge accumulation. This trapped charge can alter the electrical characteristics of the sensor, affecting its performance. Smaller CMOS nodes have smaller MOSFET gate oxide thicknesses, bringing along an improved, intrinsic radiation hardness of the integrated circuits. Nonetheless, the way this oxide growth is done has the biggest influence on the initial level of impurities. This means different manufacturers can bring along different radiation performance. The extent of ionizing damage is measured by the total ionizing dose (TID) in Gray (Gy). Although some of this damage can be mitigated by annealing—heating the sensor to repair the material—ionizing damage remains a significant challenge for the long-term reliability of silicon detectors.

Non-ionizing radiation damage, on the other hand, involves the displacement of atoms within the silicon lattice. High-energy particles can knock silicon atoms out of their lattice positions, creating vacancies and interstitials. These defects can form clusters, especially under exposure to particles like neutrons, protons, and pions, which can impart enough energy to cause multiple local displacements. Unlike ionizing damage, NIEL-induced defects are largely irreversible and can significantly degrade the ability of the sensor to generate and collect charge. The extent of this damage is typically scaled relative to the damage caused by a 1 MeV neutron flux, with particle fluences expressed in units of 1 MeV neutron-equivalent per square centimeter (1 MeV n_{eq}/cm²).

In the context of HEP experiments, such as those at the LHC, both ionizing and non-ionizing radiation damage are critical considerations. Silicon pixel detectors are located close to the interaction points where particle collisions occur, exposing them to intense radiation fields. Over time, this radiation exposure can lead to a decrease in charge collection efficiency, an increase in leakage current, and changes in the depletion voltage required to operate the sensors effectively. These effects necessitate careful design and material selection to enhance the radiation hardness of the sensors. Continued research and advancements in material science and microfabrication techniques are essential to ensure these sensors can withstand the harsh conditions encountered in hadron collider experiments.

⁴the SiO₂ is used as passivation layer on silicon sensors and as gate oxide in MOSFET transistors

1.3 Pixel sensors

Pixel sensors have become an integral component across a broad spectrum of applications, ranging from consumer electronics to HEP experiments. At their core, pixel sensors are arrays of tiny, discrete detectors — each known as a pixel — that individually register the presence of particles. Their ability to detect and precisely localize events across a two-dimensional matrix plane makes pixel sensors invaluable in scenarios where detailed spatial information is critical.

In HEP pixel sensors play a crucial role in tracking the paths of particles generated in collisions. The development of these sensors in HEP has been primarily driven by the need to measure short-lived particles through their displaced vertex (with respect to the primary interaction vertex) and handle the increasing interaction rates in modern accelerators.

The fine granularity of pixel sensors enables the precise reconstruction of particle trajectories, allowing physicists to study the properties of subatomic particles with unprecedented accuracy. This precise tracking is essential for identifying particles, measuring their momentum, and determining their origin — tasks that are critical for testing theoretical models and searching for new physics phenomena.

Beyond particle physics, pixel sensors have also revolutionized imaging technologies. In medical imaging, for instance, they are used in X-ray detectors, providing high-resolution images that are crucial for accurate diagnoses. Their ability to function in various environments, including high-radiation areas, makes them suitable for space missions, where they can be used to image distant celestial objects or monitor cosmic radiation. In consumer electronics, pixel sensors are the heart of digital cameras, capturing light to produce the images we see on our devices.

The versatility and adaptability of pixel sensors further enhance their importance. Advances in semiconductor technology have led to the development of various types of pixel sensors, with hybrid or monolithic design. Each type has its strengths, including lower power consumption, higher sensitivity, or better integration with readout electronics. These developments continue to push the boundaries of what pixel sensors can achieve, opening up new possibilities in research and technology across multiple fields.

Pixel sensors are indispensable in modern technology and research due to their ability to provide detailed spatial resolution, adaptability across different applications, and continuous improvement through advances in semiconductor fabrication. Whether in the pursuit of fundamental physics, or the development of consumer electronics, pixel sensors are central to capturing the intricate details of the world around us.

1.3.1 Hybrid vs monolithic design

Two major types of pixel detectors are prevalent in HEP today: hybrid and monolithic devices. In hybrid detectors, the pixelated sensor and readout electronics are separate components, whereas in monolithic detectors, both parts are integrated into a single device.

Hybrid pixel sensors represent a well-established technology in particle physics and other high-performance imaging applications. The hybrid design involves the separate fabrication of the sensor layer, where particle detection occurs, and the readout electronics, which processes the signal. These two components are then interconnected using bump bonding — a process where tiny metallic bumps on the sensor layer are aligned and fused with corresponding bumps on the readout chip.

The primary advantage of hybrid pixel sensors is the ability to optimize both the sensor and the readout electronics independently. The sensor layer, usually made from silicon, can be tailored for specific radiation types or energy ranges, while the readout chip can be fabricated using advanced CMOS technologies that offer high processing speed and low power consumption. This separation allows for the use of high-resistivity materials in the sensor, enhancing charge collection efficiency and radiation hardness.

However, the hybrid approach has its drawbacks. The bump bonding process is complex and expensive, and it limits pixel size, as the bump size sets a lower bound on how small the pixels can be. Additionally, this process introduces extra material between the sensor and readout electronics, which can be a disadvantage in applications requiring minimal material thickness, such as in HEP experiments where minimizing multiple scattering is crucial.

Monolithic Active Pixel Sensors (MAPS), on the other hand, offer a more integrated approach, with both the sensor and the readout electronics fabricated on the same silicon wafer ⁵ using standard CMOS processes. In this design, the sensor layer and readout circuitry are monolithically integrated, meaning they are part of the same silicon die ⁶. This integration eliminates the need for bump bonding, allowing for smaller pixel sizes and reducing the overall material budget.

One of the key advantages of the monolithic design is its potential for high spatial resolution, as the pixel size can be significantly reduced compared to hybrid sensors. This is particularly beneficial in applications requiring detailed imaging, such as biomedical imaging or vertex detectors in particle physics experiments. Additionally, the monolithic design

⁵A wafer is a circular, thin slice of silicon, used as the substrate for fabricating integrated circuits. It was cut from an silicon ingot. Usually, when referring to wafers, various authors can also refer to fully processed wafers, on which all the integrated circuits are already placed.

⁶In the context of integrated circuits, a die represents a small piece of silicon, that contains a complete functional circuit and which is diced from a full wafer

simplifies the fabrication process, leading to potential cost savings, especially in large-scale production.

Monolithic detectors also offer a reduced material budget, which is crucial in tracking detectors for particle physics. A thinner detector reduces multiple scattering of particles, leading to more accurate trajectory reconstruction. Furthermore, with the readout electronics integrated with the sensor, MAPS can provide faster signal processing and reduced power consumption, making them suitable for applications with stringent timing requirements.

Power consumption in silicon detectors comes from four primary sources: the sensor (especially after it has been irradiated), the analog circuitry, the digital components, and the energy needed to transmit data off-chip and beyond the detector. Monolithic detector designs show potential for reducing power in the sensor and analog sections [23]. Simplifying in-pixel circuitry could also lower digital power usage. In future HEP experiments, the energy required to move data off-chip and away from the detector could become the most significant factor.

However, the monolithic approach also presents challenges. The performance of the sensor layer is inherently linked to the CMOS process used for the readout electronics, which may not always allow for high-resistivity substrates ideal for charge collection. This can affect the detection efficiency and radiation hardness. Additionally, integrating all components on a single chip can complicate the design and limit the ability to optimize the sensor and readout electronics separately. An in-depth discussion on the challenges and perspectives of CMOS MAPS in HEP is provided in [24].

Both hybrid and monolithic designs offer distinct advantages and trade-offs in pixel sensor technology. Hybrid sensors provide the flexibility to optimize the sensor and readout electronics separately, making them ideal for applications requiring high performance and radiation tolerance. In contrast, monolithic sensors offer advantages in spatial resolution, material budget, and potentially lower production costs, making them increasingly attractive for a wide range of applications. The choice between hybrid and monolithic designs ultimately depends on the specific requirements of the application, such as pixel size, radiation environment, and cost considerations. As technology advances, both approaches continue to evolve, pushing the boundaries of what is possible in pixel sensor technology.

1.3.2 High-Voltage (HV) and High-Resistivity (HR) CMOS MAPS

Among the various technologies developed to enhance pixel sensor performance, high-voltage and high-resistivity CMOS MAPS are particularly noteworthy for their significant contributions. These technologies have been at the forefront of recent advancements in HEP

vertexers, enabling CMOS MAPS to achieve superior charge collection efficiency, improved radiation hardness, and very good spatial resolution.

High-voltage CMOS MAPS are designed to operate at higher voltages than standard CMOS sensors, which typically operate at lower voltages due to the limitations of conventional CMOS processes. By increasing the operating voltage, high-voltage CMOS MAPS achieve several critical advantages. Firstly, the higher voltage allows for a stronger electric field within the sensor, enhancing the speed and efficiency of charge collection. This is especially important in environments requiring rapid response times, such as high-rate particle detectors.

Moreover, the stronger electric field in high-voltage CMOS MAPS facilitates full depletion of the active area of the sensor, leading to improved charge collection from deeper within the silicon. Full depletion ensures that even charge carriers generated far from the collection node are efficiently swept towards it, reducing signal loss and improving the overall sensitivity and dynamic range of the sensor. Operating at higher voltages also enhances radiation hardness, allowing the sensors to maintain performance even after exposure to significant radiation levels, which is crucial for long-term use in harsh environments. This choice is being investigated or pursued by collaborations such as Mu3e [25], LHCb [26], and ATLAS [27].

Most CMOS processes featuring an epitaxial layer typically involve low resistivity, usually in the order of $10\ \Omega\text{ cm}$ to $100\ \Omega\text{ cm}$. While the depletion depth in such a standard resistivity epitaxial layer is a fraction of a micrometer, it may reach several micrometers in a high-resistivity epitaxial layer, making it comparable to the pixel dimensions. This increased depletion depth significantly enhances the signal-to-noise ratio, making high-resistivity epitaxial layers particularly attractive for MAPS.

High-resistivity silicon refers to silicon with a low concentration of free charge carriers in its intrinsic state, achieved through careful control of the doping process. This low concentration of free carriers results in a material with high electrical resistivity, beneficial for several reasons.

The resistivity of the silicon in which the CMOS pixel is placed defines the depth of the depletion region into the epitaxial layer formed by the n-type implant creating the diode (for a given bias voltage). When electron-hole pairs are generated within silicon by MIPs, the electrons typically diffuse through the epitaxial layer. If they are sufficiently close to the depletion region of the diode, they will be collected. If generated far from the diode, they will travel within the epitaxial layer for longer distances than those generated close to the diode and ultimately recombine, degrading the signal magnitude collected by the pixel the MIP passed through.

In the ideal scenario, the entire epitaxial layer underneath the diode would be completely depleted, changing the main charge transport mechanism from diffusion to drift, where the increased electric fields from the larger depletion region attract more charge than in the case of a smaller depletion region. As the depletion region width increases with the resistivity of the epitaxial layer, one way to extend the depletion region further into the epitaxial layer and improve charge collection efficiency is to use an epitaxial layer with a high resistivity between 1 k Ω cm to 10 k Ω cm. The use of a high-resistivity epitaxial layer should increase charge collection efficiency. Moreover, changing from a diffusion dominated to a drift dominated charge collection, the minority carrier lifetime degradation is mitigated by the larger velocity of the carriers under a stronger electric field. This enhancement attempts to make high-resistivity CMOS MAPS also suitable for environments with high radiation exposure.

This approach has been adopted by various high-energy physics collaborations, such as STAR at RHIC [28] and ALICE [29], and is being considered for future upgrades like CBM at FAIR [30] and DECAL at the ILC [31].

Both CMOS technologies show great potential for the most demanding applications in high-energy physics. The choice between these technologies will depend on factors such as the achievable depletion depth, availability, cost, production yield and maturity of the process. Addressing the associated challenges is crucial to fully realizing the benefits of these advanced pixel sensor designs.

In HV and HR CMOS MAPS the concept of fill factor plays a crucial role in determining the performance of the detectors. The fill factor refers to the fraction of the pixel area that is sensitive to charge collection, which is influenced by the design of the charge collection node and the placement of the CMOS electronics.

A large fill factor, usually chosen for HV MAPS design, means that a larger portion of the pixel area is dedicated to charge collection, which can enhance the radiation hardness of the detector and improve the efficiency of charge collection. This configuration typically involves embedding the entire readout electronics within the charge collection well, which ensures that the majority of the charge signal generated by incident particles is effectively captured and processed. However, the downside of this approach is that it leads to larger sensor capacitance, which can introduce penalties in terms of noise and power consumption.

On the other hand, a small fill factor design like the one in most HR-CMOS MAPS involves placing the CMOS electronics outside the charge collection well (see Fig. 1.14). This design significantly reduces the sensor capacitance, leading to lower power consumption and reduced




Figure 1.14 – (left) Large fill factor CMOS MAPS design. A larger sensor capacitance forms. (right) Small fill factor CMOS MAPS design. Capacitance as small as a few fF can be achieved. Adapted from [32].

noise, which is advantageous for high-resolution and low-power applications. However, the smaller fill factor also results in longer drift distances for charge carriers, which can negatively impact radiation hardness and the efficiency of charge collection, especially in high-radiation environments. As such, the choice between a large and small fill factor involves trade-offs between radiation tolerance, power efficiency, and noise performance, and must be carefully considered based on the specific requirements of the application.

1.3.3 Modified process for HR CMOS MAPS

The ITS2 upgrade has integrated HR MAPS using the standard TowerJazz 180 nm CMOS imaging sensor process. A critical enhancement in this process was the inclusion of a deep p-well, which permits full CMOS circuitry within each pixel, enabling more sophisticated readout functionality. However, despite the application of moderate reverse substrate bias, the sensitive layer of the sensor is not fully depleted, leading to charge collection that is predominantly governed by diffusion, an intrinsically slow process.

In preparation for the ITS3 upgrade and the transition to the 65 nm CMOS technology node, significant efforts have been made to enable full depletion of the epitaxial layer through various process modifications [33]. The motivation behind these changes is the thinner epitaxial layer in the 65 nm technology ($10\ \mu\text{m}$ compared to $25\ \mu\text{m}$ in the 180 nm process), which inherently restricts the extent of depletion.

A major breakthrough in this development was the introduction of a deep, low-dose n-type implant beneath the deep p-well, covering the entire pixel array (see the middle panel of Fig. 1.15). This modification effectively shifts the pn-junction from the diode deep into the epitaxial layer. Unlike the standard process where depletion began from a small region starting at the n-well diode interface, the new process initiates depletion from the uniform

n-type blanket that spans the entire pixel area. As a result, the electric field is now similar to that found in large fill-factor MAPS. With this design, minimal reverse bias is sufficient to fully deplete the sensor [34].

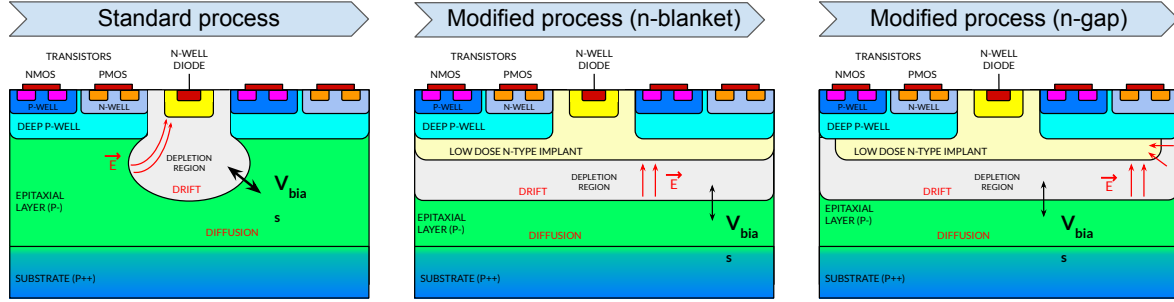


Figure 1.15 – Cross-sections through one pixel in the matrix of the three different process designs implemented. Adapted from [33].

This modification ensures that charge collection is now governed by drift rather than diffusion. Fully depleting the epitaxial layer is essential for efficient charge collection, leading to a significantly improved signal-to-noise ratio. Moreover, with the sensor now dominated by drift, charge collection is one to two orders of magnitude faster, achieving timing resolutions on the order of a few nanoseconds [35]. This improved timing reduces the likelihood of signal overlap in high-rate environments, thereby enhancing the accuracy of particle tracking. Additionally, drift-dominated charge collection is less vulnerable to radiation-induced degradation, making the sensor more robust in high-radiation environments.

To further enhance the lateral field near the pixel corners, where it is typically the lowest, an additional modification was introduced where a gap in the deep low-dose n-type implant was added (see the last panel of Fig. 1.15).

With these changes to the design of the sensors, the tolerance to NIEL damage was increased from the 10^{13} 1 MeV n_{eq}/cm^2 level [33], at which the ALPIDE chip was qualified, to levels as high as 10^{15} 1 MeV n_{eq}/cm^2 [34], although with some performance loss. Additionally, the TID tolerance has been improved from 27 kGy for the ALPIDE sensor [36] to 3 MGy [34].

These process modifications are implemented without requiring substantial changes to the pixel circuit design or layout, offering flexibility in fabrication. This means that the same pixel design can be produced using either the standard or the modified process, depending on the specific requirements of the application. This adaptability is particularly advantageous for large-scale experiments where sensors need to be mass-produced while maintaining high performance standards.

1.4 ALICE upgrades

Since its inception, ALICE has been instrumental in advancing our understanding of QCD through the study of heavy-ion and proton collisions. During Run 1 (2009–2013) and Run 2 (2015–2018) ALICE allowed researchers to push the energy frontier of ultra-relativistic heavy-ion collisions far beyond what had previously been achieved in earlier experiments at facilities like the Alternating Gradient Synchrotron (AGS), the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory, and the SPS at CERN. These experiments established the basic picture of the quark-gluon plasma (QGP) as an almost-perfect liquid, from its observations of the hot matter reaching new extremes in temperature, density, and volume during the heavy-ion collisions [37].

As part of a comprehensive overhaul of its subdetectors, making them fit for a continued scientific journey at higher luminosity, plans were set in motion to upgrade the ALICE detector during Long Shutdown 2 (LS2) [38, 39]. The upgrade aimed to enhance the capabilities to further explore the QGP, focusing on the physics of heavy quarks (i.e. charm and beauty), assessing thermal photon and low-mass dilepton production, and evaluating the production of hypernuclear states, among other objectives.

To accomplish this, it was essential to conduct measurements with both high precision and high statistical significance. The end goal was to record Pb-Pb collisions at a rate of 50 kHz, an increase by an order of magnitude with respect to Runs 1 and 2. Moreover, the goal is to collect 13 nb^{-1} from Pb-Pb collisions in Runs 3 and 4, from the 1 nb^{-1} collected during Runs 1 and 2. These upgrades were designed to improve vertexing and tracking capabilities, while retaining the unique particle identification capabilities of ALICE, all key elements for accessing rare physics channels, particularly in the challenging low transverse momentum region.

Key upgrades include the new ITS2 and Muon Forward Tracker, both utilizing HR CMOS MAPS, as well as new Gas Electron Multiplier-based readout chambers for the Time Projection Chamber. Moreover, all existing detectors have been upgraded to support a higher readout rate, with some now operating in continuous mode. With these upgrades from LS2, ALICE is now poised to take full advantage of the improved performance of the LHC.

In the past years preparations started and are now underway for further detector upgrades in the next Long Shutdown (LS3, 2025–2027). An R&D program has been started to develop a novel system of wafer-scale silicon sensors that can be bent to form truly cylindrical tracking layers to replace the three inner layers of the current ITS2. The program also aims to reduce power consumption to the point where air cooling can be used to further reduce the material

budget to approximately 0.05% X_0 per layer. Together with the reduced radius of the first layer to about 19 mm, this also results in a significant improvement of the vertexing capabilities.

Further down the line, post LS4 (2033-2034), a completely new, high-speed experimental apparatus, called ALICE 3 is planned. It will have very high resolution and ultra-low mass, with detection layers positioned very close to the interaction point and will use advanced silicon detectors. At its core there will be an all-silicon tracker, with the innermost layers located inside the LHC beam pipe, providing a very large rapidity acceptance and complemented with very good particle identification capabilities over a very broad range of transverse momenta.

1.4.1 ITS2

A key aspect of the LS2 upgrade was the installation of the ITS2 - a new Inner Tracking System characterized by high resolution and reduced material budget [40]. This upgrade was expected to significantly enhance vertexing and tracking capabilities, especially for particles with low transverse momentum.

Together with a reduction in the beampipe radius (as shown in Table 1.1), the integration of this highly granular and lightweight inner tracker, positioned closer to the interaction point, significantly enhances pointing accuracy. The ITS is instrumental in reconstructing primary and secondary vertices, as well as boosting the angular and momentum resolution for particles detected in the TPC.

Table 1.1 – Parameters of the central section of the beryllium beampipe used in ALICE [39, 40].

Parameters	Run 1 & 2 (up to 2018)	Run 3 (2022 – 2025)	Run 4 (2029 – 2032)
Outer radius (mm)	29.8	19	16.5
Thickness (mm)	0.8	0.8	0.5
Material budget (% x/X_0)	0.23	0.23	0.14

To enhance the impact parameter resolution and boost tracking performance and momentum resolution, the material budget of the first detection point had to be significantly reduced. This necessitated a technological shift in the ITS upgrade. Previously, the ITS used three silicon technologies: the first two layers comprised hybrid Silicon Pixel Detectors (SPD), the middle layers Silicon Drift Detectors (SDD), and the outer layers, with lower track density, utilized double-sided Silicon microStrip Detectors (SSD). The outer four layers featured analogue readout and offered standalone particle identification (PID) through dE/dx measurements.

For the ITS2, the decision was to use seven layers of silicon pixel detectors with digital readout, concentrically arranged around the beam pipe (see Fig. 1.16), thereby increasing the number of measurement points along particle trajectories. Monolithic Active Pixel Sensors (MAPS), which integrate both the sensor and read-out electronics into a single device, were selected for their high granularity and minimal material thickness.




Figure 1.16 – The ITS2 installation: The bottom half of the IB has been inserted around the beam pipe. The first layer of the outer barrel can be seen with the golden color. Adapted from [40, 41].

The flagship pixel sensor of the ITS2 is the ALPIDE chip, a CMOS MAPS fabricated in the 180 nm technology node (see Section 1.4.2). The chip has total dimensions of 30 mm × 15 mm, and includes the active matrix with almost square pixels and a digital periphery. Owing to the monolithic design, the chips are notably thin, with silicon sensors in the inner and outer barrel thinned to 50 μm and 100 μm , respectively, aligning with material budget requirements. The finer segmentation of pixels in ITS2, critical for small intrinsic position resolution and maintaining low occupancy, enables excellent spatial resolution ($\sim 5 \mu\text{m}$) across a wide operational range.

Some of the differences between the former ITS and the ITS2 are shown in Table 1.2.

The upgraded ITS2 significantly enhances our ability to measure heavy-quark hadrons with higher precision and efficiency. This will allow the study of heavy quark dynamics within the QGP. Its improved track reconstruction efficiency and considerably better pointing

Table 1.2 – Difference in select parameters between ITS and ITS2 [40].

	ITS	ITS2
Material budget (% x/X_0) per layer	SPD: 1.14 / 1.14 SDD: 1.13 / 1.26 SSD: 0.83 / 0.86	IB: 0.36 OB: 1.1
Number of layers	6	7
Pixel size ($r\varphi \times z$) (μm^2)	50×425 (SPD)	26.88×29.24
Spatial resolution ($r\varphi \times z$) (μm^2)	12×100	5×5
Distance to the innermost layer from the IP (mm)	39	23

resolution ⁷ facilitates the reconstruction of secondary vertices from weakly decays of charm and beauty hadrons. This is especially important in the low p_T regime, which is essential for accurately determining the total charm and beauty quark production cross-section – measurements that are not yet available in heavy-ion collisions at the LHC.

In ITS1, the pointing resolution in the transverse plane ($r\varphi$) was moderate, but it was severely limited in the longitudinal direction (z), due to the large asymmetry in the pixel pitch ($50 \mu\text{m} \times 425 \mu\text{m}$). The upgraded ITS2 addresses these challenges with a much finer pixel granularity of $29.24 \mu\text{m} \times 26.88 \mu\text{m}$, resulting in a threefold improvement in $r\varphi$ resolution and a fivefold improvement in z resolution [42] (see Fig. 1.17). This enhancement enables precise three-dimensional vertex reconstruction, a significant step forward from the two-dimensional plane projections used during ITS1.

The impact of these improvements is particularly evident in the context of heavy-flavor physics. For example, the decay of the Λ_c baryon, which has a decay length in the order of $60 \mu\text{m}$, posed a significant challenge in ITS1. The poor pointing resolution in ITS1 meant that Λ_c decay tracks were often indistinguishable from those originating from the primary vertex, especially at low p_T ($< 1 \text{ GeV}/c$), where the pointing resolution was larger than the Λ_c decay length.

⁷Defined as the width of the distribution of the distance of closest approach (DCA) of the primary tracks to the interaction vertex; it is an important parameter to evaluate the capability of a vertex detector to separate secondary vertices (from weak decays of heavy flavor particles) from the interaction vertex

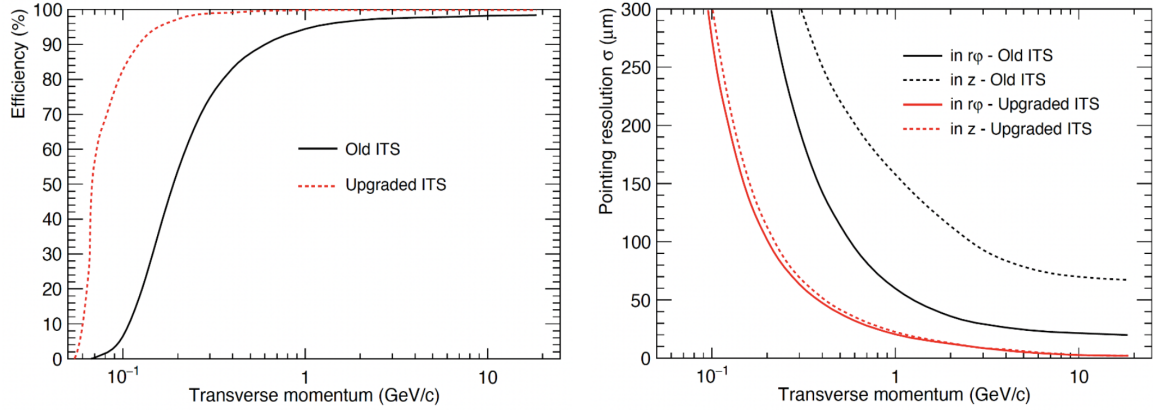


Figure 1.17 – **(left)** Tracking efficiency as a function of the particle transverse momentum for the upgraded ITS 2 (red) compared with the previous ITS (black). **(right)** Pointing resolution to the vertex of charged pions as a function of the transverse momentum for the ITS 2 (red) and the previous ITS (black). Adapted from [43].

While the production of Λ_c baryons has been measured down to very low p_T with ITS1 in Runs 1 and 2 [44], these measurements were conducted in a regime characterized by moderate track reconstruction efficiency and poor primary/secondary vertex separation, making it nearly impossible to reconstruct the secondary vertex accurately, resulting in the inability to distinguish between signal and background.

The improved pointing resolution of the ITS2 allows for the separation of tracks coming from the weak decay vertex of the Λ_c from soft particle production from the primary vertex. This capability will significantly enhance the accuracy of measurements in the low p_T range, reducing uncertainties and improving the overall precision of heavy-flavor studies.

The importance of precise secondary vertex reconstruction and effective separation of prompt and non-prompt charm components also extends to the study of beauty-containing hadrons.

The changes in ITS2 will thus enable more accurate measurements of both charm and beauty quarks, contributing to a more detailed understanding of their interactions with the QGP. Coupled with the unprecedented precision achieved with the data samples planned to be collected during Run 3, the ITS2 will be instrumental in advancing the field of heavy-ion physics and deepening our understanding of the QGP.

1.4.2 ALPIDE

The ALPIDE sensor, the CMOS MAPS used in the ALICE ITS2 upgrade, is developed using the TowerJazz 180 nm CMOS imaging sensor process. This advanced process allows for the

implantation a deep p-well, effectively isolating all n-wells, except the collection diodes, from the sensitive volume, which is the epitaxial layer [29, 45].

The ALPIDE sensor is composed of three fundamental layers:

- The substrate: a highly p-doped (p++) substrate serves as the mechanical backbone of the sensor
- The epitaxial layer: This thin, p-type epitaxial layer, just 25 μm thick constitutes the sensitive volume where charge collection occurs
- Implants and circuitry: On top of the epitaxial layer, 11 μm thick [46] CMOS implants form the transistor logic. The n-type implants that form pn-junctions with the epitaxial layer function as collecting diodes for electrons, while the deep p-wells house the necessary circuitry. Above these implants, metal layers facilitate signal propagation and bias distribution across the chip.

This architecture, produced on a high-resistivity (greater than 1 $\text{k}\Omega\text{ cm}$) p-type epitaxial layer, is crucial for efficient charge collection and sensor operation.

The sensor measures 30 mm \times 15 mm and contains a matrix of 1024 \times 512 pixel cells (see Fig. 1.18). Key operational circuits, including analog biasing, control, hit-driven readout, and interfacing functions, are compactly implemented in a narrow periphery region of approximately 30 mm \times 1.2 mm. This design results in a substantial active area of about 30 mm \times 13.8 mm.

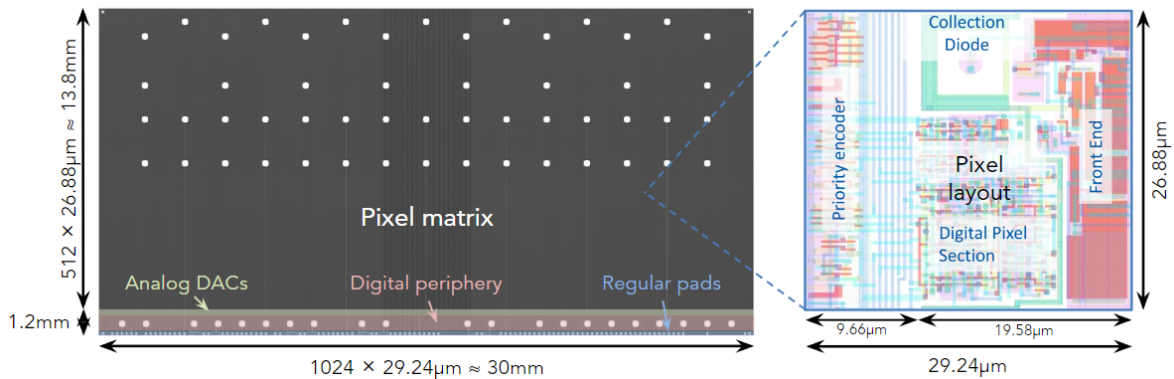


Figure 1.18 – The ALPIDE matrix dimensions and pixel sensor layout. Adapted from [47].

Each pixel within the ALPIDE sensor measures 29.24 μm \times 26.88 μm and includes a collection diode surrounded by essential readout circuitry. The pixel architecture integrates an amplifier, a discriminator, and a multi-event buffer. The multi-event buffer is a critical

digital component, housing three hit storage registers that store data until a trigger command initiates a readout.

The modest timing requirements of the ALICE experiment allow the use of a low-power (≈ 40 nW) full CMOS in-pixel frontend circuit, with a shaping time of a few microseconds and an effective input capacitance of approximately 2.5 fF.

The capacitance of the diode is a crucial parameter influencing the signal formation. The input capacitance is determined by the junction capacitance of the collection diode and the parasitic capacitance of the readout circuit and the neighbouring pixels. A lower input capacitance is desired to maximize the signal, which can be achieved by optimizing the readout circuit and applying a reverse bias.

To meet these stringent power density requirements, a novel hit-driven readout architecture with zero-suppression in the matrix has been developed. This architecture ensures that only active pixels — those that have registered a hit — are read out. Hits stored in the multi-event buffer are processed by Priority Encoders, specifically designed with an Address-Encoder and a Reset-Decoder circuit based on an arbitration-tree with priority logic. These encoders transmit the addresses (column and row) of hit pixels to the peripheral region, subsequently issuing a reset signal to the in-pixel memory element.

This hit-driven, zero-suppressed readout architecture is highly efficient, as it eliminates the need to propagate a clock signal across the matrix. Consequently, time and energy consumption are directly proportional to the number of hits, ensuring optimal performance with minimal power expenditure.

The charge collection in ALPIDE is primarily governed by two mechanisms: drift and diffusion. A small voltage of 1.8 V is applied to the collection diode, creating a small depletion region even when no reverse bias is applied. However, most of the epitaxial layer remains undepleted. To further extend the depleted region, a reverse bias voltage can be applied to the substrate. Although the maximum reverse bias before breakdown is -8 V, the operational bias is typically set at -3 V. Having a bias voltage larger than the nominal -3 V one was shown to not have any added benefits — the depletion width increases only minimally afterwards [48] and the sensor detection efficiency has only a marginal gain.

When a charged particle traverses a pixel within the ALPIDE sensor, it loses energy through ionization with both the silicon lattice and impurity atoms, resulting in the creation of electron-hole pairs along its trajectory (see Fig. 1.19). The electrons generated within the epitaxial layer undergo thermal diffusion until they encounter the stronger electric field present in the depleted region. Once within this field, the electrons are swiftly directed toward the n-well collection diode.




Figure 1.19 – Schematic cross-section of a single pixel cell in the ALPIDE sensor. At the center, an n-well implant functions as the collection diode, forming a pn-junction with the surrounding p-type epitaxial layer. This pixel structure is repeated in both directions to create the full pixel matrix. The depletion region, indicated in grey, takes on a balloon-like shape, extending only slightly into the epitaxial layer. This region defines the boundaries between the drift and diffusion zones.

When a charged particle traverses a pixel of the ALPIDE sensor, it will lose energy due to ionization with the silicon lattice and impurity atoms and will create electron-hole pairs along its path. The electrons produced in the epitaxial layer will diffuse thermally until they feel the stronger electric field in the depleted region and get swept away towards the n-well collection diode. The electrons created in the epitaxial layer are confined here because the p- – p++ junctions between the epitaxial layer the substrate or the deep p-well act as reflective barriers. Charge collection from the first few μm of the substrate is also possible [49]. Electrons originating from deeper regions of the substrate layer can potentially cross the concentration gradient barrier into the epitaxial layer. However, their lifetime in the substrate is notably short due to the high doping levels. Specifically, the carrier lifetime ranges from approximately 10 μs in the epitaxial layer to just 20 ns in the substrate [50].

The ionization process typically generates approximately 66 electron-hole pairs per μm [13, 49], resulting in the generation of about 1600 electron-hole pairs within the epitaxial layer.

The shape of the depletion region in the ALPIDE sensor is “balloon”-like, confined to the vicinity of the collection electrode, without extending too much towards the edges of the pixel. This configuration contrasts with classical planar sensors, where depletion can extend across the entire sensor thickness. For ALPIDE, the depletion radius is a critical factor, particularly given the small collection diode geometry, which requires modified calculations

for accurate modeling. Charge generated outside the depleted region primarily undergoes slow diffusion, while charge generated within the depleted region experiences faster drift (owing to the stronger electric field) in its path towards the collection diode.

The classical formula ($d \sim \sqrt{\rho V}$ [9]) used to calculate the depletion radius in planar sensors does not apply to the ALPIDE sensor due to its small collection diode setup. If used it will wrongly give a depletion larger than the thickness of the sensor since it assumes an infinite (1D-model) junction area. Instead, a more appropriate model assumes a spherical junction, where the depletion region expands in a spherical manner from the collection diode. The depletion shape in the ALPIDE can be approximated with a sphere of an inner radius R_1 (given by the doped diode itself, which does not contribute to the depletion volume) and a radius R_2 for the outside extension of the depletion into the epitaxial layer. In the case of the ALPIDE, the collection diode is an octagonal shape with a side length of approximately $1 \mu\text{m}$ [29]. The formula for the depletion radius is given by [48, 51]:

$$R = \sqrt[3]{\frac{2\epsilon_{Si}}{eN_A} \frac{3R_1}{2} (V_{bi} - V_{ext})}$$

Here, V_{bi} is the built in potential of the pn-junction ($V_{bi} = k_B T / e \cdot \ln(N_A N_D / n_i^2) \approx 0.75 \text{ V}$), where N_A and N_D are the doping concentrations of the n-well diode and the p-epitaxial layer) and V_{ext} is the applied reverse bias.

For an ALPIDE sensor that is not biased the depletion depth is around $7 \mu\text{m}$, increasing to about $10 \mu\text{m}$ at the nominal -3 V reverse bias. Despite this increase, only a small fraction of the epitaxial layer is depleted (approximately 5% at 0 V and about 13% at -3 V), highlighting that charge collection is predominantly driven by diffusion rather than drift.

The sensor has demonstrated exceptional performance. Under nominal operating conditions, it consistently achieves a detection efficiency exceeding 99%, with noise rates significantly lower than 10^{-6} per event per pixel. The sensor also benefits from an intrinsic spatial resolution of approximately $5 \mu\text{m}$, largely due to the effects of charge sharing [42].

1.4.3 ITS3

As it currently stands, the ITS2 12.5 Gpx detector marks a significant advancement in CMOS detector technology, offering high integration density, low power consumption, and a reduced material budget compared to previous MAPS generations.

Nonetheless, recent years have seen remarkable progress in CMOS imaging technology, particularly in MAPS fabrication, driven largely by commercial growth. These developments

have notably improved particle detectors, enabling greater miniaturization and the use of larger wafers, which enhance functionality while further reducing the material budget.

Aligning to these CMOS innovations, the ITS collaboration initiated an R&D program towards a next-generation inner-tracker, known as the ITS3 [52], which will replace the three innermost layers of the current ITS2 during LS3.

The new detector design will feature wafer-scale silicon sensors as large as $26\text{ cm} \times 10\text{ cm}$, thinned to less than $50\text{ }\mu\text{m}$ and achieving an impressively low material budget of $0.07\% X_0$ per layer [53]. The sensors are curved into truly cylindrical layers, secured in place using carbon foam spacers, that also help define and maintain the design radii, and cooled by forced airflow from the periphery. The pixel pitch, still under investigation, is expected to be in the order of $20.8\text{ }\mu\text{m} \times 22.8\text{ }\mu\text{m}$, based on the next sensor iteration [54].

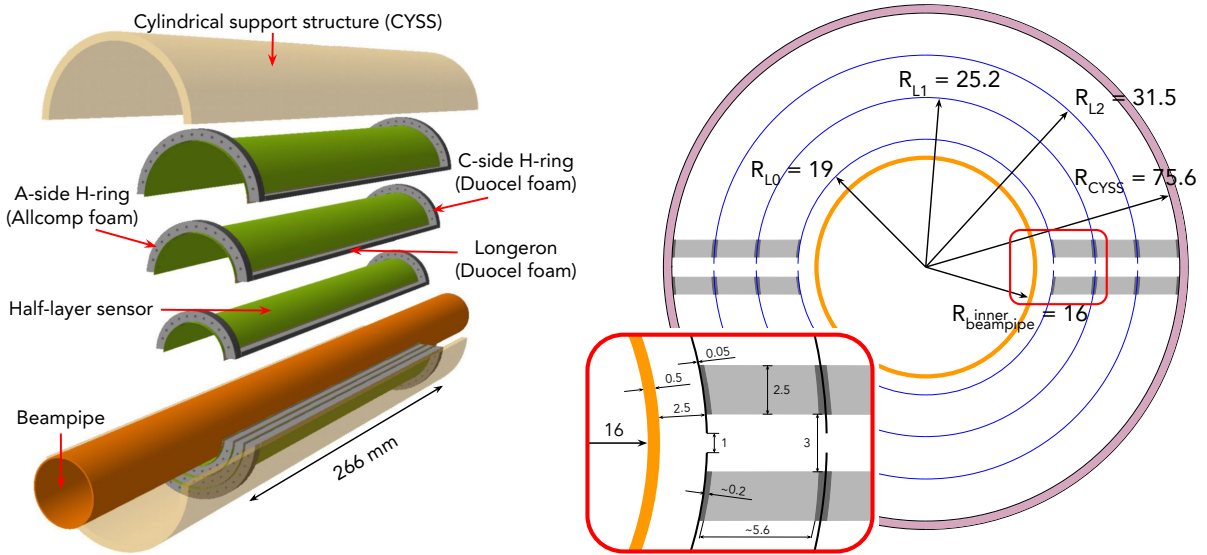


Figure 1.20 – (left) Exploded view of the ITS3 detector wrapped around the beam pipe. Image adapted from [55]. The sensors (shown in green) are supported by carbon foam longerons within the active area and half-rings at the periphery. The entire assembly is enclosed within a cylindrical support structure. **(right)** Cross-sectional view of the radial layout of the detector. Dimensions are provided in millimeters. The zoomed-in box highlights additional details, including a 1 mm mechanical gap between the top and bottom half-layers, as well as visible details of the longerons. The darker grey area on the longerons indicates places where glue infiltrates the carbon foam.

With the introduction of a new, smaller, and thinner beam pipe (see Table 1.1), the first layer of the detector can be positioned closer to the interaction point. The new beryllium beam pipe, with a thickness of $500\text{ }\mu\text{m}$ and a radius of 16 mm, allows the first layer to be placed as close as 19 mm to the interaction point (see Fig. 1.20), with the subsequent layers at 25.2 mm and 31.5 mm. The sensors have uniform dimensions along the z-axis but differ in

azimuthal dimensions due to the increasing radii. A 1 mm gap at the equator separates the top and bottom half-layers.

The entire setup is enclosed within a cylindrical support structure (CYSS), similar to that of the ITS2, which acts as an exoskeleton to provide structural support and prevent external loads from being transferred to the setup. The CYSS of the ITS2 is a lightweight sandwich structure, composed of two carbon fiber skins separated by an Airex core ($X_0 = 1380$ cm) [40]. The baseline mechanical support design for ITS3 is based on a similar structure, with variants using carbon-fiber-reinforced polymers currently under evaluation [55].

Table 1.3 captures some of the differences between the IB of ITS2 and the future ITS3.

Table 1.3 – Difference in select parameters between ITS2 and ITS3 [55].

	ITS2 (IB)	ITS3
Material budget (% x/X_0) per layer	0.36	0.07
Total number of sensors	432	6
Sensor thickness (μm)	50	$\lesssim 50$
Sensor length ($r\phi \times z$) (mm)	15×30	Layer 0 / 1 / 2 $58.7 / 78.3 / 97.8 \times 266$
Pixel size ($r\phi \times z$) (μm)	26.88×29.24	$\mathcal{O}(20 \times 22.5)$
Spatial resolution ($r\phi \times z$) (μm)	5×5	$\lesssim 5 \times 5$
Distance to the innermost layer from the IP (mm)	23	19

The technological advancements in the ITS3 are built on three core pillars: fully integrated wafer-scale MAPS produced using the 65 nm CMOS node, the bending of silicon sensors, and the development of ultra-light support and cooling structures. Each of these innovations required extensive R&D and marks a significant evolution from the design and capabilities of ITS2.

One of the pioneering aspects of ITS3 is the transition to a more advanced submicron node compared to the ALPIDE process. The new sensor will be fabricated using the 65 nm technology node from TPSCo⁸, leveraging the benefits of a smaller feature size. This transition allows for the use of larger wafers, specifically 300 mm in diameter, compared to the 200 mm wafers used in ITS2. The smaller feature size also enables more circuitry to be integrated at lower power consumption within the sensor, improving functionality and allowing for more

⁸Tower Partners Semiconductor Co., Ltd. (TPSCo)

sophisticated data processing directly within the sensor. This increase in wafer size is crucial as it facilitates the production of larger, wafer-scale sensors that can cover more area with fewer components, reducing the overall material budget and boosting the performance of the detector.

A photolithographic process known as stitching [56] is being applied for the first time to such sensor designs, enabling the creation of sensors as large as the wafer itself. Stitching is a manufacturing technique used in CMOS sensor production to achieve larger sensor sizes without requiring costly and complex photolithography equipment.

Typically, the size of a CMOS sensor is constrained by the size of the reticle, a mask used in the photolithography process to define circuit patterns on the silicon wafer. The reticle acts as a photomask in the photolithography process, where light is shone through its detailed patterns and openings to project and transfer the circuit design onto the silicon wafer coated with a photosensitive material. The exposed areas are then chemically treated, and subsequent steps such as etching, chemical vapor deposition or ion implantation ultimately define the patterns.

The wafer is incrementally shifted, and the pattern from the reticle is exposed again to create the same circuit over different areas of the wafer. This process is repeated until the entire wafer surface is patterned, creating the intricate structures of the CMOS sensor. This procedure is repeated for multiple layers, each using a different reticle, to build up the complex circuitry of the sensor. However, the size of the reticle is limited by the optics of the photolithography equipment, which can only focus light over a specific area of the wafer.

This size constraint on the reticle limits the maximum size of a single exposure, thus restricting the maximum sensor size. Stitching overcomes this limitation by combining multiple smaller exposures, or *stitches*, to create a larger, continuous sensor array on a single wafer. This process involves precisely aligning and merging the patterns from several reticle exposures to form a seamless, larger sensor. An example of how stitching is used to create a sensor similar to the one used in the ITS3 is shown in Fig. 1.21. The final sensors for ITS3 will be as large as 26 cm × 10 cm, making stitching the only viable solution for creating a single, seamless sensor, thereby reducing the need for mechanical interconnections between multiple sensors.

To qualify sensors produced in the 65 nm CMOS node for physics use, a substantial R&D effort was led by the EP R&D group at CERN, with significant support from the ALICE ITS3 team. Three key submissions were made within the 65 nm process. The first was a multi-reticle-layer run (MLR1) involving 55 different pixel test structures, both analogue and digital, including the DPTS, APTS, and CE65 sensors. The goal of this initial submission




Figure 1.21 – The principle of sensor stitching (not to scale). The design reticle contains individual sub-frames of the photomask, including a left and right endcap, as well as a repeated sensor unit housing the pixel matrix. A stitched backbone element, which integrates power, control, and data readout lines, ensures connectivity between the different regions. The sub-frames are selectively exposed after precise translation and alignment of the wafer, resulting in seamless transitions between regions. Several segments, as indicated in the figure, are later diced to produce the different wafer-scale layers of ITS3. Figure inspired by [54].

was to gain experience in the 65 nm node and optimize the charge collection properties of the sensor. The second submission was an engineering run (ER1), featuring the first large 1D-stitched sensors (MOSS and MOST). The third submission, another engineering run (ER2), targets a 2D stitched sensor known as MOSAIX.

Over the past three years, the technology has been thoroughly evaluated, and its performance parameters have been extensively investigated to determine its suitability as a particle detector for ITS3. Tests were conducted on the charge collection properties under radiation levels expected for ITS3 (as well as beyond), along with efficiency and position resolution measurements. These sensors were produced using a modified process with an n-gap (see Fig. 1.15) that enables full depletion of the sensor with minimal biasing [55]. It was discovered that the sensors could operate without reverse bias voltage, as the electric field in the epitaxial layer is sufficiently strong to ensure optimal charge collection without further biasing [55]. Currently, the first stitched sensors are available and are undergoing tests to address yield-related questions and to evaluate the performance across different regions.

The sensors will be thinned to approximately 50 μm and bent into truly cylindrical detector layers (see Fig. 1.22). The ability to bend silicon sensors into cylindrical shapes is a key

innovation that allows ITS3 to achieve its low material budget and close proximity to the beam pipe. The first layer will be positioned closer to the interaction point, at a radial distance of just 19 mm from the beam pipe, compared to the ITS2 configuration.




Figure 1.22 – Simplified radial view and differences between bottom half of ITS2 and bottom half of an engineering model (EM2) for the ITS3. The 432 ALPIDE sensors of the IB of ITS2 are replaced by six wafer-scale bent sensors in the new ITS3. Figures from [41, 55].

The bending is made possible by thinning the silicon wafers to about 50 μm , which not only reduces their mass but also makes them flexible enough to be curved without compromising their mechanical or electrical properties. This innovation was validated through extensive testing, including bending ALPIDE chips and prototypes from the 65 nm process, ensuring that the sensors maintain their performance even when curved. The author has made significant contributions to this R&D direction and the next chapter is dedicated to these bent sensors.

These bent sensors are supported by carbon foam structures and cooled using forced airflow, entirely eliminating the need for additional mechanical support structures, such as staves, or electrical connections within the active area. The carbon foam support is both lightweight and rigid, effectively holding the curved sensors in place. Cooling is achieved through forced airflow, which eliminates the need for traditional but bulkier cooling methods like water cooling. This air-cooling approach is not only lighter but also minimizes mechanical vibrations that could affect the precision of measurements. As a result, ITS3 achieves a nearly massless detector design with a material budget of just 0.07% X_0 per layer, a significant reduction compared to ITS2.

The choice of air cooling as the primary method for the ITS3 sensors is driven by the stringent material budget constraints, which prohibit the use of cooling pipes and coolant in the active area. This system is essential to maintain the sensor temperature below 30 $^{\circ}\text{C}$

and the temperature gradient within the matrix under 5 °C, key requirements for ITS3 [55] to ensure stable and reliable operation.

The mechanical and cooling concepts were validated through a series of prototypes using dummy silicon and heating elements. Detailed studies of the cooling efficiency and the induced mechanical vibrations demonstrated a large safety margin relative to the design power requirements. Full-size models confirmed the feasibility of detector integration using carbon foam structures and air cooling alone, entirely avoiding stiff carbon fiber structures and water cooling.

The proximity to the interaction point and the reduced material budget will provide ITS3 with higher efficiency for track reconstruction and a twofold improvement in impact parameter resolution in the z and $r\varphi$ directions compared to the current ITS2 (see Fig. 1.23). This enhancement is crucial for reconstructing the trajectories of low-momentum particles, which are particularly sensitive to material in their path and require high precision for accurate tracking. At large values of the transverse momentum the impact parameter resolution converges to about 4 μm , a value primarily determined by the intrinsic resolution of the sensor.

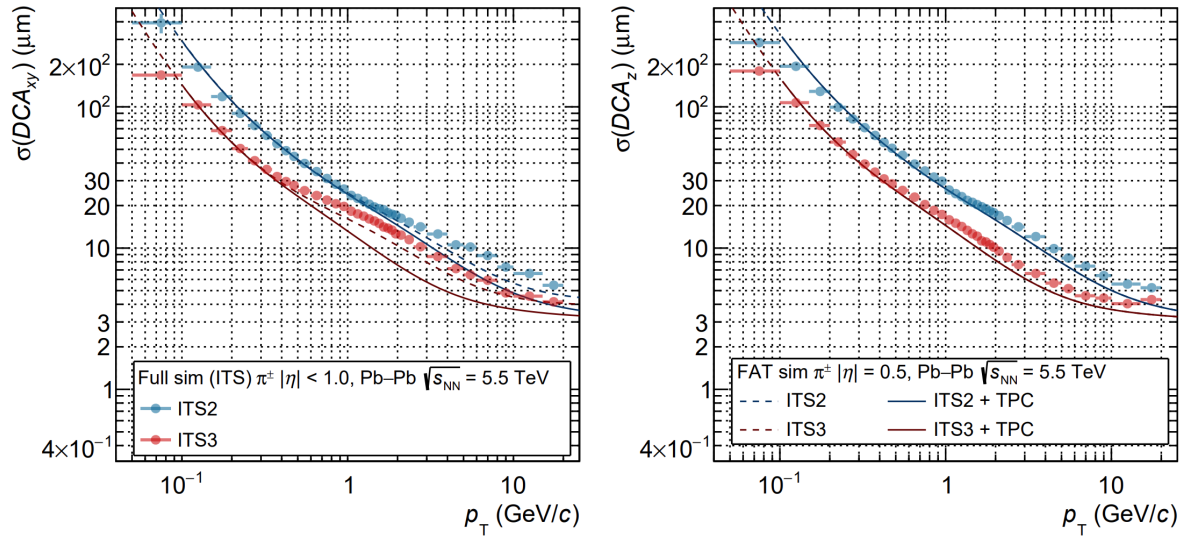


Figure 1.23 – Impact-parameter resolution in the $r\varphi$ and z direction for primary charged pions with $|\eta| < 1$ as a function of the track transverse momentum for ITS2 and ITS3 detectors. Taken from [55].

The reduction in material budget is not merely a technological achievement; it has significant implications for the physics that ITS3 and ALICE as a whole can accomplish. By minimizing the material that particles must traverse, one reduces the multiple scattering with the detector material, which in turn improves the precision of particle tracking.

The primary physics objective of ITS3 is to significantly enhance the precision in measuring heavy-flavor hadrons and dileptons, which are crucial for investigating the properties of the QGP [55]. Building on the success of ITS2, ITS3 introduces several technological and design advancements aimed at improving spatial resolution, tracking efficiency, and overall data quality, particularly for low-momentum particles (<1 GeV/ c). These enhancements open the door to exploring new physics observables [57].

The enhanced spatial resolution and tracking efficiency of ITS3, especially at low transverse momentum, will enable more precise reconstruction of decay topologies of heavy-flavor hadrons like Λ_c^+ , B_s^0 , and Ξ_c^+ . These particles have short lifetimes ($c\tau \sim 100$ μm) and small decay channel branching ratios, necessitating excellent pointing resolution to effectively resolve the combinatorial background.

For example, ITS3 is expected to reduce the statistical uncertainty in the measurement of Λ_c^+ production in central Pb-Pb collisions by up to a factor of four for $p_T < 4$ GeV/ c [55]. This improvement will enable more detailed studies of charm quark diffusion and hadronization within the QGP, providing stringent constraints to the theoretical models. Similarly, the ability to measure the B_s^0 down to $p_T = 2$ GeV/ c will offer the possibility to further study beauty quark hadronization in the QGP, assessing possible phenomena like coalescence in the beauty sector.

Moreover, one new technique that has already begun to be explored in the ITS2 and will see even greater benefits with the ITS3 upgrade is called “strangeness tracking”. The goal of this method is to determine the flight path of charged strange baryons (for example Ξ^- and Ω^-) with high spatial precision, by matching the decay particles with the clusters produced by the strange baryons themselves in the inner ITS layers. An example with the Ξ^- decay is shown in Fig. 1.24.

With this technique the strange-baryon trajectories can be extrapolated to the primary vertex with an uncertainty of few tens of microns rather than several hundreds of microns (see the right panel of Fig. 1.24). This opens the possibility to measure in heavy-ion collisions the Ξ_c^- and Ξ_c^0 baryons in decay channels that include a Ξ^- ($c\tau = 4.91$ cm), exploiting the enormous background rejection allowed by a precise reconstruction of the charm baryon decay vertex, displaced tens to hundreds micron from the primary vertex. An ideal case is the $\Xi_c^+ \rightarrow \Xi^- \pi^+ \pi^+$ decay channel. Strangeness tracking will also allow the first measurement in heavy-ion collisions of the production of the Ω_c^0 baryon, reconstructed in the $\Omega^- \pi^+$ decay channel, as well as to improve significantly the measurement of the production and lifetime of hypertriton (${}^3_\Lambda\text{H}$, $c\tau = 5.4$ cm), which also can be tracked before decaying.

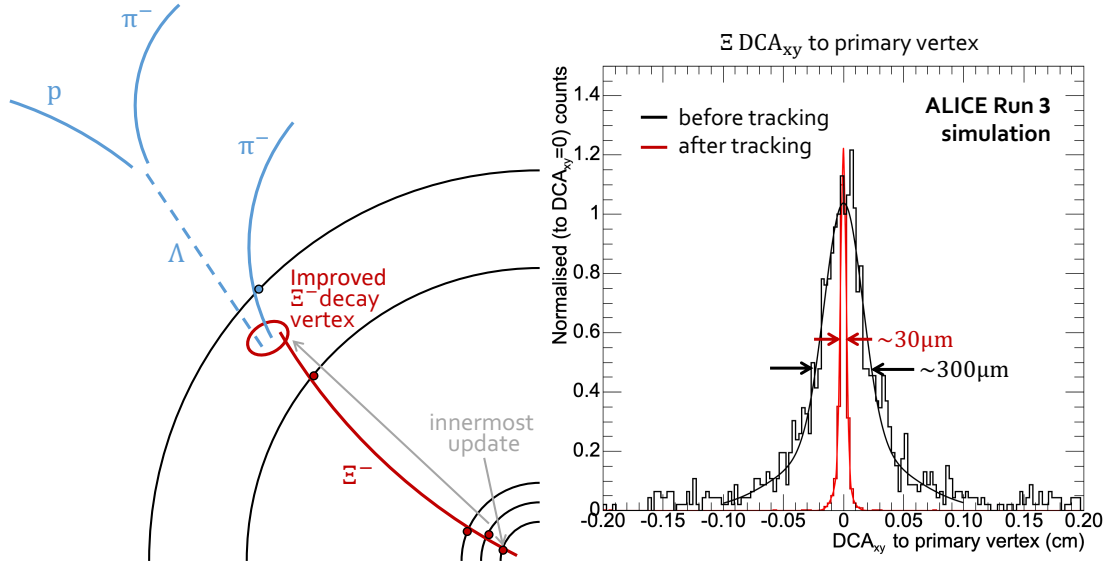


Figure 1.24 – (left) The decay and reconstruction of the Ξ^- . The ITS tracklet is shown in red and the reconstructed decay tracks are shown in blue. The Λ is neutral, so it cannot be tracked. (right) The DCA to the primary vertex of the Ξ^- is greatly improved using tracking in the innermost layers of the ITS. Figures taken from [58].

The achievable performance, in terms of strange-baryon impact parameter resolution and tracking efficiency is highly sensitive to the detector layout. Considering that Ξ^- and Ω^- on average fly for few cm only before decaying, the ITS3 setup offers an undoubtable advantage with respect to the ITS2.

Beyond heavy-flavor physics, ITS3 will enhance measurements of thermal radiation from the QGP, particularly through the study of dielectron production. The systematic uncertainty in estimating the medium temperature via the analysis of the invariant-mass spectrum of electron-positron pairs is expected to be halved with the ITS3, providing more accurate information about the temperature and evolution of the QGP, thereby deepening our understanding of this early-universe state of matter [55, 57].

The development of ITS3 required overcoming several significant challenges, particularly in validating the new 65 nm CMOS technology for high-energy physics applications. The smaller feature size and thinner epitaxial layers in the 65 nm process produce smaller signals compared to the 180 nm process, necessitating process optimizations to ensure effective charge collection. Additionally, the new technology had to be tested for radiation hardness. Extensive R&D was conducted to optimize pixel design, improve charge collection efficiency, validate the new technological node, and ensure reliable sensor operation under the challenging conditions anticipated during LHC Run 4.

The R&D phase also involved developing and testing prototypes to confirm that the sensors could be bent without losing mechanical integrity or electrical performance. These tests verified that both the mechanical properties and charge collection efficiency of the sensors were maintained even after bending. Additionally, the air cooling system underwent rigorous testing to ensure effective heat dissipation while maintaining detector stability and minimizing vibrations.

In summary, ITS3 represents a major leap forward in detector technology, building on the success of ITS2 while introducing significant innovations that will enhance the capability of the ALICE experiment to study the QGP. By reducing the material budget, improving spatial resolution, and enabling more precise measurements of heavy-flavor hadrons and dileptons, ITS3 will provide crucial insights into the fundamental properties of the QGP and the early universe. These improvements are not just incremental; they represent a transformative step forward in the capabilities of the ALICE experiment, ensuring that it remains at the forefront of heavy-ion physics research.

1.4.4 ALICE 3

Even with the advancements anticipated in LHC Runs 3 and 4, certain critical measurements — such as the production of multi-charmed baryons, the elliptic flow of electron-positron pairs, and photon production at very low momentum — are likely to remain beyond the reach of ALICE [59].

To address these challenges, ALICE 3 has been proposed as a next-generation experiment (see Fig. 1.25) at the LHC [60], specifically designed to tackle unresolved questions in heavy-ion physics, particularly in understanding the QGP and its properties. ALICE 3 is planned for installation at LHC Point 2 during the fourth long shutdown (LS4) in 2033-2034.

ALICE 3 will leverage advanced detector technologies to achieve its ambitious goals. The experiment will feature a large silicon tracker composed of CMOS MAPS arranged in 11 barrel layers, supplemented by 12 forward and 12 backward endcap disks, covering a total area of approximately 60 m^2 and spanning a wide pseudo-rapidity range ($|\eta| < 4$). The design includes a combination of cylindrical detection layers and traditional stave-mounted chips.

A Vertex Detector (VD), which includes the three innermost layers and disks, will be housed inside the beam-pipe in a secondary vacuum and will be retractable, similar to an iris diaphragm, allowing it to be positioned as close as 5 mm from the interaction point once the beams are optimized. This will provide a track pointing resolution of better than $10 \text{ }\mu\text{m}$ and will require a new sensor design, still based on the 65 nm CMOS node, featuring micro-channel cooling.




Figure 1.25 – The ALICE 3 experimental apparatus. Taken from [61].

The middle layers and outer tracker will employ more traditional methods, using staves on top of which pixel sensors are mounted, though options for using curved sensors in the innermost sensors of the middle layers are also being considered. The outer tracker will have an area of about 54 m^2 , accounting for the acceptance overlap of the staves. Due to the large scale of production, module fabrication is expected to be outsourced to industry.

In addition to the tracker, ALICE 3 will incorporate time-of-flight (TOF) detectors and a Ring Imaging Cherenkov (RICH) system for particle identification, enabling precise separation of electrons, photons, and hadrons across a broad momentum spectrum. The integration of these technologies will provide unparalleled resolution and efficiency in detecting and reconstructing particles, particularly at low transverse momenta, which are crucial for studying soft QCD processes and the early stages of heavy-ion collisions.

Building on the success of ALICE, the ALICE 3 experiment is designed to push the boundaries of precision in measuring the production of heavy-flavor hadrons and conducting detailed studies of electromagnetic radiation, among other investigations into the strongly-interacting matter in its deconfined phase. Among its primary physics objectives are the detailed studies of heavy-flavor hadron production in a large rapidity window and the investigation of states containing multiple charmed baryons (for example $\Xi_{cc}^{++} \rightarrow \Xi_c^+ \pi^+$ and $\Omega_{cc}^+ \rightarrow \Omega_c^0 \pi^+$) [62]. Strangeness tracking will be relevant in this and other physics channels. Ongoing studies are focused on optimizing the layer configuration (which will be more than those of ITS2 and ITS3), which is expected to significantly enhance the performance of strangeness tracking [57]. These studies are expected to improve our understanding of hadronization processes and help

resolve discrepancies between theoretical models such as the statistical hadronization and coalescence models.

1.5 Sensor characterization

1.5.1 Testbeams and particle telescopes

Sensor testing and characterization are critical steps in the development and deployment of any sensor technology, ensuring that each chip performs according to its design specifications under various conditions. This process identifies any defects or inconsistencies that could compromise the functionality of the sensor in real-world applications. By thoroughly characterizing the sensor, researchers gain a deep understanding of its behavior, which is essential for optimizing performance, improving reliability, and ensuring safety.

Prior to the sensor being physically built, a series of simulations helps refine key parameters related to its geometry and electronic response. For pixel sensors, these simulations typically include PSPICE electronic circuit optimization, which evaluates the functionality of both analog and digital circuits. Additionally, tools like Synopsis TCAD are used to simulate doping profiles and electric field characteristics, while frameworks like the Allpix² model charge generation and transport. These simulations guide the designers in making informed decisions, ensuring that the sensor will be efficient, have a fast charge collection time, exhibit good radiation hardness, and provide excellent timing and spatial resolutions.

The first stage of sensor testing usually begins in the laboratory, where controlled conditions allow for precise measurements and systematic experimentation. In this setting, sensors undergo a range of electrical and physical tests to verify their basic functionality. Key parameters such as operating voltages, power consumption, and noise levels are measured. Various calibrations are performed, often involving pulsing capacitances in the pixel sensors to relate DAC values to thresholds in electrons, or using radioactive sources for direct calibration.

While laboratory tests are crucial, they cannot fully replicate the complex and dynamic environments in which sensors will ultimately operate. For instance, most radioactive sources used in laboratory tests have decay particles with significantly lower energies than those encountered in high-energy physics experiments. Moreover, particles from these sources often cannot penetrate multiple sensor layers, and due to their low energy, they scatter significantly, making track reconstruction nearly impossible. This limitation underscores the importance of test beams, which provide a controlled yet realistic environment where

sensors can be exposed to high-energy particles similar to those they will detect in actual experiments.

At test beam facilities, sensors are subjected to beams of particles with well-defined properties and large momenta, allowing researchers to study the performance of chips in detecting particles, their spatial and timing resolution, and their ability to handle high particle fluxes. These facilities simulate the intense conditions that sensors will face in high-energy physics experiments, such as those conducted at the LHC.

Test beams are essential for several reasons. First, as already mentioned, they provide a realistic test environment that cannot be fully replicated in the laboratory, enabling a more accurate assessment of the sensor performance under real-world conditions. Second, test beams allow for the study of how sensors respond to specific types of particles and energy ranges, which is crucial for optimizing their design for particular experiments. Third, they help identify potential issues that may have eluded simulations or that may not have been apparent during laboratory testing, such as unexpected interactions with high-energy particles or the effects of sustained high-radiation exposure.

Test beam studies also provide invaluable data for calibrating the sensor and its readout electronics, ensuring that the data collected during actual experiments is accurate and reliable. The insights gained from test beams often lead to refinements in sensor design, improvements in performance, and enhancements in robustness, making test beams a critical component of the sensor development process.

As an analogy often quoted in the field: "If building a modern particle physics experiment is a marathon, then visiting a test beam facility is the 100-meter dash" [63]. In just one or two weeks, small teams work non-stop to gather as much data as possible on a piece of equipment they are considering for installation in an experiment.

Silicon sensors are typically tested at both test beam and irradiation facilities, allowing researchers to evaluate key performance metrics under realistic conditions. During these tests, detectors are mounted in carefully controlled environments where various parameters can be systematically varied. High-energy particles passing through the detector generate electronic signals, which are later analyzed. The data collected is critical for understanding the behavior of the detector under realistic conditions and for making necessary adjustments to its design or calibration.

Irradiation facilities are equally important because they allow for the evaluation of detectors, their electronics, data transmission components, and on-board data processing units under conditions that simulate the high radiation environments expected in modern exper-

iments. These facilities provide the high-dose-rate environments needed for accelerated testing, enabling researchers to simulate years' worth of radiation exposure in just days or weeks.

The key instrument used in all measurements with silicon sensors (and not only) is a “telescope”. Particle telescopes are sophisticated setups used in test beams to precisely measure the trajectory of particles as they pass through different layers of detectors. These telescopes typically consist of several layers of well-characterized silicon pixel detectors placed along the path of a collimated beam of particles. Each layer records the position of particles as they pass through, allowing for the reconstruction of their trajectories with high precision. An example of such a telescope is shown in Fig. 1.26.



Figure 1.26 – A particle physics telescope composed of six ALPIDE sensors is shown in this photo. The sensors are rotated along their vertical axis to allow particles to pass at an angle, in order to study the resulting effects. The sensors appear as the small terracotta rectangles positioned in the middle of the square green PCBs on the top side of the picture, which serve as carrier cards. These cards not only enable easy swapping and exchanging of sensors, but also act as an interface to which the sensors are wire-bonded. The carrier cards are inserted into data acquisition (DAQ) boards that house an FPGA chip for fast data processing. The DAQ boards are connected to a wide array of cables that transmit control signals, send data, and receive power. Each telescope plane is mounted on translation stages for easy alignment. Two scintillators, of which one is visible on the left, act as triggers for the setup.

The primary role of a particle telescope in testing silicon pixel detectors is to provide a precise reference track for the particles. A new sensor with untested performance, known as the device-under-test (DUT), is placed in the middle of the telescope, where the uncertainty due to tracking is minimized because measurement points are available both before and after the DUT.

The reference sensors in the telescope are used to reconstruct a track, with typically at least three sensors upstream and three downstream of the DUT. This arrangement ensures good tracking efficiency and provides a unique solution for fitting a track in three dimensions. These reference tracks are then used to characterize the sensor under test.

By comparing the measured position of hits on the DUT with the known positions of the reference tracks reconstructed by the telescope, one can determine crucial parameters such as detection efficiency and resolution. Position resolution refers to the ability of the detector to accurately determine the position of a particle as it passes through the sensor. In test beam experiments, this is measured by comparing the position recorded by the detector under test with the reference position provided by the particle telescope. The difference between these measurements gives an estimate of the position resolution.

Time resolution, on the other hand, measures the ability of the sensor to accurately record the time at which a particle passes through the sensor. This is particularly important in experiments where precise timing information is required, such as in time-of-flight measurements or in experiments where particle collisions occur at extremely high rates. Time resolution is typically measured by analyzing the time difference between the particle passing through the detector and the reference time provided by a fast timing detector or an external clock.

Detection efficiency is another crucial parameter evaluated during test beam experiments. It refers to the probability that a particle passing through the detector will be successfully detected and correctly recorded. High detection efficiency is essential for ensuring that the detector can reliably track particles, even in high-flux environments where many particles are passing through simultaneously.

To measure efficiency, the number of particles detected by the sensor under test is compared to the number of particles known to have passed through the detector, as measured by the particle telescope. This ratio provides the efficiency of the detector. If the efficiency is lower than expected, it may indicate issues such as dead pixels, suboptimal sensor settings, regions with low charge collection efficiency or other factors that need to be addressed before the detector can be deployed in an actual experiment.

Test beam campaigns typically last a full week, with five to six days dedicated to actual data collection. The remaining time is spent assembling the setup, making necessary electrical connections, aligning components, and sometimes debugging or adjusting parameters, samples, and detector positions. These intensive campaigns are vital for ensuring that sensors perform optimally in their final applications, providing accurate and reliable data in the challenging environments of high-energy physics experiments.

1.5.2 Track models

A number of track models are available to describe the trajectory of a particles through a particle telescope, differing in complexity based on factors such as the momentum of the particles, the number of tracking layers, the size of the sensor pixels, and the material budget encountered along the path.

In all the test beam analyses performed in this thesis, no magnetic field was present, meaning that particles are expected to travel through the setup in a near-straight line. A particle traversing the setup along the z (beam) direction can be described by a straight-line equation:

$$\vec{r}(z) = \vec{s} + \vec{d} \cdot z$$

where $\vec{r}(z)$ represents the position vector of any point on the line, \vec{s} is the state vector of the track at $z = 0$ (a reference point on the line) and \vec{d} is a direction vector indicating the direction of the line (a 2D slope).

The quality of the track fit is given by the chi-squared value (χ^2), which measures the goodness of the fit between a set of observed points r_{obs} and the straight line model:

$$\chi^2 = \sum_{i=0}^{n_{planes}} \frac{(\vec{r}_{obs,i} - \vec{r}(z_i))^2}{\sigma^2} \quad (1.8)$$

Here, i indexes of the measurement plane (sensor) in the telescope, n_{planes} is the number of planes used in the fit, $r_{obs,i} = (x_i, y_i, z_i)^T$ represents the center of a cluster of pixels formed by a track on the i^{th} plane in global coordinates and $\vec{r}(z_i)$ is the intercept of the track at position z_i (at the i^{th} measurement plane). The measurement uncertainties are taken equal in x and y and equal to the intrinsic position resolution of the ALPIDE sensor (5 μm).

Minimizing the χ^2 with respect to the track parameters (calculating the derivatives of Eq. (1.8) and setting them to zero) yields the best-fit values, a process used later in detector alignment with respect to a set of tracks. The reduced chi-squared value (χ^2/ndf) is often reported, accounting for the degrees of freedom, given by $2 \cdot n_{planes} - 4$ (with two measurements per plane and four constraints from the track fit).

For high-momentum particles in low-material environments, the primary uncertainty arises from pixel size, making a straight-line fit sufficient. However, for low-momentum particles or systems with substantial material budgets, multiple scattering becomes significant, necessitating more complex models that account for interactions with the surrounding air, pixel sensors, or test samples in the path of the particles.

The data analyzed in this thesis was collected at the DESY II test beam facility, as well as at the SPS of CERN. While high-momentum particles at the SPS generally allow for straight-line approximations, the lower momentum electrons (1 GeV to 6 GeV) available at DESY often require track models that incorporate multiple scattering to improve performance.

A model employed in this work to account for this is the General Broken Line (GBL) model implemented in the Corryvreckan framework. Unlike simpler models that assume a straight-line trajectory, GBL divides the path of the particles into segments (“broken lines”) between measurement points. Each segment can be adjusted to account for deflections caused by multiple scattering. The position and orientation of each line segment are described by parameters estimated using methods such as least squares fitting or maximum likelihood.

This approach provides a more accurate reconstruction of the trajectory of particles, particularly in environments with significant material budgets or with low momentum where multiple scattering cannot be neglected. GBL is especially useful in multi-layered detectors, where accumulated scattering effects can significantly distort the straight-line approximation, leading to inaccuracies in measuring the momentum and trajectory of particles.

In the GBL approach, the trajectory is fitted using a combination of local straight-line segments and small-angle deflections (kinks) at specific points, typically at each detector layer (see Fig. 1.27). The model employs a least-squares fitting procedure that minimizes the difference between the measured positions and the predicted trajectory while accounting for the uncertainties due to multiple scattering. This method ensures that both scattering and measurement errors are properly incorporated, leading to a more precise determination of track parameters.




Figure 1.27 – Example of two track models in a telescope where one plane is thick. **(top)** Straight line track does not accurately describes the trajectory of the particle. **(bottom)** GBL track better approximates the trajectory and has smaller residuals.

Defining a kink angle $\vec{k} = (k_x, k_y)$ at each scatterer, we can define the χ^2 function by the sum of the χ^2 for the measurement offsets χ_m^2 and kink angles χ_k^2 , which is later minimized during the track fit and sensor alignment [64]:

$$\chi^2 = \chi_m^2 + \chi_k^2 = \sum_{i=1}^{n_{\text{planes}}} (\vec{x}_i - \vec{x}_i^{\text{pred}})^\top \mathbf{V}_{m,i}^{-1} (\vec{x}_i - \vec{x}_i^{\text{pred}}) + \sum_{i=2}^{n_{\text{scat}}-1} \vec{k}_i^\top \mathbf{V}_{k,i}^{-1} \vec{k}_i \quad (1.9)$$

The predicted hit positions (\vec{x}_i^{pred}) are obtained from the track parameters and the \vec{x}_i are the measured hit positions at the i^{th} detector layer. Their difference represents the residuals.

The first term reflects how well the predicted track matches the actual measured positions at each measurement plane, weighted by the inverse of the measurement covariance matrix $\mathbf{V}_{m,i}$. The second term, which involves the scattering angles \vec{k}_i , accounts for the influence of multiple scattering along the track.

The GBL model is implemented efficiently, with the model matrix structured to allow inversion in linear time relative to the number of measurements. The GBL fit computes the full covariance matrix of the corrections to the track parameters along the trajectory, providing essential input for track-based alignment with global methods, such as Millepede II.

1.5.3 Alignment models

During the test beam, great care is taken to align the sensors as accurately as possible. This maximizes the measurement area and ensures precise data collection. Initial mechanical alignment typically achieves accuracy on the order of below a millimeter to a few hundred microns per plane. However, due to inevitable small misalignments and mechanical tolerances, the reference setup and the DUT must undergo alignment corrections offline to refine this initial setup.

In this process, it is essential to understand the different coordinate systems used for sensor alignment (see Fig. 1.28). The first is the pixel coordinate system, where the position of hits on the sensor is defined by the pixel numbers (row, column). This system is intrinsic to the sensor and does not directly correspond to physical distances. The second is the local coordinate system, which represents the sensor in physical units (millimeters) and is used to describe the position of hits in a more physically meaningful way. The pitch of the pixels is used to span the matrix in both directions. This coordinate system has its origin in the center of each pixel sensor. Finally, the global coordinate system, often referred to as the telescope coordinate system, places all sensors in a common frame of reference, allowing the trajectory of a particle to be tracked across multiple sensors. This system also uses millimeters as the

unit of measurement and is defined by a sensor chosen as reference, with respect to which distances to the other sensors are calculated.




Figure 1.28 – The three coordinate systems in a particle telescope.

Sensor alignment involves adjusting several degrees of freedom (dofs) to ensure that the measured positions are accurate across the entire telescope and reflect (but account for) the true misalignment present in the system. The most robust degrees of freedom include translations in the x and y directions, which correspond to shifts in the plane of the sensor. These are critical because any misalignment in these directions can lead to significant errors in the reconstructed particle trajectories. Additionally, rotations around the z -axis, which are in the plane of the sensor, are also robust and must be carefully corrected. Less robust degrees of freedom include rotations around x and y axes, which involve tilts of the sensor out of its plane. These tilts are harder to detect and correct because they do not significantly affect the measurement of particle positions. Nonetheless, they are less prone to be present. Alignment constants are expected to be completely insensitive to translations along the z axis for particles that travel parallel to this telescope direction. Usually, the z distances between planes are provided as direct input.

The alignment process is typically performed using iterative methods that refine the sensor positions and orientations based on the residuals between the predicted (track intercept) and measured (cluster center of gravity) hit positions. Initially, a coarse alignment is done using the correlation of the cluster positions on each plane. This is followed by a more precise alignment, often using track-based methods where particle trajectories are reconstructed assuming an initial alignment, and the residuals between the expected and observed hit positions are minimized. These residuals are then used to adjust the sensor positions and orientations iteratively. This process continues until the alignment converges, ensuring that the sensors are correctly positioned relative to each other and the global coordinate system, and that the reconstructed particle trajectories are as accurate as possible.

1.5.4 Parameters of interest

Residuals and position resolution

Residuals are key parameters in track reconstruction and sensor performance evaluation. A residual is defined as the difference between the track intercept on a sensor and the center-of-gravity of the associated cluster.

The width of the residual distribution, often quantified by the standard deviation or RMS, gives the spatial resolution of the sensor convoluted with the precision of the track reconstruction. The residuals exhibit distinct probability density functions depending on whether the measurement on the plane was used for tracking or not. Biased and unbiased residuals are commonly distinguished in this context.

Biased residuals are calculated by comparing the measured hit position on a sensor with the predicted position from a track fit that includes the hit itself. This introduces a bias because the hit being evaluated contributes to the definition of the track, typically resulting in smaller residuals.

Unbiased residuals, on the other hand, are calculated by excluding the hit in question from the track fit. The predicted position is obtained from the track extrapolation using other hits in the event. This method avoids the bias introduced by using the same hit in both the prediction and the comparison, providing a more accurate measure of the true resolution of the sensor.

Unbiased residuals are particularly important in sensor alignment and performance studies because they reflect the actual performance of the sensor independent of the fitting process. The following relation is true:

$$\sigma_{\text{unbiased}}^2 = \sigma_{\text{intrinsic}}^2 + \sigma_{\text{track}}^2$$

where σ_{unbiased} is the measured width of the residuals distribution, $\sigma_{\text{intrinsic}}$ is the intrinsic position resolution of the sensor and σ_{track} is the precision of the track reconstruction.

For small pitch sensors, like the ALPIDE it will closely resemble a Gaussian, while for large pitches it will look like a box plot, convoluted with a Gaussian smearing the edges. The uncertainty due to tracking can be calculated for appropriate track models, using online tools, like [65, 66].

For digital sensors, the width of the residuals is fundamentally limited by the binary readout process. In the absence of any other effects, this upper bound can be approximated by the pitch of the sensor divided by $\sqrt{12}$ [9], which accounts for the uniform distribution

of hits within the pixel. For example, in the ALPIDE sensor, this would correspond to about $8\text{ }\mu\text{m}$.

However, clustering techniques and the physical properties of the sensor, such as charge diffusion, can significantly improve the effective resolution. In the case of ALPIDE, the effective residual width is reduced to about $5\text{ }\mu\text{m}$, thanks to the charge diffusion process, which allows for more precise localization of the hit within the pixel grid. This improvement is critical in achieving high-precision tracking necessary for accurate particle trajectory reconstruction.

Detection efficiency

The detection efficiency, ϵ , is another crucial parameter that measures the ability of sensors to detect and register particle hits. It is defined as the ratio of the number of tracks (n) matched with a cluster hit within the geometrical acceptance of a DUT, to the total number of accepted telescope tracks (N):

$$\epsilon = \frac{n}{N}$$

The uncertainty in the efficiency, σ_ϵ , is determined considering a binomial distribution of the detection process. It is quantified as the boundaries of the frequentist Clopper-Pearson confidence interval of $\pm 1\sigma$, as implemented in the ROOT framework [67].

In analogy to the standard deviation of a Gaussian normal distribution, it corresponds to the central 68.3% of the binomial distribution, but takes the lower limit of 0 and the upper limit of 1 into account correctly yielding asymmetric error bars.

In summary, testing silicon pixel detectors at test beams is a critical step in ensuring their performance in real-world experiments. Particle telescopes play a vital role in these tests by providing precise reference measurements, allowing for the accurate determination of the position and time resolution, as well as its detection efficiency of sensors. The data gathered from these tests is invaluable for optimizing the design of the chips and ensuring their readiness for high-energy physics experiments.

Bent sensors

The LHC Long Shutdown 3, scheduled for 2026-2028, will bring significant upgrades to the Inner Tracking System (ITS) of ALICE. The current three innermost layers of the ITS2, featuring a total of 432 ALPIDE chips mounted on staves, will be replaced by six fully integrated large-area sensors, each approximately $10\text{ cm} \times 27\text{ cm}$ in size. These sensors will be thinned down to $50\text{ }\mu\text{m}$ and bent to radii as small as 2 cm .

In preparation for these next-generation bent ASICs, readily available ALPIDE sensors were bent into various configurations in order to thoroughly test the characteristics of the active CMOS matrix under bending stress and compare them to their flat counterparts. Several dedicated test beam campaigns were conducted to test these bent ALPIDE setups. The results of these tests are summarized in three documents published in [68–71] and in preparation.

The author, alongside other ITS3 collaborators, played a key role in the steps following the construction of the bent sensors. This included participation in five measurement campaigns at the DESY II and CERN SPS facilities in 2020 and 2021, covering the preparations for the test beam, the setup configuration and assembly, as well as data acquisition.

The author also contributed to various stages of data analysis and the preparation of publication documents. Contributions included adapting the Corryvreckan analysis framework to accommodate bent sensor geometry, software alignment of the sensors in specific setups, conducting follow-up laboratory studies, extracting key parameters such as position resolution and efficiency, and contributing to the writing of the publications.

ALPIDE sensors bent at the three target radii for ITS3 (18, 24, 30 mm) were shown to remain fully operational without any performance degradation in terms of efficiency and spatial resolution compared to their flat state. Moreover, they were operated standalone in a configuration that mimics the future ITS3 detector, demonstrating excellent track impact-parameter resolution for reconstructed primary tracks to the interaction vertex.

Given that the ALPIDE sensors were fabricated using a different CMOS technology node (180 nm) than the ones proposed for the ITS3 (65 nm), further studies are ongoing with the new chips, with early results confirming findings observed with the ALPIDE sensors.

Through these efforts, the ITS3 community achieved an important milestone in the development of fully cylindrical silicon tracking detectors based on monolithic active pixel sensors, showing that bent sensors are an excellent choice for the next generation of fully cylindrical ultra-low material budget tracking detectors.

2.1 Bending ALPIDE sensors

2.1.1 Bending CMOS sensors

In recent years, significant advancements have been made in the development of ultra-thin and flexible electronics, particularly with the integration of silicon IC technology into highly flexible platforms [72]. These innovations have not only enhanced the mechanical resilience of electronic components under bending stresses, but have also opened new avenues for applications in wearable technology, smart packaging, and other flexible electronics domains. Researchers have focused on understanding and improving the mechanical and electrical properties of these ultra-thin chips, ensuring that they maintain performance even under the stress of deformation. This foundational work has set the stage for further exploration into the bending capabilities of sensors and electronic assemblies, pushing the boundaries of what is possible in the realm of flexible electronics.

Following these studies, readily-available ALPIDE sensors, similar in most characteristics to the ones intended for use in the ITS3, were systematically tested with three and four-point test methods, where a force is applied and the breaking point is measured [55, 73].

Sensors thinned down to 50 μm , in both 180 nm and 65 nm technology nodes, were shown to achieve bending radii below 10 mm, but tended to break due to large stress induced forces shortly thereafter. The bending force scales inversely with the cube of the thickness. Sensors that have been thinned down to 30 μm demonstrated enough flexibility to slide into the aperture of the setup without breaking, achieving bending radii as low as 5 mm.

2.1.2 Bending along the short axis

The very first operational bent ALPIDE sensor tested by the ITS3 collaboration is depicted in Fig. 2.1. Prepared by collaborators from INFN Trieste, this sensor was bent along its short axis (along the rows), a variant referred to as $\nu 0$ in the following discussion.

A 50 μm thick ALPIDE sensor was bonded in the usual manner to a carrier card along its long edge, where the bonding pads are located. A strip of 3MTM Adhesive Transfer Tape 467MP was applied to the periphery of the sensor, securing it in place. The adhesive is placed on a 2 mm wide area at the bottom of the sensor and spans the entire length of the sensor. Considering that the periphery of the ALPIDE sensor is about 1.2 mm wide over the entire length of the chip, this results in a narrow strip of the active area, approximately 0.8 mm wide which is also glued to the carrier card.




Figure 2.1 – The first operational bent ALPIDE sensor tested by ITS3, bent along its short axis. The rows go along the x direction and the columns go along the bent y direction.

Before bending, the sensor was wire-bonded to the carrier card, with the two top corners left unattached. This ensured that the bonding area remained flat and well-secured throughout the bending process.

To bend the sensor, it was sandwiched between two $120\text{ }\mu\text{m}$ thick Kapton[®] (polyimide) films. The foils are glued with the same tape around the sensor, in order to exert some pressure and keep the chip fixed. A mechanical enclosure, consisting of an Ω -shaped aluminum fixture with two metallic wheels, was used to curve the sensor. The polyimide films were attached to these wheels, which were then moved in parallel to the short edge of the sensor using a micrometer screw, wrapping the foils around the wheels and bending the sensor. Once the desired curvature was achieved, the wheels were locked in place.

During the bending process, the top side of the sensor (the one visible in the picture), which contains the electric circuitry, was compressed. Dedicated studies [74] have shown that bending affects the analogue power supply current, with relative changes of up to 10%, less than the strain induced by bending of approximately 0.1%. This outcome is attributed to the piezo-resistive effect, which represents the changes in the electrical resistivity of the semiconductor under mechanical strain. The analogue current change, which dominates the digital one, is primarily due to shifts in the pixel transistors and, to a lesser extent, to changes in the on-chip DAC outputs.

The curvature of the sensor was measured using a Mitutoyo Coordinate Measuring Machine (CMM), both before and after the test beam, with the sensor positioned on the measuring table as shown in Fig. 2.1. The 3D CMM measurement involves probing discrete points on the surface of the sensor using the tip of a probe head. When the probe contacts the sensor, the machine samples the axis position sensors, recording the location of that point on the surface. This process is repeated across the whole measuring surface, resulting in a 3D scatter plot.

The sensor was sampled in a grid-like pattern, with the probe tip following the curvature of the sensor multiple times along the rows and columns. For each point in the y direction, the points along the x direction were averaged, and their dispersion was taken as an uncertainty. The resolution of the CMM machine is $5\text{ }\mu\text{m}$ in the table plane and $80\text{ }\mu\text{m}$ along the z direction (height).

The measurement points are shown in Fig. 2.2. The data points were fitted with the following parametric function:

$$z(y) = \begin{cases} 0 & \text{for } y \leq y_0 \\ r - \sqrt{r^2 - (y - y_0)^2} & \text{for } y > y_0 \end{cases}$$

where z represents the height, y is the horizontal axis (along the bending direction), y_0 is the glued part of the active matrix that is assumed to be flat and r is the bending radius.

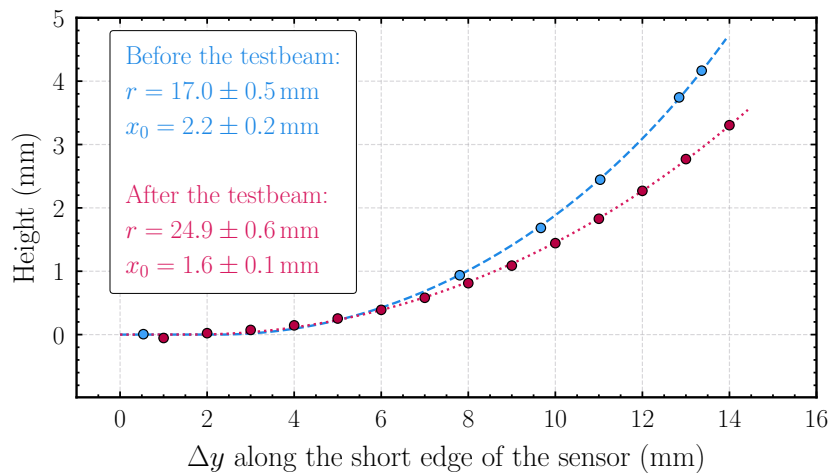


Figure 2.2 – Measurement of the curvature of the sensor using a CMM. Adapted from [68]. The blue points and fit represent data taken before the testbeam and red points and line represent data taken after the testbeam. The black line represents the fit to testbeam data at a radius of (22 ± 1) mm.

A notable feature of the plot is the observed change in the bending radius between the two measurements. The time difference between these measurements was approximately four months, primarily due to Covid restrictions in place at the time.

This change in radius is attributed to the relaxation of the Kapton holder over time. This suggests that the data collected during the test beam was at a larger radius than initially anticipated. This was later confirmed by the in-beam measurements.

2.1.3 Bending along the long axis

To correct for the relaxation of the polyimide foil and explore the effects of bending along different directions, subsequent bending experiments focused on ALPIDE sensors bent along the long axis (along the columns), where the top electronic circuitry was decompressed.




Figure 2.3 – ALPIDE sensors on 3D-printed jigs bent along their long axis. **(left)** The sensor is glued to the jig with an adhesive tape. **(right)** The sensor is held in place by a Kapton sleeve with a cutout.

Two types of sensors were produced for this purpose, as shown in Fig. 2.3. The first version (further denoted v1), depicted on the left, features an ALPIDE sensor that was glued using adhesive tape onto a 3D-printed cylindrical jig. The second version (further denoted v2) replaces the adhesive tape with a Kapton sleeve that has a cutout. Both versions include an aperture in the PCB, allowing particles to pass through this window that is smaller than the ALPIDE sensor itself, minimizing scattering, since the jig is 3.1 mm thick and made of a polypropylene-like material that could cause changes in the particle trajectory.

These sensors were produced by colleagues at CERN and INFN, and the production process is briefly described below. A more detailed description of the materials and procedures can be found in [75].

For the v1 version, shown in Fig. 2.4 (left), a 3D-printed jig with a window of 40 mm × 5 mm (in earlier designs) or 17 mm × 9 mm (in newer designs) was created. This window covers the full active length of the sensor in the bending direction, but only 5 mm out of the 13.8 mm of the matrix in the non-bent direction.




Figure 2.4 – Bending procedure for the v1 version. A strip of adhesive is placed between the jig and the sensor. The apertures are visible in the last picture and marked with red arrows. Adapted from [76].

A $4\text{ cm} \times 3\text{ cm}$ strip of 3M adhesive tape was applied to the jig around the opening, with a cutout made to remove any tape from the opening area. The ALPIDE chip was then placed on top of this adhesive tape, centered over the hole. A Kapton foil, slightly larger than the ALPIDE, was placed on top of the sensor and used to gently press the sensor into shape until it adhered to the tape, after which it was removed. Finally, a custom-made FPC (flexible printed cable) was screwed into place close to the periphery of the sensor in dedicated holes. This process is delicate, and the sensor is prone to damage. Additional tools were developed to improve this procedure.

For the v2 version, shown in Fig. 2.5, a diced ALPIDE sensor was picked from the dicing tray and was partially placed on a flat surface, leaving part of it hanging over the edge. The portion of the sensor resting on the surface was held by a vacuum chuck. The assembly was then aligned with a jig printed at a specific radius, attached to a motorized rotating stage with micrometer accuracy. An FPC cable was attached beforehand to the jig.



Figure 2.5 – Bending procedure for the v2 version. The sensor is sandwiched between the jig and a Kapton sleeve under tension and rolled in place. Adapted from [77].

Several variants of the jig, designed for different radii, were produced. A pre-cut Kapton sleeve with a large window ($25\text{ mm} \times 10\text{ mm}$) was first screwed in on one side with two hex screws, while the other side was held under tension by a weight over a pulley. The ALPIDE sensor was aligned in the desired position, and the stage began rotating, bending the sensor over the opening and the jig. Finally, the Kapton sleeve was secured on the other side, keeping the sensor in place. The window in the 3D-printed jig was enlarged and now measures $17\text{ mm} \times 9\text{ mm}$.

Once either version of the sensor was secured on the jig, the bonding process was performed (see Fig. 2.6). Because the usual carrier card could not be used with a sensor bent along its long edge, an FPC was used to interface between the sensor and the DAQ system. The custom-made FPC was wrapped around the cylindrical jig, placed very close to the periphery of the sensor, allowing a maximum distance of a few millimeters between the two.




Figure 2.6 – Bonding v1 and v2 ALPIDE sensors on cylindrical jigs. Taken from [78].

The micrometer stage rotated the jig to align the bonding pads to the needle head. One bond was performed, after which the sensor was rotated precisely until the next bond was perpendicular to the needle, and the process was repeated.

The custom-made FPC was first attached to an I-board in close proximity to the sensor itself. This board connects to a long (up to 30 cm) standard 15-pin FFC (flat flex cable), which interfaces with an adapter board. The extended length allows the bent sensors on the jig to be installed at different angles relative to the readout boards, as well as positioning the DAQ boards further away in the compact telescope designs. The entire setup for one bent sensor, connected through the entire chain, is shown in Fig. 2.7.




Figure 2.7 – The stacked ALPIDE sensors on cylindrical jigs of different radii are bonded to FPC that connect to I-boards. One I-board is connected to an adapter board via a long white FFC. Finally, the adapter board interfaced the usual DAQ board.

Discussion

The v0 variant of bent ALPIDEs showed a large relaxation along the bending direction, with the polyimide foil relaxing over approximately four months, resulting in a change in the bending radius of nearly 8 mm. Additionally, some sagging was observed in the middle of

The role of the I-board is to decouple the signal from the long FFC as close as possible to the sensor itself. It contains a number of resistors and capacitors. This is needed to have the same functionality as the carrier card that has an array of capacitors as close as possible to the bonds.

The adapter board has several functions. First, it provides decoupling for both the analog and digital domains using an array of capacitors to ground. It also enables the conversion of signals from a 3.3 V standard logic level to M-LVDS levels for the differential control port and reference clock input. Additionally, it safely translates signals from a 1.8 V logic level system to 3.3 V, facilitating communication between components operating at different voltage levels. The adapter board connects to the sensor side via an FFC and interfaces with the DAQ board through a PCI-e connector.

Using adapter boards minimizes the need for changes to existing software and provides a cost-effective solution for directly interfacing with existing DAQ boards.

the sensor along the x -direction, particularly on the top side of the Kapton foil (as shown in Fig. 2.1). This sagging caused noticeable alterations in the residuals of the data, as detailed in Fig. 8 of [68]. Specifically, a change of nearly $50\text{ }\mu\text{m}$ was observed between the corner where the Kapton foil is fixed by the wheels and the middle of the sensor along the x -axis, indicating that the cylindrical model used does not perfectly describe the surface shape.

To address these issues, the v1 version was developed, utilizing a 3D-printed cylindrical jig to ensure a well-defined bending radius. The precision of the mounting procedure and assembly was verified using a 3D optical profilometer, as shown in Fig. 2.8. The assumption of a cylindrical shape was confirmed to be accurate within $50\text{ }\mu\text{m}$ in the region where the opening is present. However, it was noted that the edges of the sensor tend to detach (as seen in Fig. 2.3 (left)), resulting in a larger radius around the edges that deviates from the nominal value of the jig.




Figure 2.8 – Cylindricity measurement of v1 bent ALPIDE sensors at three different radii. The jig aperture of $17\text{ mm} \times 9\text{ mm}$ is also marked. Measurements performed by L. Lautner.

To further improve the setup, the v2 and final version with these jigs was produced, featuring a thin Kapton sleeve with an opening that is screwed in on two sides, securing the bent sensor tightly to the 3D-printed jig. This design offers better control over the bending radius, as shown in Fig. 2.9. However, some warping is still visible, indicating that the bent

sensor does not form a perfect cylinder. This suggests that a 3D measurement would be beneficial for any alignment or calculations requiring a precise definition of the surface of the sensor.

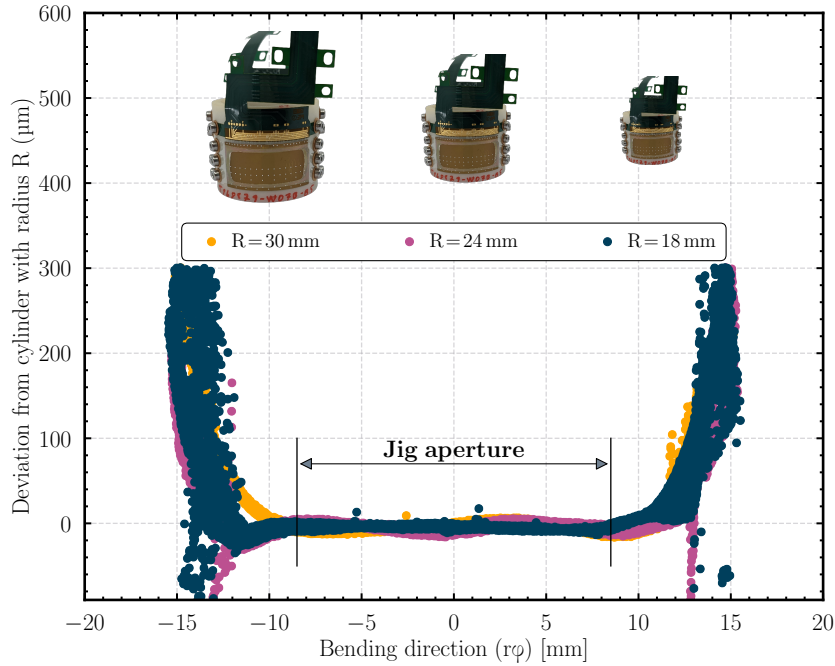


Figure 2.9 – Cylindricity measurement of v2 bent ALPIDE sensors at three different radii. Measurements performed by L. Lautner.

2.2 Alignment of bent sensors in software

In most track reconstruction and analysis software, pixel sensors are typically modeled as flat, planar surfaces. This assumption simplifies the geometry and the associated mathematics, making it easier to compute the interactions of particles with the detector surface. However, with the advent of more complex detector designs, such as those incorporating bent or curved sensors, there is a need to extend the existing models to account for these new geometries.

The curvature of these new sensors is a dimension that introduces additional complexity in determining the position where a particle intercepts the sensor surface, as well as in transforming coordinates between the detector frame and the global reference frame. Moreover, more degrees of freedom for sensor alignment are available, which are not normally available for flat sensor.

One of the primary challenges in implementing bent sensors is handling coordinate transformations. In the case of flat sensors, the transformation between the local sensor coordinates and the global coordinates is relatively straightforward, often involving simple translation and rotation operations. For bent sensors, however, the curvature of the surface means that these transformations are no longer linear.

The implementation involves defining the geometry of the bent sensor, typically through a mathematical model that describes the curvature. This could be in the form of a parametric equation or a surface defined by a set of control points. Once the geometry is defined, coordinate transformations must take into account the curvature. This involves transforming a point in the global reference frame to the local curved surface by projecting it onto the curved geometry.

This projection is not trivial, as it requires solving for the point on the curve that is closest to the given global point, which often involves solving a system of nonlinear equations. Additionally, the transformation from local to global coordinates must consider the differential geometry of the surface, including tangent vectors and normal vectors that vary across the surface.

The author was majorly involved in adapting the Corryvreckan [79] tracking reconstruction and analysis code used within the collaboration to accommodate bent sensors.

In the implementation, the bent sensor is modeled as a perfect cylinder. This choice of geometry simplifies certain aspects of the mathematics, as a cylinder has well-defined parametric equations and symmetry that can be exploited in calculations. Ideally, one would use a 3D profilometric scan mapping over the sensor surface as input to the calculations. This takes into account all local warps and distortions.

However, this is a very complex task. Instead, it was already shown, for example in Fig. 2.9, that deviations in the analysis area are in the order of $\pm 30 \mu\text{m}$ over about 20 mm. Therefore local deviations are expected to be small and the approximation of a perfect cylinder is good without introducing a big bias.

A new class called BENT PIXEL DETECTOR has been introduced at the core of Corryvreckan. This class extends the functionality of the base PIXEL DETECTOR class, adding the specific logic required for handling the cylindrical geometry.

The cylindrical geometry of bent detectors is set up by defining the radius of a cylinder and its axis of curvature. Currently, the code supports bending in only one direction, with the radius being fixed. However, future updates aim to make the radius a variable parameter to enhance alignment flexibility.

The data from the detectors consists of row and column information for each event and sensor plane. In the `BENT PIXEL DETECTOR` class, methods are provided to transform these coordinates into local sensor coordinates. The local reference frame representing the position in Cartesian coordinates is derived by converting the pixel indices into cylindrical coordinates (ρ, z, ϕ) , representing the radial distance from the origin, the axial coordinate (along the axis of the cylinder), and the angular coordinate (around the cylinder).

In the context of flat sensors, spatial resolution manifests as uncertainties in the x and y directions, corresponding to the plane of the sensor. However, for bent sensors, these uncertainties change due to the curvature. The spatial resolution matrix, an error matrix used to quantify cluster position uncertainties, is crucial in the track fitting and alignment processes.

For flat sensors, when a particle track hits the sensor perpendicularly (at an angle of 0° to the normal of the sensor), the uncertainty is confined to the x and y directions. This happens because the z direction, orthogonal to the plane of the sensor, does not contribute to the uncertainty of where the particle hits the surface (which by default is defined at a specific z specified by the user). From the perspective of the incoming particles, the flat sensors are perceived undistorted. The pixels appear with their true, unaltered dimensions relative to the incoming particles, meaning that the spatial resolution directly reflects the inherent geometry of the pixels. In this scenario, the uncertainties in the x and y directions are solely dependent on the pixel size and the resolution of the detector, without any distortion or contraction effects caused by curvature. The spatial resolution error matrix in this case is diagonal, with non-zero entries only in the x and y directions:

$$\begin{pmatrix} \sigma_x^2 & 0 & 0 \\ 0 & \sigma_y^2 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

For bent sensors, the situation changes. The curvature alters how particles perceive the pixel matrix. When a particle track intersects the sensor at an angle, particularly in the bending direction, the pixels appear “contracted” in that direction (see Fig. 2.10). This contraction reduces the uncertainty along the x -axis, as the particle traverses a pixel which is seen as shorter along this axis. Consequently, the uncertainty typically associated with the x direction is now transferred to the z direction, which is now accessible due to the curvature.

The clusters on a 18 mm bent sensor in global coordinates are illustrated in Fig. 2.11. This figure presents data from a testbeam where only a single v1 bent sensor was tested. The




Figure 2.10 – The pixel matrix as seen by incoming tracks appears contracted in the direction of bending. The projected pixels have a smaller pitch in this direction than their real one.

beam direction is oriented towards the positive z values, and the sensor is rotated by 90° around the global z axis. In addition to the clusters on the bent sensor, clusters from the two innermost flat reference ALPIDE sensors are shown in blue and orange. These reference clusters are shifted along z and brought closer in the software to the bent sensor to highlight its curvature. Additionally, 30 reconstructed tracks are depicted, each associated with clusters on all three sensors.




Figure 2.11 – Clusters associated to tracks on the bent sensor (red-blue) and on the innermost two reference flat ALPIDEs (orange and blue). With grey lines, 30 tracks are matched between the corresponding clusters of these three sensors.

When a particle traverses a flat sensor, determining the intercept of the trajectory of the particle with the sensor plane is straightforward. However, for bent sensors, the intercept calculation requires finding the intersection of a straight line in space with a curved surface.

In Corryvreckan, this intersection is computed by parameterizing both the trajectory of the particle and the surface of the sensor. Depending on the curvature of the sensor, multiple possible intercept points can arise, requiring additional logic to determine the correct intercept. This is guided by physical considerations such as the direction of the particle travel and the specific geometry setup.

The calculation of the intercept between a particle track and the cylindrical surface involves solving a quadratic equation derived from the geometry of the problem (see Fig. 2.12).




Figure 2.12 – Example of a particle intersecting a cylinder of radius R . The state and direction vectors of both the track and the cylinder are shown.

The position of the particle track in the global coordinate system can be represented by a position (state) vector $\vec{s}_0 = (x_0, y_0, z_0)$ and a direction vector $\vec{d}_{\text{track}} = (d_x, d_y, d_z)$. The parametric equation of the track can be written as:

$$\vec{r}(z) = \vec{s}_0 + z\vec{d}_{\text{track}}$$

where z is chosen as the direction the particle travels along.

The cylinder is represented by a central axis, characterized by a position vector $\vec{s}_{\text{cylinder}}$ and a direction vector $\vec{d}_{\text{cylinder}}$. The center of the cylinder can be shifted from $(0, 0)$.

The equation for a point on the cylinder surface can be written as:

$$\left\| \vec{r} - \vec{s}_{\text{cylinder}} - \left[(\vec{r} - \vec{s}_{\text{cylinder}}) \cdot \vec{d}_{\text{cylinder}} \right] \vec{d}_{\text{cylinder}} \right\| = R$$

where R is the radius of the cylinder, and r a point on the cylinder.

To find the intersection, the parametric equation of the track is substituted into the equation of the cylinder. This will result in:

$$\left\| \vec{s}_{\text{track}} + z\vec{d}_{\text{track}} - \vec{s}_{\text{cylinder}} - \left[\left(\vec{s}_{\text{track}} + z\vec{d}_{\text{track}} - \vec{s}_{\text{cylinder}} \right) \cdot \vec{d}_{\text{cylinder}} \right] \vec{d}_{\text{cylinder}} \right\|^2 = R^2$$

When expanded it represents a quadratic equation in z :

$$\alpha z^2 + \beta z + \gamma = 0$$

where:

$$\begin{aligned} \alpha &= \vec{d}_{\text{track}} \cdot \vec{d}_{\text{track}} - \left(\vec{d}_{\text{track}} \cdot \vec{d}_{\text{cylinder}} \right)^2 \\ \beta &= 2 \left[\left(\vec{s}_{\text{track}} - \vec{s}_{\text{cylinder}} \right) \cdot \vec{d}_{\text{track}} - \left(\vec{d}_{\text{track}} \cdot \vec{d}_{\text{cylinder}} \right) \left(\left(\vec{s}_{\text{track}} - \vec{s}_{\text{cylinder}} \right) \cdot \vec{d}_{\text{cylinder}} \right) \right] \\ \gamma &= \left(\vec{s}_{\text{track}} - \vec{s}_{\text{cylinder}} \right) \cdot \left(\vec{s}_{\text{track}} - \vec{s}_{\text{cylinder}} \right) - \left(\left(\vec{s}_{\text{track}} - \vec{s}_{\text{cylinder}} \right) \cdot \vec{d}_{\text{cylinder}} \right)^2 - R^2 \end{aligned}$$

The solutions for z are given by the quadratic formula:

$$z_{1,2} = \frac{-\beta \pm \sqrt{\beta^2 - 4\alpha\gamma}}{2\alpha}$$




Figure 2.13 – Possible cases for track intercepts with a cylinder

The solutions correspond to the two points where the track intersects the cylinder (see Fig. 2.13). Depending on the direction of the track, one of these points will be the first intercept as the track enters the cylinder. The correct intercept is determined based on the geometry of the setup, typically by selecting the smaller positive z value. For example, in scenarios where two bent sensors are placed on opposite ends of the 3D-printed jigs, those two intercepts would represent hits on the two different sensors.

The code also accounts for cases where the track intercepts just one of the sensors or none at all and treats those cases accordingly, by performing a few other checks.

The intercept point is calculated by substituting the selected value of z back into the parametric equations of the track:

$$\vec{r}_{\text{intercept}} = \vec{s}_0 + z\vec{d}_{\text{track}}$$

This equation provides the position vector of the intercept point on the cylinder surface.

For alignment, the current implementation supports rotations around the center of the cylinder, as well as shifts and rotations in all three spatial directions. However, modeling of sensor detachment at the corners is not included.

Other potential degrees of freedom for misalignment exist, which are not accounted for in the current implementation. Systematic shifts of the sensors and changes in the track parameters can counterbalance each other without affecting the overall χ^2 of the track fit. As such they can introduce a bias in the track parameters. These effects, known as weak modes, may include shearing and torsion, or shifts along the z (beam) direction of the different planes. Using testbeam data these are hard to model. Nonetheless, the implemented degrees of freedom are sufficient for achieving the required alignment presented in this study.

The implementation of bent sensors in Corryvreckan enhances the capability for more accurate detector modeling of ITS3-like geometries. By extending the existing framework to accommodate curved surfaces, this development provides the necessary tools to model a broader range of detectors with curved designs. Further improvements are anticipated and will be pursued in future work.

2.3 Data analysis of bent sensors at different radii

2.3.1 Sensors bent along the short edge

The v0 sensor underwent laboratory characterization to assess its electrical functionality before and after bending. The analysis revealed that the number of non-responsive pixels and the fake-hit rate showed negligible changes compared to the flat state of the sensor. Specifically, six hot pixels were detected both before and after the bending process, and a total of 1024 dead pixels (equivalent to a double column, the functional readout unit of an ALPIDE sensor from the matrix to the periphery) were present prior to bending, with the same number remaining afterward. The regions surrounding these dead pixels were carefully considered during efficiency measurements. The pixel threshold distribution, illustrated in Fig. 2.14 shows this.

This distribution was obtained by injecting a test charge on the VPULSE line, located immediately after the collection diode and reset stages. This charge traverses the entire front-end and digital pixel circuitry, propagating through the matrix via the priority encoder combinatorial digital circuit to reach the periphery, thus being sensitive to both the analog and digital processing chains.




Figure 2.14 – Threshold scan over the surface of the sensor before and after bending.

The threshold distribution demonstrates negligible differences from the flat state of the chip, indicating that the integrated electronics on the top side of the ALPIDE sensor were unaffected by bending stress. This procedure was also repeated on several other bent devices, yielding consistent results.

After verifying the performance in laboratory tests, the sensor was then tested in a testbeam campaign at the DESY II testbeam facility.

Testbeam setup

The testbeam was carried out in June 2020 at the DESY II accelerator, beam line 24, using a beam of 5.4 GeV electrons. The telescope setup used for these studies is depicted in Fig. 2.15.

A sketch of the setup showing the distances involved can be found in Fig. 2.16.

The reference telescope consisted of six ALPIDE sensors, with three positioned upstream and three downstream, forming two reference arms. These sensors were operated under nominal settings ($\geq 99.9\%$ detection efficiency, $5\text{ }\mu\text{m}$ intrinsic spatial resolution) to efficiently and precisely track the electrons from the beam. The sensors were housed in metallic enclosures for safe handling, which allowed for secure attachment to fixed positions on an optical table. The DAQ boards, used for readout and interfacing with the PC, were connected via PCI-e on the top side.

The bent sensor, shown in Fig. 2.1, was positioned at the center of the setup, between the two reference arms. It was connected to the DAQ board in the same manner as the other




Figure 2.15 – The telescope used for measuring the ALPIDE bent along the short axis. The three upstream and three downstream reference ALPIDE sensors are housed in metallic enclosures, which are then screwed at fixed positions on an optical table. Also visible are the DAQ boards and the bent sensor from Fig. 2.1, placed as a DUT in the middle of the telescope. The whole setup was later light-shielded by placing a metallic box on top. A second telescope box is visible in the background and was operated in parallel.




Figure 2.16 – Sketch of the telescope setup used to test the v_0 bent ALPIDE sensor. Six reference ALPIDE sensors, three upstream and three downstream, reconstruct tracks that are then interpolated at the position of the bent sensor.

ALPIDE sensors. The device under test (DUT) was securely held in place by two L-shaped angle brackets, which also allowed for vertical adjustments. Finally, a cover box was placed over the entire setup to provide light shielding and protection for the sensitive components.

The bent sensor was tested both without (0 V) and with (−3 V) reverse bias. However, the reverse bias data is limited and will not be shown in the following analysis.

The entire setup was operated using an external trigger system. The trigger was generated by the coincidence of the discriminated outputs from two PMTs connected to scintillators, each measuring 4 cm × 5 cm (larger than the ALPIDE sensors), positioned in front of and behind the telescope. The trigger logic was implemented using NIM crates, such that when both scintillators detected a particle within a 100 ns time window, a trigger signal is forwarded.

Each discriminated scintillator output initiated a 50 μs post-protection time window, which prevented event pile-up by excluding triggers from beam particles arriving in close succession. This was necessary because the in-pixel amplifier pulse of the ALPIDE sensors could extend to several tens of microseconds, depending on the operating settings [80]. In addition to the post-protection, an event separation of 100 μs was enforced after each accepted trigger.

These decisions were combined with a BUSY signal from the telescope in a logic OR configuration. The BUSY signal was raised when at least one DAQ board was processing an event and was released when the system was ready to take new data. A detailed diagram of this trigger logic is available in Fig. 4.6 of [81].

The data collected during the testbeam was taken with EUDAQ2 [82] and processed using the Corryvreckan reconstruction software [79], and the results are included in [68].

Initially, pixels that trigger more than 20 times the average rate are identified as noisy and are masked from further analysis. Cluster positions from hits on the six reference ALPIDE sensors are used to fit straight lines, forming tracks. Due to the larger area of the scintillators compared to the ALPIDE sensors, an average track multiplicity of approximately 0.7 was observed.

Only one track per event is allowed in the analysis. However, due to the strict trigger logic, the low noise level in the ALPIDE sensors, and the relatively low event rate (approximately 500 Hz) at beamline 24 for the momentum of 5.4 GeV/c, the probability of assigning a cluster to an incorrect track is minimal.

These tracks are then used to align the reference telescope. The ALPIDE planes are translated along the horizontal and vertical axes and rotated around the beam direction (the three degrees of freedom most susceptible to misalignment) relative to a reference plane (picked as plane 0) until the track χ^2/ndf is minimized. Following this alignment process, only tracks with a $\chi^2/\text{ndf} < 3$ are retained for further analysis.

The aligned tracks are then interpolated to the position of the bent ALPIDE sensor. For each event, a circular window with a generous radius of $250\text{ }\mu\text{m}$ is defined around the point where the track intercepts the sensor. If a cluster is found within this window, it is considered to be matched to the track.

Given that the sensor is bent along the row direction, an increase in the size of clusters is expected here due to the longer path that particles traverse through the sensor. This expectation is confirmed by the data shown in Fig. 2.17, where the average cluster size is plotted across groups of 30 rows. The corresponding incident angle, assuming a perfect cylindrical model, is also shown.




Figure 2.17 – Average cluster size in bins of 30 rows as a function of the chip row for the v0 bent ALPIDE sensor measured in beam.

Near the periphery of the sensor (row=512), where it is attached to the PCB of the carrier card, the tracks pass through the sensor perpendicular to its surface. At row 0, where the sensor has the highest curvature, particles pass through the sensor at an angle of almost 36° .

The data shows that cluster size decreases from row 0 to row 512, consistent with the decreasing angle at which particles pass through the sensor. The difference in average cluster size between row 0, where the angle is 36° , and row 512, where the tracks are perpendicular to the sensor, is approximately 25%. This is compatible with the expected increase of 23.6%

in the path length of particles at an angle of 36° , which traverse an average distance of $25 \mu\text{m}/\cos(36^\circ) = 31 \mu\text{m}$, aligning well with the observed increase in cluster size.

The DUT was aligned independently of the reference telescope using a residual minimization method. In the first iteration, the sensor was allowed to move along three degrees of freedom: translations along the x and y directions, and rotations around the beam axis. In subsequent iterations, the other two rotational axes and the translation along the z direction were also optimized. Finally, the bending radius of the sensor itself was adjusted to optimize the setup, resulting in a data-driven estimate of the radius, which was determined to be $(22 \pm 1) \text{ mm}$ [68]. This procedure was not done by the author, but by a colleague from the collaboration. Still, the author contributed to the results presented here and in the corresponding publication.

The detection efficiency across the full surface of the sensor is calculated by taking the ratio of reference tracks that have an associated cluster on the bent sensor to the total number of reference tracks reconstructed by the six ALPIDEs in the telescope. This is then expressed as inefficiency $(1 - \epsilon)$ for better visualization and plotted as a function of the sensor threshold in Fig. 2.18. For better clarity, three efficiency lines are also included in the plot, below which values larger than those indicated are achieved.




Figure 2.18 – Sensor inefficiency on a lin-log scale binned over 64 rows as a function of the chip threshold and the row for the v0 bent ALPIDE sensor. Older data from flat ALPIDE sensors, taken with a different chip than the one being bent here, are shown for reference.

The data points represent an average over 64 rows and are color-coded. At least 8k tracks The matching between row and incident angle is also indicated on the plot, assuming a cylindrical surface with a radius of $(22 \pm 1) \mu\text{m}$ and that tracks passing row 511 are perpendicular to the sensor surface.

For comparison, older data taken with a flat ALPIDE sensor, other than the one being bent here, is overlayed and matches well with the values from the bent sensor where the incident angle is close to zero.

The nominal operating point of an ALPIDE sensor is around a threshold of 100 e^- , where it can be seen that the sensor maintains an efficiency of over 99.9%. However, at higher thresholds, the bending effect becomes more apparent, with the most bent regions of the sensor (closer to row 0) showing higher inefficiency.

Below a threshold of about 40-50 electrons, noise levels increase, leading to a higher likelihood of mistakenly associating noise clusters to tracks. Therefore data below 40 electrons is not considered.

At a threshold of 300 e^- , there is a significant difference in inefficiency between regions where tracks are perpendicular to the sensor surface and regions where tracks pass through at an angle of 36° . This can be attributed to the fact that in the most curved parts of the sensor, tracks traverse about 25% more epitaxial layer (active volume), leading to a corresponding increase in charge deposition.

These results demonstrate that the ALPIDE sensor, when bent along the short axis, retains performance characteristics (such as noise level, dead pixels, and efficiency) comparable to those of a flat ALPIDE sensor. Bending was found to have minimal impact on these parameters, although some changes in analog currents were observed, likely due to piezo-resistive effects. These changes remain within the operational tolerances for ALPIDE chips.

However, the bending method used here is not ideal. The sensor mounting allowed relaxation of the polyimide foil used to sandwich the sensor, resulting in an unstable bending radius. Additionally, this method introduces extra scattering material into the path of beam particles in the form of the two Kapton tapes, which can interfere with particle trajectories, particularly at the electron energies used at DESY, where multiple scattering effects are more pronounced.

2.3.2 Sensors bent along the long edge

Building on the findings from the ALPIDE sensor bent along its short axis, a new series of measurements was conducted, this time focusing on ALPIDE sensors bent along their long axis.

This approach allows to study the effects of decompressing the circuitry on top of the sensor, as well as examining the impact of bending in this alternate direction.

To accommodate sensors bent along the columns, a different mounting procedure was required. The standard carrier card could not be used, necessitating the redesign of the interface between the sensor and the carrier card. The solution involved the use of an FPC cable, an I-board connected to an FFC, and an adapter board, as outlined in Sec. 2.1.3.

2.3.2.1 Investigating various radii

The subsequent tests focused on the v1 and v2 bent ALPIDE sensors, which were bent to various radii and tested in both laboratory and testbeams. The laboratory results mirrored those observed with the v0 bent sensor, showing no discernible difference from the flat state: the number of noisy or non-functional pixels remained consistent before and after bending.

For a first testbeam campaign three v1 ALPIDE sensors were bent to the three radii initially planned for the ITS3¹ of 18, 24, and 30 mm. Each sensor was mounted on a separate 3D-printed jig and positioned along the same side of the cylinder. Behind each sensor, an aperture of 17 mm × 9 mm was created, providing a region where particles could pass without being affected by the material of the jig, encountering only the silicon sensors themselves and the surrounding air volume.

The stacked setup is shown in Fig. 2.19. On the left side, the outermost bent ALPIDE is depicted with its aperture marked by the blue rectangle, while on the right an X-ray CT scan is presented.

In this initial version of the ALPIDE sensor bent along the long edge, the sensor was attached to the jig using bi-adhesive tape. Due to the natural tendency of the sensor to return to its original shape, the edges began to slightly detach, a phenomenon visible in the CT scan. The degree of edge detachment was more pronounced in sensors with a greater degree of bending.

The FPC of each sensor was connected to the I-board. For this first in-beam measurement in such configuration, this small decoupling board lacked capacitors. The I-board was then connected to an adapter card via a 15 cm long FFC. The adapter boards were equipped with 100 nF capacitors for the digital and analog domain voltages, while the bias voltage was decoupled through a separate line using a low-pass Pi filter board.

¹At the time of the testbeams, the envisioned radii for ITS3 were 18, 24, and 30 mm. These values were later slightly adjusted to 19, 25.2 and 31.5 mm due to stitching constraints (see Ch. 1.4.3 for more details).




Figure 2.19 – Outermost v1 type bent ALPIDE on jig. Other two similar ones are stacked one inside the other. **(left)** X-ray CT scan of the stacked sensors, present on the same side of the cylinder. Scan taken by L. Lautner.

The bent sensors were integrated into a particle telescope consisting of four ALPIDE chips, used as reference planes, with two sensors upstream and two downstream of the bent chips. The setup is illustrated in the sketch in Fig. 2.20. Due to limited availability of DAQ boards, only two reference ALPIDE sensors were used per reference arm, as the telescope operated in parallel with two others. The jigs holding the bent sensors were rotated counterclockwise by 90° around the z axis.




Figure 2.20 – Sketch of the experimental setup where three bent sensors of various radii were tested.

The exercise took place at the DESY II testbeam facility at beamline 24 where a beam of 5.4 GeV electrons was directed at the setup.

A similar trigger logic to that described in Sec. 2.3.1 was employed. Upon receiving a valid trigger signal, a trigger ID was assigned to each event, facilitating reconstruction later. The scintillators had a size of 4.5 cm \times 2.5 cm, slightly larger than the surface of the reference ALPIDEs. On average about 0.7 tracks per event for nominal settings (chip threshold \sim 100 electrons) are recorded as a consequence.

The data acquisition was done using the EUDAQ2 software, with subsequent analysis performed using the Corryvreckan reconstruction framework.

Events are built from hits corresponding to the same trigger ID on all planes. Pixels that recorded activity more than 50 times the average across the entire active matrix are considered noisy and are masked.

When a particle interacts with the sensor, the deposited charge can spread due to the diffusion-dominated charge collection mechanism of the ALPIDE sensor. Adjacent pixels to the one being hit can also collect charge if it exceeds the set threshold, leading to the formation of clusters. The position of these clusters is determined by the geometric center-of-gravity of the associated hits.

Initial corrections for gross misalignments in the x and y directions were made by analyzing the spatial correlations between clusters on all telescope planes.

Track candidates were then reconstructed by fitting clusters in the reference planes using a GBL fit. This fit accounts for uncertainties arising from multiple scattering within the sensor layers, jigs, and the surrounding air volume, as well as measurement uncertainties due to the finite pixel sizes. Typically, only one cluster per track is needed to form a track, which is true for the majority of events given the size of the reference ALPIDEs, the trigger logic, and the low event rate (approximately 500 Hz) at the chosen momentum at the DESY II beamline.

Tracks from the reference ALPIDEs were subsequently processed using the Millepede II algorithm for alignment. The alignment parameters were incrementally unfrozen to maintain better control over the procedure. Initially, only shifts in the x and y directions were allowed. This was followed by permitting only rotations around the z axis in the second iteration. In the third iteration, both x and y translations and z rotations were enabled. Finally, translations along the two main axes and rotations around all three axes were allowed until the fit converged.

With the reference planes aligned, only tracks with a $\chi^2/\text{ndf} < 3$ were selected for further analysis. The mean values of the residuals were below 1 μm , and the residuals were Gaussian-shaped with an RMS in the range of 5 μm to 6 μm .

In the subsequent step, the bent ALPIDE sensors were sequentially added, one at a time, to the tracking and alignment algorithms. For each event containing a track, the cluster closest to the track intercept on the bent sensor surface was associated with the track. Initially, the association window was a circle of radius 500 μm around the track intercept point, but through successive alignment iterations, this was reduced to 50 μm .

The distance between the associated cluster and the track intercept per plane was used to align the bent DUTs. These bent sensors were iteratively shifted and rotated to minimize this unbiased residual. Initially, only translations in the x and y directions were allowed. Subsequently, rotations around the z axis and, later, around the x and y axes were also corrected. Finally, the radius of the sensor was varied by up to ± 1 mm, and the procedure was repeated. While significant changes were not observed for the radius in the middle region of the sensor in terms of residuals, some improvements were noted at the optimal radius near the sensor edges.

Ultimately, tracks with clusters on all telescope planes were used to perform a final alignment of the translations and rotations of all sensors using the Millepede II algorithm. This final step was conducted to eliminate any remaining local minima in the alignment of the individual sensors.

Since the bent setup is rotated by 90° to accommodate the FPC connection (as explained in Sec. 2.1.3), this rotation reduces the area where tracks can pass through all the sensors, as illustrated in Fig. 2.21.

Rotating the DUTs by 90° around the z axis results in a fiducial area where tracks pass through all chips (assuming perfect alignment) of approximately $13.76 \text{ mm} \times 13.76 \text{ mm}$. This area corresponds to a square with a side length equal to the short edge of the active region of an ALPIDE sensor, as the periphery regions, which do not record particles, are inactive. In the actual setup, slight misalignments of the sensors further reduce this analysis area.

Moreover, as is evident in the CT scan, parts of the bent ALPIDE sensors are shown to have one or more of the jig layers present behind a small part of the active matrix. This material serves as a scattering medium, potentially altering the trajectory of the tracks.

To visualize the scattering effects from the surrounding material more clearly, the material budget of the sample is probed using a procedure later detailed in Ch. 3. In summary, tracklets are reconstructed separately using the first three upstream devices (two reference and the first bent sensor) and the last two downstream reference planes. The first bent sensor is included in the upstream tracklet reconstruction because it does not introduce bias into the measurement, as it is positioned before any interaction with the jig material.




Figure 2.21 – Due to a 90° rotation around the beam axis of the bent sensors (that accommodates the fragile FPC connections), the fiducial area where tracks intersect both the reference ALPIDEs and the bent sensors is smaller.

These tracklets are then propagated to the middle of the jigs, where an analysis plane is defined, perpendicular to the z axis. At this plane, an arbitrary “kink” in the direction of the two matching tracklets is allowed. The magnitude of this kink angle provides information about the material budget encountered between the last upstream measurement point and the first downstream measurement point. The greater the angle, the more material was traversed along the path. The tracklets in space are projected on the orthonormal x and y axes and two projected kink angles are recorded. The analysis area is divided into bins, and tracklets intersecting a specific bin in the $x - y$ plane are assigned to that bin. After all the tracks are matched, each bin will contain a distribution of kink angles, with the width of the distribution being shown in Fig. 2.22. The width is measured using the absolute average deviation (AAD) of the inner 95% of the distribution (see Section 3.3.2). Essentially, the higher the value on the z -axis, the more scattering occurs at that specific global position.

From the 2D distribution, it is immediately apparent that there is a region of reduced scattering (indicated by a smaller width of the kink angle distribution within each bin) in the center of the analysis area shown in Fig. 2.21.

The dimensions of this region along the global x axis correspond well with the aperture in the 3D-printed jigs, which measures $17\text{ mm} \times 9\text{ mm}$. The 9 mm region with reduced scattering in this direction aligns with the rotation of the bent sensors by 90° relative to the reference




Figure 2.22 – Material budget of the structures forming the bent sensor jigs. The aperture in the 3D-printed jigs of $17\text{ mm} \times 9\text{ mm}$ is clearly visible. Not the full extent along the y_{global} is available, since the bent sensors are rotated by 90° along the z axis with respect to the flat reference ALPIDE sensors.

ALPIDE sensors, resulting in a swap of their global x and y axes. The figures refer to the global coordinate system defined by the flat reference ALPIDE sensors.

In the global y direction, the entire 13.76 mm width corresponding to the short side of the reference ALPIDE sensors is available. A small region of approximately 1 mm , where increased scattering occurs, is visible at the top of the plot. This region corresponds to the area near the bonding pads of the reference ALPIDE sensors, where part of the active matrix is glued to the carrier card, which in turn acts as a scattering medium.

This region of reduced scattering is designated as a region of interest (ROI) and is subsequently used to determine parameters of interest, such as spatial resolution and efficiency.

The first parameter examined was the average cluster size as a function of the sensor threshold. The threshold was measured both during the testbeam and afterward by injecting test charges using a pulsing capacitor into each pixel across the entire matrix and recording the values, as described in Sec. 1.4.2.

As expected, the cluster size decreases with increasing threshold since fewer pixels will have sufficient shared charge to exceed the threshold. However, a discontinuity in this trend was observed for all the data collected at various radii, as shown in Fig. 2.23. This discontinuity was later understood upon investigating the timestamps of the runs where the discrepancy occurs.

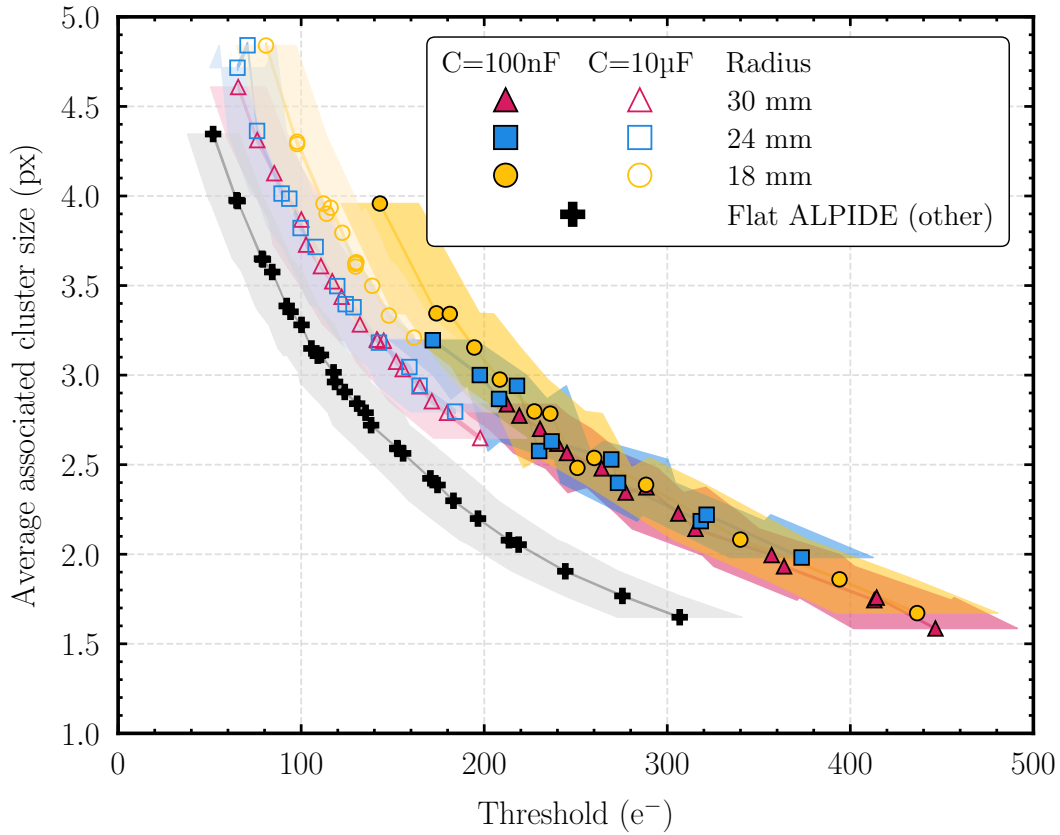


Figure 2.23 – Average associated cluster size as a function of the set matrix threshold for the three bent sensors, as well as a comparison with a flat ALPIDE, other than the three sensors bent here.

At the time of the testbeam, the setups with ALPIDE sensors bent along the long edge were still in the early design stages. Significant levels of noise were observed at low threshold values during the testbeam, exceeding what was expected for the operating settings of the sensors. In an attempt to mitigate this noise, the capacitances on the adapter boards were adjusted. Initially, the capacitances were approximately 100 nF. An additional 10 μ F capacitor was placed in parallel to both the digital and analog domains. This adjustment moved the noise floor to lower threshold values, thereby allowing access to the domain of nominal ALPIDE settings, which are around 100 electrons.

For this reason, the data in Fig. 2.23 was divided into two groups: before and after the capacitance change. After this adjustment, the data aligned more closely with the expected trend, confirming that the observed discrepancy was primarily due to the change in system capacitance. The addition of the 10 μ F capacitor shifted the curve by approximately 20–40 e^- towards lower threshold values.

In the case of flat ALPIDE sensors, the carrier card features an array of 10 μF capacitors located in very close proximity (1 mm) to the bonding pads of the sensor. For the bent ALPIDE sensors, even though decoupling was implemented at the adapter board level, noise could still be picked up along the 15 cm FFC cable, the I-boards, and the FPC.

It was only later during this analysis and when comparing to a flat ALPIDE, that it became evident that a shift still persists despite using the same capacitor network as for the flat sensors. This indicates that the true value of the threshold is not yet fully understood and cannot be precisely correlated with the DAC settings used at the testbeam.

This issue is currently under investigation. Numerous tests have been conducted at CERN by the other primary author of [69]. The results of these tests will be detailed in [83]. Among the tests, different capacitances were evaluated for their decoupling effectiveness at various positions along the readout line, including positions very close to the sensor (e.g., on the I-board). Additionally, sensors were bent and then returned to their flat state consecutively [84], yet no significant change in the mean threshold was observed that would account for the difference in the cluster size plot.

To decouple other performance metrics from the effects of the threshold, the following plots are presented as a function of the average cluster size, effectively absorbing the threshold shift.

The first parameter investigated was the detection efficiency of the sensor. As before, for clarity, the inefficiency is plotted instead, with visual markers indicating the corresponding upper limits of the efficiency decade, as shown in Fig. 2.24.

It is evident that the three bent ALPIDE sensors, each bent at different radii, show compatibility not only with each other but also with a flat ALPIDE sensor when plotted as a function of the average associated cluster size.

Previous measurements with both flat and bent ALPIDE sensors along the short edge (mounted on the normal carrier card) indicated that at the nominal operating thresholds (100–200 e^-), an average cluster size of 2.5–3.5 px is expected. At these values, an efficiency exceeding 99.9% was measured, consistent with known results from flat ALPIDE sensors [85] and with results from ALPIDE sensors bent along the short edge (see Figs. 2.17 and 2.18).

These results suggest that the efficiency of the sensors does not depend on the bending radii, regardless of the operating point of the chip. As shown for the ALPIDE sensor bent along the short edge (see Fig. 2.18), larger cluster sizes can be measured at positions corresponding to larger incident angles due to the curvature of the sensor. Here, an increase in efficiency is expected due to enhanced charge sharing and increased charge deposition from the longer




Figure 2.24 – Detection inefficiency as a function of the average associated cluster size for the three bent sensors at different radii, as well as the data from a flat sensor, other than the ones tested during this testbeam. Dotted lines mark the upper limits of the efficiency.

path of particles through the epitaxial layer. This effect is more pronounced for sensors bent at smaller radii, where larger incident angles are accessible.

The next parameter of interest is the position resolution. Like the efficiency, this performance metric is also evaluated within the ROI, where no scattering from the window is present. The unbiased residuals are calculated as the distance between the associated cluster position on the bent sensor and the track intercept position on the surface of that sensor, excluding the DUT from tracking. Consequently, the distribution of all residuals represents a convolution of the actual space point resolution and the intrinsic uncertainty from the propagation of reference tracks.

This tracking error is estimated from Monte Carlo simulations using online calculators for the respective track models, yielding a value of $(3.15 \pm 0.30) \mu\text{m}$ [65, 66]. The uncertainty accounts for the momentum spread of approximately 150 MeV at the DESY II beamlines [86] as well as the uncertainty of the alignment procedure.

This value is subtracted in quadrature from the standard deviation of the residuals distribution of each sensor in order to obtain the resolution. The position resolution in both the bending direction and the non-bending direction is shown in Fig. 2.25.

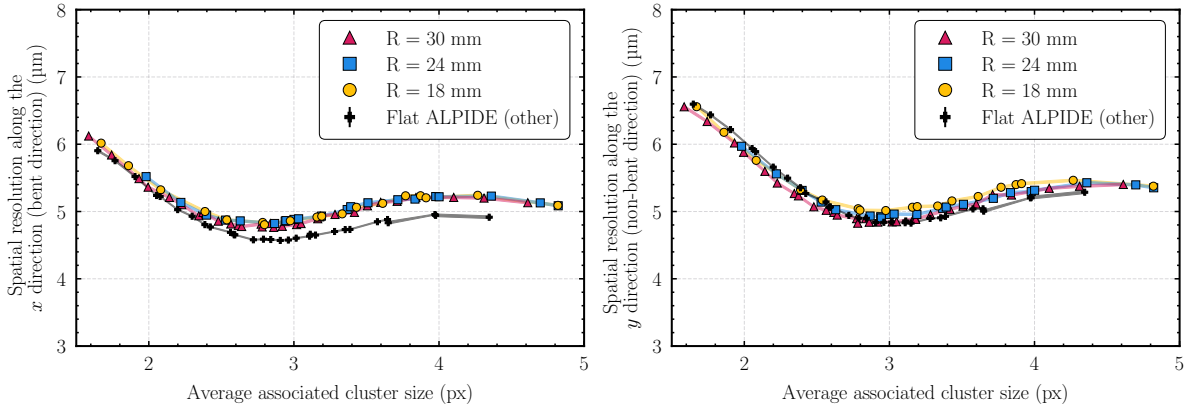


Figure 2.25 – Position resolution as a function of the average associated cluster size along the **(left)** bent direction and **(right)** the non-bent one.

Similarly to the detection efficiency, no significant dependence on the bending radius is observed. Moreover, the results are comparable to data obtained from flat ALPIDE sensors. The minimum resolution in both cases is observed for an average associated cluster size between 2.5 and 3.5 pixels, as expected from a diffusion-dominated digital sensor. At this regime, the ALPIDE sensor was shown to be more than 99.9% efficient.

A small average cluster size represents the case where the chip threshold is high. This leads to a worsening of the resolution because the efficient pixel area starts to shrink (the pixel edges become less efficient). Conversely, a large average cluster size represents the case where the chip threshold is low. In this case the center of gravity of the cluster no longer accurately constrains the impact point and a slight worsening is visible.

In order to have a direct comparison with the sensor threshold, another testbeam campaign was conducted and analyzed.

In this campaign, six ALPIDE sensors were bent to the three radii initially envisioned for the future ITS3 (18, 24, and 30 mm). Two sensors were placed on each jig, positioned on opposite ends. Each sensor had an aperture behind it measuring 17 mm × 9 mm, creating a region where particles from the beam could pass without being influenced by the material budget of the jig, interacting only with the sensors and the surrounding air volume.

This compact setup, dubbed the μ ITS3, is shown in Fig. 2.26, alongside an X-ray computer tomography scan performed with a ZEISS METROTOM 1500 CT scanner.




Figure 2.26 – (left) The μ ITS3: six ALPIDE sensors, bent two per layer, to the planned radii of the ITS3 (18, 24 and 30 mm). (right) An X-ray CT scan of the setup done by L. Lautner. The sensors are indicated by the red arrows. Sensor edge detachment is still present and is indicated with an orange arrow.

The six FPCs extending from the sensors are attached to a second version I-board in the form of a wheel, which can accommodate all six devices and is visible at the bottom of the picture. The CT scan shows the six sensors, marked by red arrows, along with best-fit radii given by the CT machine software, where the radius of each sensor and the center point of the cylinder are free parameters. The sensors still show a tendency to detach (as indicated by the orange arrow in the image), causing the radii to be slightly larger than the nominal values.

The μ ITS3 was tested at CERN SPS using a beam of mixed hadrons (60–70 % pions, 25 % protons, and 5–15 % electrons and muons) with an energy of 120 GeV impinging on the setup. A render of the sensors without any surrounding material is also provided in Fig. 2.27.

The setup consists of twelve sensors operating at the same time. Six flat reference ALPIDE sensors – three in the upstream arm and three in the downstream arm – reconstruct the passage of particles upon receiving a trigger from the coincidence of two scintillators. A trigger logic similar to that described in Sec. 2.3.1 is used here. Upon receiving a valid trigger signal, a trigger ID is assigned to each event, simplifying the subsequent reconstruction process. The scintillators have a size of 4.5 cm \times 2.5 cm, slightly larger than the surface of the reference ALPIDEs.

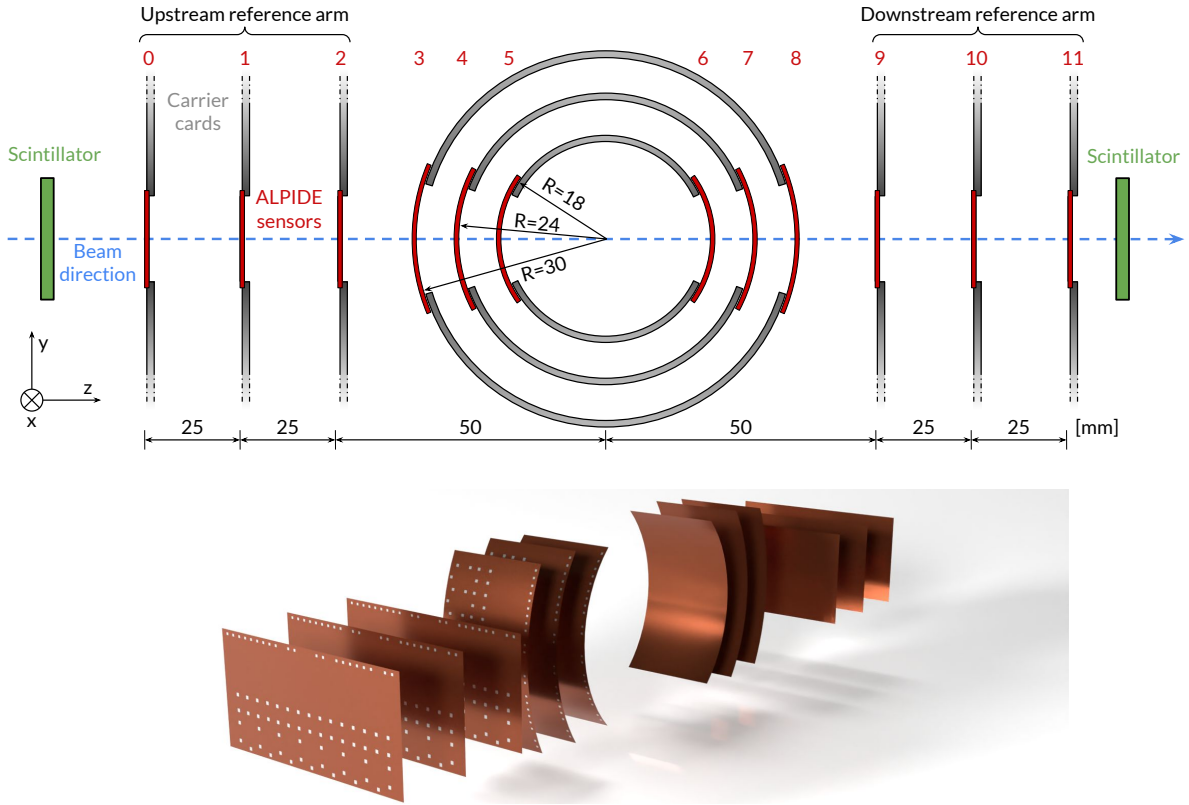


Figure 2.27 – Experimental setup sketch showing the 12 ALPIDE sensors, of which six serve as reference planes and six are bent and treated as DUT. The 3D render was adapted from M. Mager.

Events are reconstructed from hits that share the same trigger ID across all ALPIDE sensors. Pixels that trigger more than 20 times the average rate per chip run are classified as noisy and are excluded from further analysis. Typically, around two pixels per matrix are masked as a result.

Adjacent pixels registering the passage of a particle during an event on each plane are grouped into a cluster. The cluster position is then determined by calculating the geometric center of gravity of these associated hits.

An initial rough alignment is performed by examining the spatial correlations of clusters between planes, focusing on the global x and y coordinates, which are most susceptible to misalignment.

The displacement of the correlation distribution from zero indicates the physical shift of the sensors relative to the reference detector (in this case, ALPIDE 0). This shift is then corrected in the software.

The subsequent alignment of the sensors is conducted in four steps. Aligning all 12 sensors simultaneously, each with some degree of uncorrected misalignment, proved challenging

due to the complexity of the bent geometry, causing the algorithm to converge to a local minimum. Moreover, some unaccounted degrees of freedom contributed to this difficulty. Therefore, a new strategy was adopted, where parts of the telescope are aligned separately before the entire setup is aligned together, allowing the algorithm to perform better with updated initial conditions.

In a first step, only the reference ALPIDEs are aligned relative to the reference plane. Straight-line track candidates are reconstructed from clusters in the six reference planes, with each reference plane required to have only one cluster per event. Tracks that do not pass through the region of interest (ROI) of $13.7 \text{ mm} \times 9 \text{ mm}$ – defined by the low scattering region in Fig. 2.22 – are rejected.

The goal of the track-based alignment procedure is to minimize the χ^2 of track fits and reduce the bias and uncertainty in the fitted track parameters. This is achieved by feeding the tracks to the Millepede II algorithm implemented in Corryvreckan, which performs a simultaneous fit of all tracks and determines the final alignment corrections. Tracks with a $\chi^2/\text{ndf} < 5$ are retained (corresponding to an average distance of less than $13 \mu\text{m}$ between the track intercept per plane and associated cluster position).

In a second step, the first half of the telescope (sensors 0 through 5 from Fig. 2.27) is aligned, with ALPIDE 0 as the reference. Since the first three detectors have already been aligned, the algorithm only needs to align the last three sensors in the upstream arm (sensors 3, 4 and 5).

In the third step, only the second half of the telescope (sensors 6 through 11) is considered, with ALPIDE 9 serving as the reference detector. The algorithm aligns the first three bent sensors (6, 7, 8) while the downstream reference arm remains fixed.

The roto-translation matrix between two pivots – ALPIDE 0 and ALPIDE 9 – is then used to combine the upstream and downstream telescope alignments. This step accounts for potential different local minima in the alignment between the upstream and downstream arms.

Finally, after the separate alignments are done and combined, the alignment constants for all 12 sensors are fed into Millepede II, which, with the improved initial conditions, converges to a better minimum.

For each alignment step, a first iteration aligns only the translations in x and y . A second iteration aligns both translations in x and y and rotations around z . In the final iteration, rotations around x and y are also included. Translations along z are fixed based on the CT scan fit values. The sensor radius is adjusted by up to $\pm 1 \text{ mm}$ from the CT scan fit value until the profile of the mean and RMS of the spatial residuals on the sensor surface is minimized.

To ensure this alignment minimum is correct, a second alignment strategy is used to cross-check the final alignment constants derived from Millepede II.

In this approach, after aligning the reference sensors, each bent sensor is treated as a DUT and aligned individually with respect to the reference tracks. The DUT is iteratively rotated and translated until the unbiased residuals are minimized. As before, this process is repeated for different radii around the CT scan fit value. Initially, only translations in x and y are allowed, followed by rotations around x , y , and z . Finally, both translations and rotations around these axes are allowed simultaneously.

Clusters are associated with the corresponding track if they fall within a circular search window of $50\text{ }\mu\text{m}$ radius around the track intercept on the sensor plane. For a non-aligned sensor, the initial association window is set at $500\text{ }\mu\text{m}$ and is reduced after each alignment step.

After all sensors are aligned, tracking is performed through the entire telescope, yielding similar results to the first alignment method. The normalized track χ^2/ndof , along with the average distance in both x and y directions between the track intercept per plane and the associated hit for one of the bent sensors, is shown in Fig. 2.28, providing a measure of the alignment quality achieved.




Figure 2.28 – Tracking and alignment quality checks. **(left)** The track χ^2/ndof and the **(right)** distance between tracks and clusters.

The cluster size associated with the most upstream bent sensor (ALPIDE 5, with an 18 mm radius) is illustrated in Fig. 2.29. The left side of the plot shows the cluster size as a function of the column (aligned with the bent direction), while the right side shows it as a

function of the row (along the non-bent direction). The data is further divided based on the matrix threshold, binned in increments of 50 e^- , for both the case where the sensor is reverse biased (bottom two subplots) and when it is not (top two subplots).

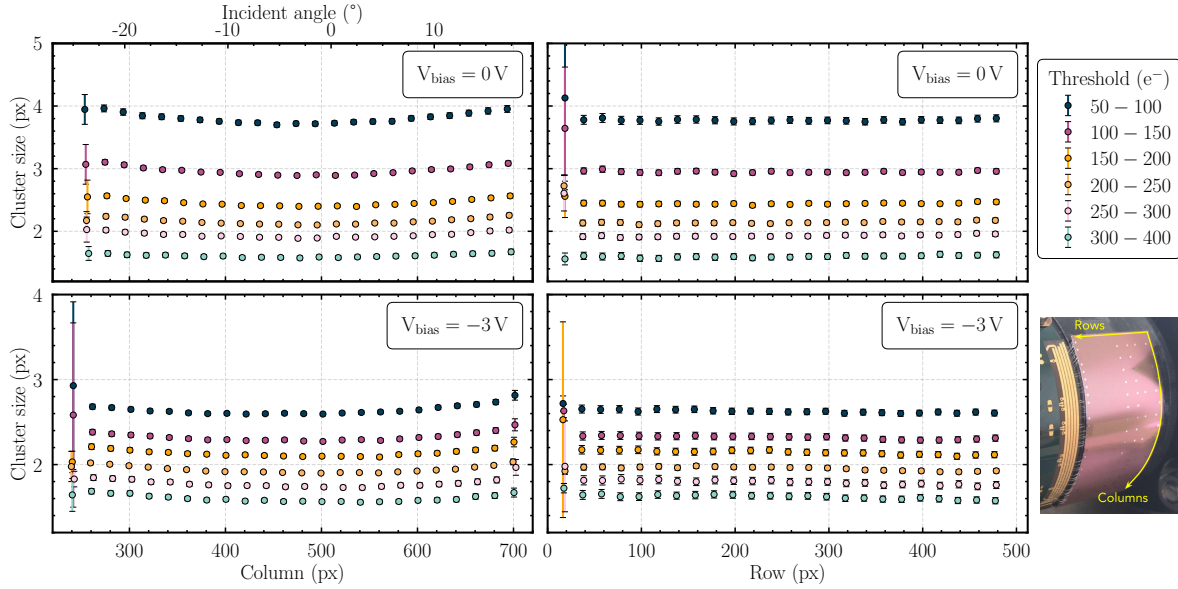


Figure 2.29 – Cluster size as function of the column and row for the ALPIDE 6 sensor bent at radius of 18 mm shown as a function of the threshold and reverse bias value. The incident angle of the particles is also shown for the columns.

It is evident that the curvature of the sensor affects the cluster size along the columns (bending direction), as previously discussed in relation to Fig. 2.18. The point where the incident angle is 0° is identified as the average minimum of the different curves. Due to the presence of the region of interest (ROI), the largest incident angle reaches a maximum of about 23° , at which a particle traverses approximately 8.6% more of the epitaxial layer. This difference is reflected in the average cluster size, which increases by 7–8 %. When the sensor is reverse biased, the average cluster size is about 25% smaller due to the larger depleted area, and the difference between the maximum incident angle and perpendicular tracks reduces to 4–5 %.

Along the rows (non-bent direction), no curvature effect is expected, and this is confirmed by the constant average cluster size across rows, regardless of the applied threshold.

The efficiency of the sensors displays a similar pattern (see Fig. 2.30). For the sensor bent at a radius of 18 mm, the curvature directly influences the efficiency, with tracks passing at an angle detected with higher efficiency compared to those perpendicular to the sensor surface. This effect is most pronounced at high detection thresholds, with a difference of about 5.7% for the unbiased sensor, which is less pronounced when reverse bias is applied.

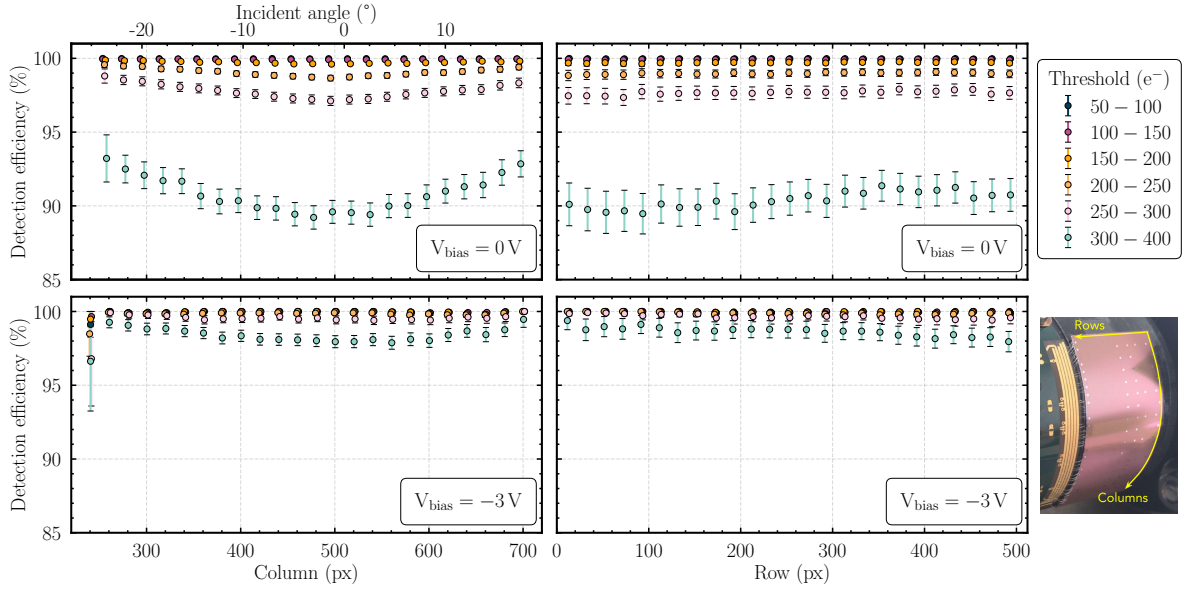


Figure 2.30 – Detection efficiency as a function of the column and row, for different thresholds and reverse bias voltages. Since ALPIDE 5 did not allow to be biased, the results are shown for ALPIDE 6, the next sensor in the path of the particle, also bent at 18 mm.

Figure 2.31 demonstrates the impact of different radii on the efficiency of the sensors at two bias voltages as a function of the matrix threshold. The inefficiency is plotted on a lin-log scale for the three different radii of the upstream bent ALPIDEs. The data is taken from the following sensors: ALPIDE 3 (with a radius of 30 mm), ALPIDE 4 (24 mm), ALPIDE 5 (18 mm) — only the 0 V data is available for this sensor as it could not be back-biased and ALPIDE 6 (18 mm) — which includes only the 3 V back-bias data. Moreover, older data from flat ALPIDE sensors, taken with a different chip than the one being bent here, are shown for reference.

It can be seen that in general the bent sensors are slightly more efficient than the flat sensors, a trend especially noticed at higher chip threshold values. At nominal settings, all sensors are more than 99.9% efficient. Biasing the sensors will increase the depletion of the sensor and bring a boost in efficiency across the whole threshold range.

Figure 2.32 shows the other major parameter of interest, namely the position resolution. Here, the tracking uncertainty estimated using [65, 66] at the position of each sensor under study has been quadratically subtracted from the unbiased residuals.

The position resolution is at or below 5 μm around the nominal operating point of the sensor, where the chip threshold is approximately 100 electrons. At threshold values above 150 electrons, all curves for the bent sensors converge and follow a common trend, showing improved position resolution compared to the flat ALPIDE.




Figure 2.31 – Sensor detection inefficiency on a lin-log scale as a function of the chip threshold for sensors bent at three different radii.




Figure 2.32 – Position resolution along the bent direction (x -axis) as a function of the chip threshold for sensors bent at three different radii.

However, at lower sensor thresholds, the reverse-biased sensors exhibit poorer position resolution that increases sharply. This is thought to be due to the lack of decoupling in the electrical line supplying the reverse-bias voltage during the test beam, potentially creating an additional capacitance that could explain this behavior. Further investigation is underway to confirm this assumption.

To study more exotic geometries, the μ ITS3 was rotated by 90° around the global x axis and placed as shown in Fig. 2.33a.

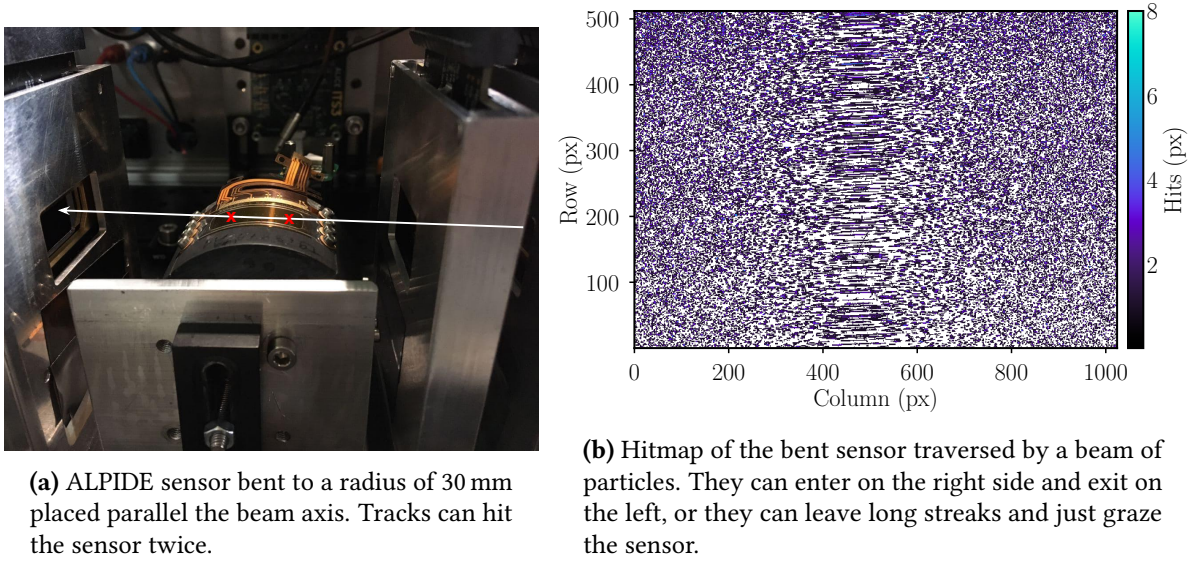


Figure 2.33 – Particles that traverse the bent sensor in a grazing topology.

Particles passing through the sensor in this configuration can follow two distinct paths. They may enter the sensor on the right side, exit at the same point through the back, travel a few millimeters underneath, and then re-enter through the back side, depositing charge before finally exiting again. The two distinct clusters left by these particles is shown in Fig. 2.33b. Also visible is the second case, where near the center of the chip, particles may graze the sensor, leaving long trails as they traverse the epitaxial layer.

Such geometries are possible and of interest for the future ITS3 detector, which will be placed in a magnetic field. Low momentum tracks can curl and hit the sensor twice or even traverse longer distances through the sensor.

Adapting the current tracking algorithm in Corryvreckan to recognize two clusters which are separated by a few mm as belonging to the same track and tracking them proved to be complicated and was ultimately not pursued. Instead, the grazing topology was examined more thoroughly.

The hits belonging to a track on the bent DUT are clustered and the following picture shown in Fig. 2.34 emerges.




Figure 2.34 – Average cluster size distribution over the surface of the chip and the projection along the bent direction. The long particle trajectories through the sensor are clustered using the center-of-gravity method, resulting in a concentration of large clusters near the middle of the sensor.

The long streaks shown in Fig. 2.33b are formed of hits which are collapsed on the central bin during clustering using the center of gravity method. This effectively produces an apparent depletion in the direct neighbouring pixels.

The mean length of these long clusters is about 100 pixels. This can be understood from simple geometric considerations, starting from the chord length $c = 2\sqrt{R^2 - (R - t)^2}$. Here, t (sagitta) is the thickness of the epitaxial layer. The sensor from which the data is shown here had a radius of 30 mm. Therefore, the average length of the chord is expected to be about 85 pixels.

Still, charge deposition also happens in the substrate, and given that there is no hard boundary between the epitaxial layer and the substrate due to out-diffusion, it can be assumed that from up to 10 μm into the epitaxial layer the charge can be collected [87].

Thus, the effective length of the tracks through the sensor approaches 100 pixels. This method could potentially be refined for further investigation and used to quantitatively assess the thickness of the epitaxial layer.

At shallower angles (greater than about 2.5°), tracks will penetrate too deeply into the substrate at the midpoint, resulting in two distinct cluster blobs with no recorded information between them. Beyond this point, tracks will enter on one side of the sensor and exit on the opposite side.

2.4 Inelastic interactions in the μ ITS3 with a target

2.4.1 Experimental setup

A 1.16 mm thick and 5 mm diameter copper (Cu) target housed in a plastic holder was placed at the very center of the μ ITS3 setup. The detectors were exposed to the same mixed hadron beam at the SPS, consisting of 60–70 % pions, 25 % protons, and 5–15 % electrons and muons. Given the composition of the beam, inelastic interactions with the Cu target were anticipated. The objective of this study was to demonstrate the capability of performing track and vertex reconstruction with bent Monolithic Active Pixel Sensors (MAPS) in a geometric setup that closely resembles the future ITS3 detector, comprising three bent layers followed by three flat ALPIDE sensors (out of four in the final ITS3 detector). These results are pending publication [70], where the author played a significant role. The following analysis results follow closely what is already written in this publication.

As with the setup discussed in Fig. 2.27, the μ ITS3 was rotated by 90° around the z axis to facilitate the connection of the FPCs to the adapter boards, minimizing strain on the connectors. The geometry of the setup is illustrated in Fig. 2.35, with acceptance regions indicated by the η lines, as will be later explained.

The trigger system was similar to the one used in previous experiments, initiated by the coincidence of two scintillators positioned before and after the setup. The scintillators, each measuring $4\text{ cm} \times 3\text{ cm}$, had an active area slightly larger than that of the reference ALPIDE sensors. The trigger logic incorporated a past protection time of $30\text{ }\mu\text{s}$ and an event separation time of $100\text{ }\mu\text{s}$, preventing pile-up from closely spaced beam particles and allowing the in-pixel amplifier, which can take several microseconds to reset, to return to baseline. Each event was defined by assigning a unique trigger ID to all the data collectors.

If an inelastic interaction occurred within the Cu target, resulting in multiple tracks emerging from the π -Cu interaction point, at least one track needed to pass through the downstream scintillator to record the event.

Like before, data analysis was performed using the Corryvreckan reconstruction framework, complemented by custom C++ ROOT scripts.




Figure 2.35 – Experimental setup which consists of 12 ALPIDE sensors: six serve as reference planes and six are bent, two by two, to the anticipated radii of the ITS3 (18, 24 and 30 mm). A copper target is placed in the middle of the bent layers which, upon bombardment with hadrons, will produce more tracks downstream. The sensors in this setup are different from the ones used in the analysis of the setup from Fig. 2.27, but otherwise the setup is very similar having only the target inserted.

The alignment process for the setup mirrored previous procedures. Initial alignment was achieved using tracks that recorded a single cluster in each of the 12 planes, assuming these events were primarily without interactions or involved interactions where particles did not significantly alter their trajectory through the downstream sensors (in which case the multiple scattering effects were accounted for in the tracking process). The alignment was refined iteratively, adjusting shifts and rotations to minimize the χ^2 distribution of these tracks until an optimal alignment was achieved.

Figure 2.36 displays the distribution of hit residuals along the bent direction (global y axis) for each of the 12 sensors in the setup. A similar distribution was observed in the non-bent direction. The left panel shows unbiased residuals, where tracks were fitted using only clusters from the flat reference ALPIDE sensors. In contrast, the right panel shows biased residuals, where tracks were fitted using clusters from all 12 sensors.




Figure 2.36 – (Unbiased) and (biased) spatial residuals along the bending direction of all 12 ALPIDE sensors, showcasing the quality of the alignment of the setup.

The residual distributions are centered around zero, have a Gaussian shape and minimal tails. The standard deviation for each distribution is approximately $5\text{ }\mu\text{m}$, with only a few outliers excluded from the fit. This value aligns with the known performance of flat ALPIDE sensors. For biased residuals, the mean width across the 12 sensors is about $4.2\text{ }\mu\text{m}$, comparable to that of the reference planes.

The tracking error, estimated using [65, 66], ranged from $1.6\text{ }\mu\text{m}$ to $1.5\text{ }\mu\text{m}$ for biased residuals (corresponding to layers with radii of 30–15 mm, respectively) and from $2.2\text{ }\mu\text{m}$ to $2.1\text{ }\mu\text{m}$ for unbiased residuals (for the same respective radii order).

Since the tracking uncertainty is excluded from biased measurements, a narrower residual distribution is expected. The unbiased residuals for the flat layers alone are presented in Fig. 2.38 for the bent direction and Fig. 2.37 for the non-bent direction. The single hit resolution is determined by subtracting the tracking error in quadrature from the standard deviation of the residuals, depending on if the quantity is biased or not:

$$\sigma_{\text{unbiased}}^2 = \sigma_{\text{intrinsic}}^2 + \sigma_{\text{track}}^2 \qquad \sigma_{\text{biased}}^2 = \sigma_{\text{intrinsic}}^2 - \sigma_{\text{track}}^2$$

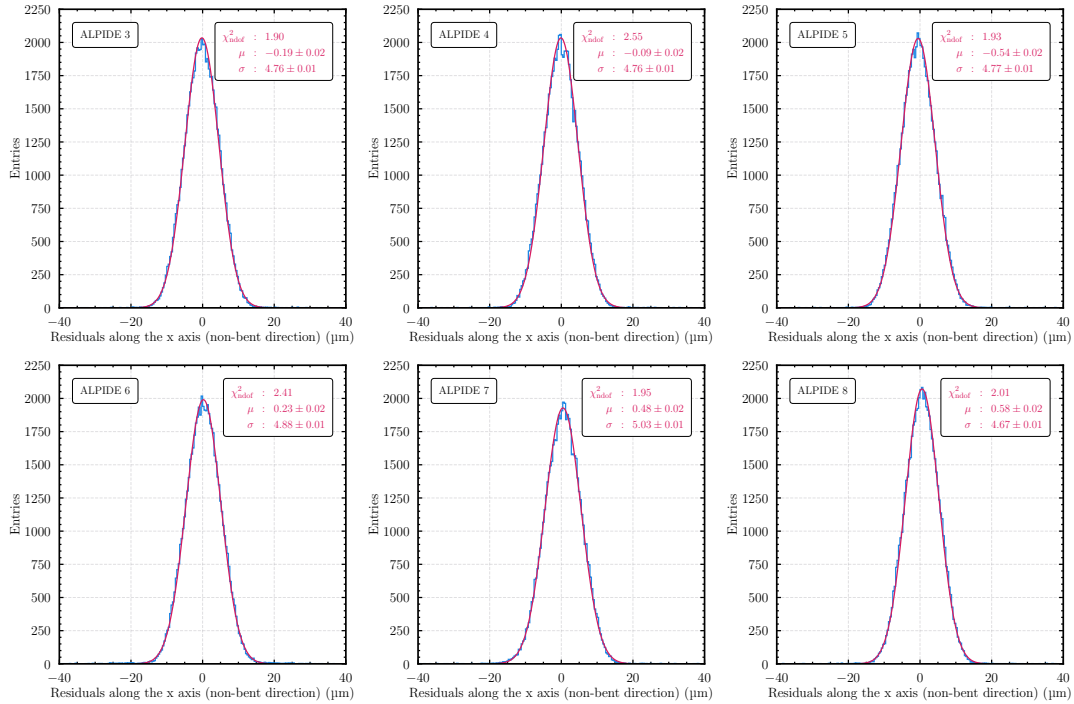


Figure 2.37 – Unbiased residuals on the bent sensors along the rows (non-bent direction).

With the setup now aligned, the analysis strategy shifts focus from particles that pass through the setup without interaction to those that interact with the Cu target. The interesting

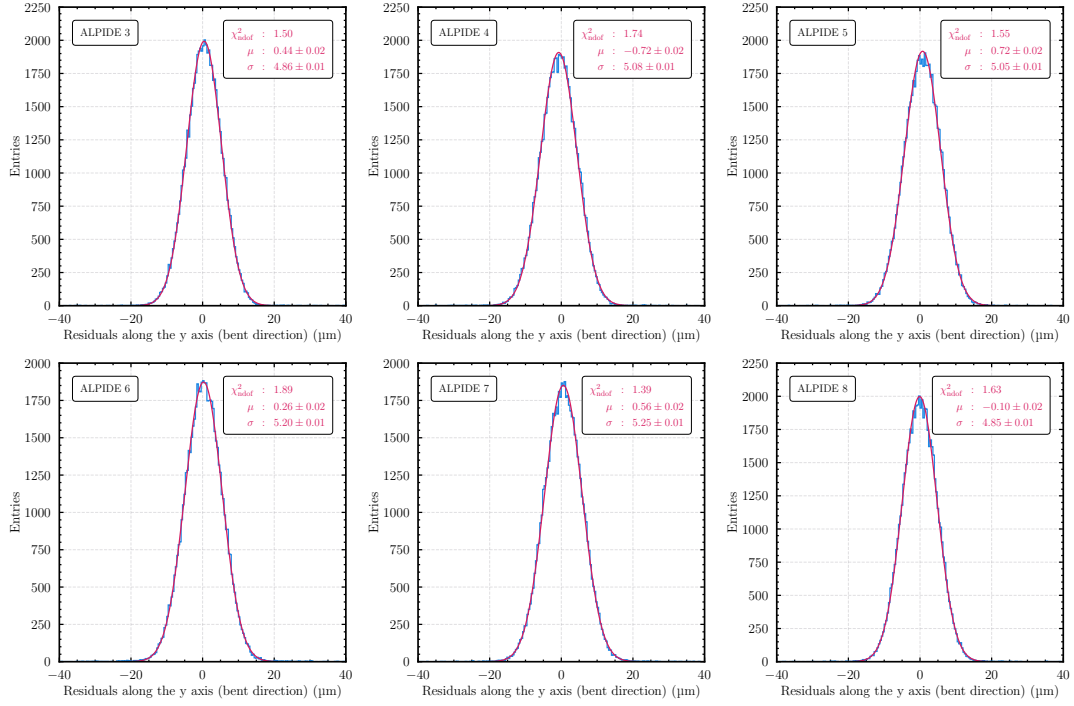


Figure 2.38 – Unbiased residuals on the bent sensors along the columns (bent direction).

events are selected by requiring that there is only one reconstructed track in the upstream arm and at least two tracks downstream of the target. For the downstream portion, each track must have at least three associated clusters in consecutive detectors. If multiple track candidates are possible, the one with the best χ^2 is retained. There are four possible track topologies, depending on their angle with respect to the beam direction, as shown in Fig. 2.35 and explained below.

Tracks emerging from the target may have hits in all three bent sensor layers downstream but might not register hits in the downstream reference ALPIDEs. This occurs because the bent ALPIDEs are rotated by 90° around the z axis, extending their active area further along the global y axis.

Track topologies are categorized according to the pseudo-rapidity (η) ranges depicted in Fig. 2.35. The η observable is defined based on the angle θ that the tracks form with the global z axis:

$$\eta = -\ln \left[\tan \left(\frac{\theta}{2} \right) \right]$$

The η ranges also account for measured misalignments of the sensors and the fact that the Cu target is shifted by about a millimeter in the positive global y direction, as will be shown

later. This asymmetry is reflected in the slightly different ranges between the top and bottom with respect to the beam axis.

Tracks that pass only through the three bent sensors can do so for $1.1 < \eta < 3$. This topology is further referred to as a 3-hit track. If tracks have an η value smaller than 1.1, they will pass through only two sensors and will not meet the 3-cluster-per-track requirement imposed during track fitting. If they have an η value greater than 3, they will hit ALPIDE 9 and register an additional hit, becoming a 4-hit track (or more).

In most cases, these 3-hit tracks are accompanied by another track at large η values, necessary to trigger the downstream scintillator and ensure the coincidence trigger and event readout.

Additionally, 4-hit and 5-hit track topologies are possible, occurring within an η range of 2.42 to 3.8.

Tracks with an η greater than 3.8 will pass through all six downstream sensors. The number of tracks for each topology previously discussed is shown in Fig. 2.39. Most of the tracks are the ones that pass through all six downstream sensors. Tracks with 4 and 5 hits are suppressed, because of the small acceptance in which such a topology is possible.

The χ^2/ndf distribution of all the upstream (incoming) and downstream (outgoing) tracks is shown in Fig. 2.40. All track topologies are included.

In this plot, no cuts are applied to the tracks. For the incoming beam particles, the distribution peaks at a χ^2/ndf of 1, with very few entries exceeding a value of 2. For the outgoing particles, the distribution also peaks at a χ^2/ndf of 1, but it exhibits a long tail. Approximately half of the events have a χ^2/ndf greater than 4, with overflow included in the total entries shown in the legend of the plot. The reason for this skewed distribution is that many outgoing tracks have low momentum and undergo further interactions with some or all of the jig layers downstream. This extra material budget is not accounted for in the tracking algorithm. The only material that is accounted for is the target itself, the silicon sensors and the air between them.

The relationship between η and the momentum of the particles was derived from a PYTHIA8 simulation, where protons and pions (representing the hadronic component of the incoming beam) interact with Cu nuclei. After the interaction, the average momentum was estimated in η bins, as shown in Fig. 2.41. This is an average computed from all tracks from all simulated events contained in a given η bin.




Figure 2.39 – Number of outgoing tracks for each topology.

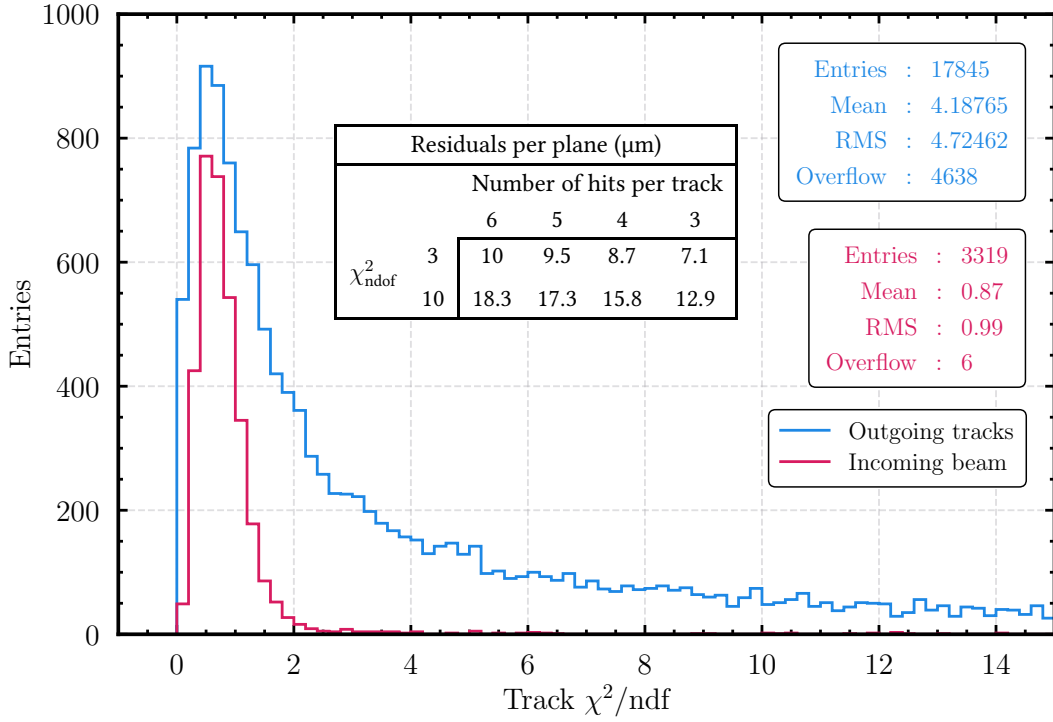


Figure 2.40 – The track χ^2/ndf distributions for incoming and outgoing (from the Cu target) tracks. A table that shows the correspondence between the χ^2/ndf for the example values of 3 and 10 and the residuals value per plane, for a number of hits per track is also provided (for example, a track that is formed by six downstream hits and has a χ^2/ndf cut of 3, will have on average a distance per plane between the track intercept and the hit of $10\mu\text{m}$).

The origin of the tracks with a high χ^2/ndf was investigated by applying a cut-off at a value of 5 and selecting only tracks that are above this selection criterion. The intercepts of these tracks are plotted as a function of their position on the $x - y$ plane at a position along the z axis corresponding to the first downstream flat sensor (ALPIDE 9), as shown in Fig. 2.42.

This figure includes two superimposed rectangles: the larger red rectangle represents the active area of ALPIDE 9 (approximately $30\text{ mm} \times 13.8\text{ mm}$), while the smaller blue rectangle indicates the opening in the jigs projected onto this flat sensor, which is about $19.5\text{ mm} \times 9\text{ mm}$. This size is slightly larger than the opening mentioned in Section 2.1.3 for the v0 and v1 3D-printed jigs. The variation occurs because each jig has slightly different openings, and the inner and outer openings differ by a few millimeters, a fact better illustrated in the figure provided in Appendix F).

It can be observed that most of these tracks originate from regions where the tracks intercept the 3D-printed jig layers of the μ ITS3. This observation gives confidence in the




Figure 2.41 – (simulation) Average track momentum as a function of the pseudorapidity (η) for events with more than two outgoing tracks. The detector acceptance starts at an $\eta = 1.3$. An example slice for $\eta \in 4.3 - 4.4$, displayed as a figure inset, shows the distribution of the momentum.




Figure 2.42 – Position of tracks with $\chi^2/\text{n df} > 5$ on the surface of ALPIDE 9. A red rectangle shows the active area of the flat sensor and the blue rectangle shows the opening on the jig projected onto the surface of ALPIDE 9.

chosen track quality cut to provide a statistically significant sample while rejecting tracks known to originate from interactions with the jig layers, which could otherwise impact the measured parameters.

The shape and size of the target are reconstructed by selecting events where at least two outgoing tracks are associated with one incoming track. All these tracks are propagated onto an analysis plane at global $z = 0$. The result is depicted in Fig. 2.43.




Figure 2.43 – The target shape and size is reconstructed by propagating incoming (left) and outgoing (right) tracks at the position of the target.

The target, with its known radius of 5 mm, is clearly identifiable. As previously noted, the target is slightly shifted relative to the global coordinate system, with a shift of approximately 1.3 mm in the positive global y direction and 0.25 mm in the negative global x direction.

The thickness of the target can also be determined by calculating the point (vertex) along the incoming track that simultaneously minimizes the distance between this vertex and the point of closest approach (PCA) on each outgoing track. This minimization is achieved using the following equation:

$$\chi^2 = \sum_{\text{tracks}} \frac{d^2}{\sigma_d^2} = \sum_{\text{tracks}} \frac{(\vec{pca} - \vec{V}) \cdot (\vec{pca} - \vec{V})}{\sigma_d^2} \quad (2.1)$$

where $\vec{V}(x, y, z)$ is the vertex position, \vec{pca} is the point of closest approach and σ_d^2 is the expected error evaluated with the error propagation of the line equation by using the fitting parameters and errors.

The minimization is performed on events with at least three outgoing tracks. The distribution of the z positions of the reconstructed vertices is shown in Fig. 2.44b.

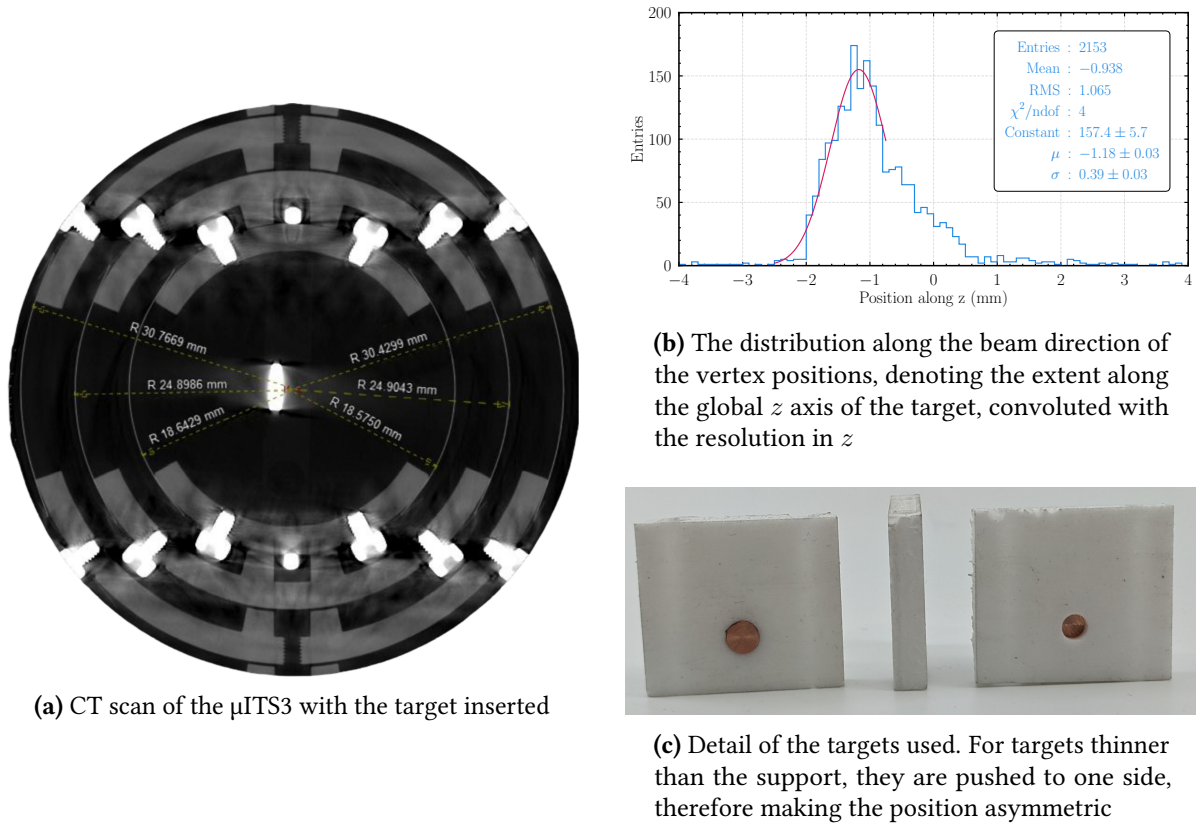


Figure 2.44 – Details about the Cu target position inside the μ ITS3

The distribution indicates that the target is slightly shifted towards the negative global z direction by about 1.2 mm. This finding is corroborated by the X-ray CT scan on the right (see Fig. 2.44a), where a slight displacement and tilt are visible. This displacement occurs because the target holder, designed to accommodate targets of various thicknesses, did not have the target placed in the middle. The holder is housed along rails in the middle of the setup, but the target was pushed to one side, as shown in Fig. 2.44c. The CT scan reveals a displacement of 1.26 mm from the center of the target to the geometrical center of the setup. A similar displacement was measured using a caliper.

The thickness of the target can be inferred from the standard deviation of a Gaussian fit between -2.5 and 0 . A longer tail is observed on one side of the distribution, which is currently not fully understood and under study, but it appears to be stemming from the vertexing algorithm.

The standard deviation is a convolution of the target thickness and the tracking resolution along z . The resolution in z is not known, but is assumed to be in the order of a few hundred

μm . Such a worse resolution is expected from tracks passing at shallow angles with respect to the beam axis.

Given that the interaction probability is uniform across the entire thickness of the sample, the standard deviation is multiplied by $\sqrt{12}$. From this, a thickness of (1.35 ± 0.11) mm is calculated. Considering the resolution in the z direction, this result is consistent with the visually inspected sample, the CT scan images, and the measured target thickness of (1.16 ± 0.01) mm.

An example of an event where an incoming hadron interacted with the Cu target, producing multiple outgoing tracks, is depicted in Fig. 2.45. From this single incoming track, eight outgoing particles are generated.




Figure 2.45 – Example of inelastic interaction. One incoming hadron interacts with the Cu target and generates eight outgoing tracks. The hits on the individual sensors are shown, as well as the sensors themselves and the track fits.

One of these outgoing tracks passes exclusively through the three downstream bent layers, at an approximate η of 2.9. The remaining seven tracks traverse all six downstream sensors, with η values ranging between 4.4 and 5.3.

The η distribution of the tracks in the μ ITS3 setup is shown in Fig. 2.46.

In yellow, the reconstructed incoming beam tracks are shown. These tracks are reconstructed in the first six sensors, before the particles interact with the Cu target. Due to a small, non-zero incident angle of the incoming beam with respect to the global z axis, these particles exhibit large η values.

All outgoing tracks generated by inelastic interactions with the Cu target are displayed in blue. The majority of these tracks have η values below 8. The lower limit of the range, starting around an η of 1.5, represents the acceptance limit of the three downstream bent ALPIDE sensors. Tracks with η values below this threshold typically scatter upon hitting the 3D-printed jigs, causing them to either fail the track quality selection criteria or pass through only two ALPIDE sensors, which disqualifies them from being considered valid tracks.




Figure 2.46 – Distribution of incoming and outgoing tracks as a function of their pseudorapidity (η).

A smaller blue peak, occurring at η values between 8 and 11, represent tracks that pass with minimal scattering through the target. These tracks are the ones that accompany the small η tracks, triggering the downstream scintillator, as previously discussed.

The tracks shown in green are those that pass the track quality selection criterion. These primarily include the large η tracks that pass through the silicon sensors only, as well as parts of the larger blue peak at smaller η values. Two distinct peaks are observed here.

The first green peak, between η values of 1.8 and 3, consists entirely of tracks reconstructed from hits in all three downstream bent layers. This is consistent with the η ranges provided in Fig. 2.35.

The second peak, from η 4 to 7, mainly consists of tracks with a hit in all six downstream sensors. The particles in this case scattered in the target and exited at a slightly larger angle than particles that do not interact at all.

The absence of tracks that pass the selection criteria between η values of approximately 3 to 4 can be attributed to their poor quality. It is suspected that some rotation of the jigs and misalignment of the downstream reference ALPIDEs causes this, resulting in the low number of 4 and 5-hit events observed in data (see Fig. 2.39), coupled with their rejection due to track quality cuts. This issue is currently being investigated together with the authors of [70] in order to reduce the number of rejected tracks without significantly impacting the distance of closest approach measurements.

The final green peak at $\eta > 8$ is entirely contained within the batch of selected tracks and mainly consists of tracks interact only minimally with the target.

Events with tracks reconstructed from hits in all 12 ALPIDE layers are overlaid in red. They have a similar shape and range as the incoming beam peak (yellow). In this case no interaction occurred between the particle and the target. Minor differences are visible between the last green peak and the red one, because particles forming the last green peak come out of the interaction with less momentum and therefore have a shift to a smaller η .

In a final step, the distance of closest approach (DCA) is estimated to assess the pointing resolution capability of the μ ITS3 system.

The DCA is a crucial metric in particle physics, used to determine the proximity between either a track and a specific point (typically the primary vertex of a collision) or between two tracks, or to positively identify primary tracks. It serves as a key indicator of the resolution of the tracking system, in terms of the impact parameter of a track. This is vital for studying physics signals that involve a secondary vertex close to the interaction vertex, such as those arising from the weak decays of beauty and charm hadrons. The DCA is often used as a topological selection criterion to reduce combinatorial background in such studies. It is influenced by the material budget, any misalignment in the tracking system, the intrinsic point resolution of each tracking layer, the number of planes and the lever arm.

To demonstrate the tracking capabilities of the μ ITS3, the outgoing tracks generated from interactions with the Cu target are used to estimate the DCA distribution. A momentum-averaged DCA resolution is computed by comparing the η results with PYTHIA simulations to infer momentum, since the absence of a magnetic field means no direct momentum information is available for the tracks passing through the setup.

The DCA distributions to the primary vertex of tracks reconstructed from the standalone downstream bent sensors, as well as tracks with hits in all six downstream planes, are shown in Fig. 2.47. Only tracks that have a $\chi^2/ndf < 5$ are used to compute the DCA.

On the left side of the figure, the distribution for the beam remainder component (the third green peak from Fig. 2.46 at $\eta > 7.5$), which contains tracks reconstructed from hits in all downstream planes, is displayed. This distribution is fitted with a Gaussian, and the standard deviation is provided. These tracks are mostly high-momentum tracks, so the distribution is primarily influenced by the spatial resolution of the detector.

Since the DCA resolution for the selected events is the same for both incoming and outgoing tracks – representing two independent measurements – the final value is reduced by a factor of $\sqrt{2}$. The measured resolution for 6-hit tracks in the μ ITS3 is approximately $(3.71 \pm 0.09) \mu\text{m}$. This result aligns with the impact parameter formula (eq. 17) from [88], which predicts a value of $4.76 \mu\text{m}$.

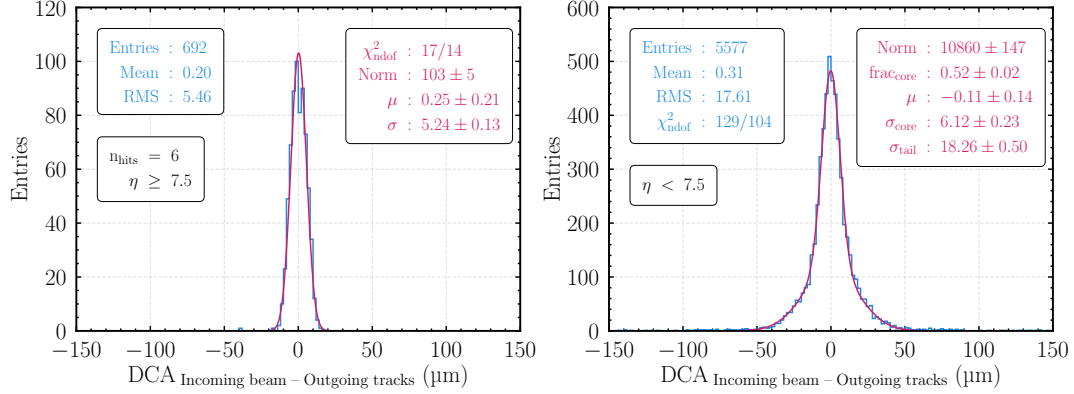


Figure 2.47 – DCA distributions to the primary vertex (**left**) for tracks that pass through all six downstream sensors at an $\eta > 7.5$ and (**right**) for tracks that have an $\eta < 7.5$.

Excluding the high- η tracks, a similar DCA distribution is constructed on the right side of Fig. 2.47 for the two main blue peaks in Fig. 2.46. This distribution is fitted with a two-component Gaussian function to capture both the central and tail resolution components.

The resolution, averaged over the momentum of the tracks, track length, and hit multiplicity, is approximately $(4.87 \pm 0.24) \mu\text{m}$ for 52% of the outgoing tracks forming the core of the distribution, and $(17.9 \pm 0.5) \mu\text{m}$ for the remaining tracks forming the tails. From this distribution, the DCA resolution is estimated for each track hit multiplicity to separate the influence of the track topology. For tracks that register hits in all three bent layers, the DCA resolution with the $\chi^2/ndf < 5$ is measured to be $(15.78 \pm 0.29) \mu\text{m}$, while for events with hits in all the downstream planes, it is $(4.29 \pm 0.14) \mu\text{m}$. Very few 4- and 5-hit tracks survive the cut, and this is currently under further investigation.

These results are qualitatively compared to the ITS3 standalone simulations from [55] and shown in Fig. 2.48.

The measured DCA of approximately $4 \mu\text{m}$ for tracks that pass through all six downstream sensors corresponds to tracks with η values between 4 and 11 (as seen in the second and third green peaks in Fig. 2.46). The relation between these η ranges and the average momentum is derived from the Toy MC shown in Fig. 2.41. These tracks generally have an average momentum above 10 GeV. At this momentum, the DCA evaluated using a fast analytic tool (FAT) in the ITS3 TDR plot is also around $4 \mu\text{m}$, which is fully consistent with the findings presented here.

For tracks that pass through only the three bent layers downstream, the measured DCA is $(15.78 \pm 0.29) \mu\text{m}$. These tracks fall within an η range of 1.8 to 3 (first green peak in Fig. 2.46). The average momentum for these tracks ranges from 1 to 5 GeV. In this momentum regime,




Figure 2.48 – (simulation) DCA (impact-parameter) resolution in the $r\phi$ direction for primary charged pions with $|\eta| < 1$ as a function of the transverse momentum for the (entire) ITS2 and 3 standalone tracking, obtained with a fast analytic tool (FAT). Plot taken from [55] (Fig. 7.5).

the ITS3 TDR MC simulation reports DCA values between 8 to 20 μm . The measured value in this study is found within this range.

A second qualitative comparison can be made with results from the strangeness tracking sector, which involves the direct detection of charged baryons before their weak decay, emphasizing the importance of secondary vertex reconstruction.

Simulations of the ITS3, combining tracks formed in its inner layers, report pointing resolution values in the order of 30 μm for strange baryon tracks formed by three or more hits (as seen, for example, on page 7 in [89]). These values are comparable to the measurements obtained with the μ ITS3 in the case of tracks formed using hits in only the three downstream bent layers. The primary difference arises from the fact that in the strangeness tracking simulation, strange baryons are tracked using tracks formed by three or more hits in consecutive layers [90]. Additionally, the simulation accounts for the beam pipe contribution and tracks under a magnetic field.

Although these comparisons are qualitative, the consistency of the absolute numbers reinforces confidence that the setup configuration enables precise and efficient particle tracking and vertex reconstruction in hadronic environments.

Further improvements in the alignment of the setup are planned to reduce the number of tracks rejected after the χ^2/ndf cuts. Many of the current limitations (e.g. the 3D printed jigs, the bent sensors being rotated with respect to the reference flat sensors) were in place to ensure the mechanical stability of the system during the testbeam. The final ITS3 detector will not be constrained by these mechanical limitations.

2.5 Discussion

Recent advancements in silicon imaging technology for consumer applications have opened new avenues for innovative detector designs, significantly enhancing the potential for physics research.

Building on these state-of-the-art developments, the ITS3 project proposes a new vertex detector design featuring curved, wafer-scale, ultra-thin silicon sensors arranged in cylindrical layers. This would be a first in the HEP community and sets the foundation for nearly-massless detector designs that can be positioned in very close proximity to the collision point.

The new detector will replace the three inner layers of the current ITS2. The first layer of the ITS3 detector will be situated just 19 mm away from the interaction point, with a material budget six times less than that of the ITS2 inner layers. This proximity and reduction in material are expected to significantly improve tracking precision and efficiency at low p_T . With these enhancements, the detector will improve the measurement precision of several observables in the heavy-flavour sector, have a significant impact at low p_T , and bring a new set of fundamental observables within reach [55].

With such a rich physics objective, it is essential to ensure the detector performs as required. Since bent sensors represent a novel approach in the HEP field, an extensive array of measurements is necessary to thoroughly test their characteristics.

In preparation for this, 50 μm thick ALPIDE sensors, which share some of the characteristics with the final ITS3 sensor (excluding changes in the technology node with a resulting lower analogue current, and different pixel pitch), were bent in various configurations and tested both before and after the bending process.

The author played a significant role in this scientific journey, as part of a broader effort within the ITS3 community. While the bending and bonding of the ALPIDE sensors were carried out by ITS3 colleagues, the author was deeply involved in all subsequent tests, including preparation, data acquisition, and data analysis across five test beam campaigns. These tests were conducted during a critical period (2020–2021), marked by lockdowns due to the Covid pandemic, so the author had fully assumed this responsibility. The results presented in this

chapter are summarized in three documents, some of which are pending publication [68–71], to which the author contributed in various capacities.

The author played a key role in implementing the bent sensor geometry into the Corryvreckan framework, facilitating the analysis of data collected from test beams. To this end coordinate transformations from local sensor coordinates to global coordinates, the correct handling of roto-translation matrices and uncertainties, adjustments to track intercept parameters for sensors with an extent in the z -direction, tracking and alignment processes, and the validation of results were some of the changes needed to adapt the reconstruction framework to bent sensor geometries.

Among the publications, the first significant document [68] served as the initial proof of concept, utilizing an ALPIDE sensor bent along the short edge at a radius of 22 mm. The author actively participated in the test beam preparations and data acquisition, and contributed extensively to the data analysis and paper writing. The primary objective of this paper was to demonstrate the feasibility of using bent MAPS without compromising performance.

This work confirmed that bent sensors remain fully functional after bending. Electrically, bent sensors did not exhibit any additional non-functioning or noisy pixels. The threshold distribution measured before and after bending remained consistent. However, there were changes in current consumption, approximately $\pm 10\%$ and $\pm 5\%$ for chips bent along the short and long edges, respectively. These deviations, attributed to piezo-resistive effects, were largely influenced by the rotation of any two current mirror FETs with connected sources and gates within the architecture of the chips [74]. Despite these changes, they remained within the nominal operating margin of the ALPIDE sensors.

These sensors were also subjected to extensive testing during test beam campaigns. An ALPIDE sensor mounted on a carrier card was bent by sandwiching it between two Kapton foils and rolling it into position using a mechanical device. The test beam results demonstrated that the sensor maintained the same efficiency as in its flat state and could even achieve an increase in efficiency over a wide range of thresholds when particles passed through at an angle of approximately 36° , the maximum angle achievable with such a bent device.

Over four months, this device relaxed, resulting in a change in the radius of the sensor by approximately 8 mm. To address this issue, subsequent iterations employed 3D-printed cylindrical jigs designed to match the nominal radii of the ITS3 at the time of the tests. The ALPIDE sensors were either glued to the jig using bi-adhesive tape or secured in place by a Kapton sleeve. Both the jigs and the Kapton sleeve featured apertures that allowed particles from the beam to pass through the sensors without obstruction.

For the second publication [69], the author made similar significant contributions to the preparation and data acquisition, with an even greater focus on the analysis side, working alongside fellow colleague L. Lautner as the two primary authors. The objective of this publication was to explore ALPIDE sensors bent along the long edge, with an emphasis on achieving better control over the curvature of the sensor. The newly designed sensors were mounted on 3D-printed jigs and connected to the readout boards using flexible printed cables and adapter boards.

Several of these sensors were bent to the three radii planned for the ITS3 (18, 24, and 30 mm) and were subjected to testing during multiple test beam campaigns. These tests were conducted in conjunction with flat reference ALPIDE sensors that provided tracking.

As in the previous studies, the bent sensors demonstrated efficiency comparable to their flat counterparts. At nominal thresholds of around 100 electrons, an efficiency exceeding 99.9% was observed. Additionally, when the sensors were reverse biased with -3 V , a further noticeable increase in efficiency was recorded.

Results of the bent sensors with and without bias at the three different radii showed no changes with respect to the flat state, in terms of efficiency and spatial resolution. Moreover, a slightly better efficiency was observed due to particles impinging on the sensor at a non-zero incident angle.

Configurations in which the beam enters from the side of the bent sensors, allowing it to either pass twice through the sensor or simply graze it, were also studied. These scenarios are of particular interest for the ITS3, where, due to the presence of a magnetic field, low-momentum particles could curl and pass through the same sensor multiple times. By analyzing these grazing configurations, valuable information about the depth of the depletion region can be inferred.

Sensors bent along the long edge at different radii were arranged in a nested configuration, creating a mini-telescope of six bent sensors, referred to as the μITS3 . A copper target was placed at the center of the μITS3 , and the entire setup was integrated into a larger telescope. The primary goal was to assess the tracking performance and demonstrate the vertex reconstruction capabilities of such a setup in hadronic environments. The results of this testbeam campaign and subsequent analysis form the basis of the third paper [70], in which the author played a significant role, alongside colleagues from INFN Bari.

Aligning this twelve-plane telescope presented considerable challenges, but the author developed two alignment methods that produced consistent results. In a configuration analogous to the future ITS3, which will include three bent sensors and three (out of four)

flat ALPIDE sensors, a hadron beam was directed at the target and multiple outgoing tracks could be observed for each incoming track

In the initial phase, the target reconstruction yielded results that matched expectations across all three major axes. The shape and size of the target in the plane perpendicular to the beam direction were accurately reconstructed, and the target thickness measurements were in very good agreement with manual measurements and observations from an X-ray CT scan.

Subsequently, the analysis focused on the proximity between the incoming and outgoing tracks relative to the vertex of the interaction with the Cu target, with the impact parameter resolution being evaluated through the calculation of the distance of closest approach.

The data analysis revealed that for outgoing tracks passing through only the three downstream bent ALPIDEs, a pointing resolution of approximately $16\text{ }\mu\text{m}$ was achieved. For tracks at smaller angles relative to the beam, which passed through all six downstream sensors, a pointing resolution of about $4\text{ }\mu\text{m}$ was measured. These results align well with similar values obtained from Monte Carlo simulations conducted for the ITS3 Technical Design Report and are consistent with simulations performed for strangeness tracking, which employed a similar configuration of the setup and focused on secondary vertex determination.

While some aspects of the results remain under discussion and further investigations are ongoing pending the publication of the papers, these findings provide confidence in the alignment and tracking methods used, as well as in the ability of the setup to efficiently track particles and reconstruct vertices.

This series of measurements served to prove that CMOS MAPS subjected to bending forces do not experience a significant degradation in performance compared to their nominal flat state. Although some effects of curvature are observed (i.e. increased current consumption), these remain well within the nominal operating parameters of the ALPIDE sensors.

Nonetheless, it is important to note that these tests were conducted on ALPIDE sensors, which are produced using the 180 nm CMOS imaging node, a sensor architecture different from the one proposed for the ITS3. The ITS3 will feature sensors manufactured using the 65 nm technology. This smaller node allows access to larger wafers and offers benefits such as lower current consumption and potentially improved performance.

Currently, the collaboration is investigating the effects of bending on 65 nm analogue and digital chips. Preliminary results indicate that, similar to the ALPIDE sensors, bending does not have a significant impact on the efficient operation of these new sensors.

The ITS3 is more susceptible to sensor deformation than to positioning errors because of the thin and large sensors. X-ray inspections of the detector prototypes have revealed

deviations from the ideal cylindrical geometry and testbeam data showed that sensors can detach and the position resolution is impacted. Common deformations include: a curvature radius different from the nominal one, deviations from the cylindrical shape at the edges of the sensors, and compression causing elliptical distortion of the cylinder. Simulations have studied the impact of each deformation. Alignment procedures have been shown to restore the original performance [55].

These findings represent an important milestone, reinforcing the potential of bent MAPS as a robust and reliable technology option for detectors, especially in application where proximity to the interaction point and reduced material budget are prioritized. The successful integration and operation of the ITS3 in the ALICE experiment in the next years has the possibility to give rise to a paradigm shift in low-mass vertex detector design.

Material budget imaging

The upcoming Run 4 upgrade of the ITS of ALICE will feature the ITS3, a device which, once built, can be considered the most advanced ultra-light particle detector in HEP. This CMOS pixel camera, operational at the heart of the ALICE experiment and featuring many innovative concepts in HEP, will lack services in the active area. Like this it will reach unprecedented values of material budget, approaching a nearly-massless design.

Six wafer-scale sensors will be thinned down to about $50\,\mu\text{m}$ and concentrically bent around the beampipe. They will be held in place by open-cell carbon foam support structures positioned along the midsection (long side) of the cylinder and by carbon foam rings at both endcaps. On one of the endcaps, the foam will also serve as a heat exchange radiator. These will be the only support structures present in the active area of the detector, providing stability and ensuring the sensors maintain their targeted radii. Cooling will be facilitated through forced airflow from one of the endcaps. Altogether, the total material thickness for the first layer for tracks with $|\eta| < 1$ will be $0.09\% X_0$, of which the sensor itself is responsible for $0.07\% X_0$.

Understanding the little amount of scattering introduced by the carbon foam supports and silicon sensors is key to enhancing spatial resolution and ensuring high-efficiency track reconstruction for low-momentum particles. Dedicated measurements were performed using two foam samples mounted on ALPIDE sensors provided by CERN colleagues. The author was involved in all consequent steps of measurement and analysis.

During a dedicated testbeam campaign using electrons in the GeV-energy range, the author collected data and analyzed the two samples using a technique known as material budget imaging. This method involves positioning the samples in the middle of a tracking telescope, the former effectively acting as devices under test. Track reconstruction is performed separately upstream and downstream of the samples, with precise matching of the two tracklets at the position of the sample. The projected track angles are then used as a measure of the scattering that took place in the sample and are later related to the material budget contribution.

Corrections for telescope-induced effects are applied to the scattering distributions, and the results are then compared with theoretical predictions. The analysis revealed a good agreement between the expected and the observed material budget, confirming the validity of the measurement technique, as well as the small expected multiple scattering. This outcome validates the usage of low density carbon foam within the active area of the ITS3, ensuring no significant degradation in single point resolution.

3.1 The road to nearly massless detectors

3.1.1 Material budget changes from ITS2 to ITS3

Each layer of the ITS2 is segmented in the azimuthal direction into elements called staves. Extending over the entire length of their respective layers, these staves are the fundamental building blocks of the detector. They provide power distribution and routing of data lines, house cooling pipes and ensure the mechanical robustness required for the safe operation of the chips.

The ALPIDE sensors are mounted on these stave structures (see Fig. 3.1), each comprising three components: a hybrid integrated circuit (HIC), a coldplate, and a space frame.




Figure 3.1 – Exploded layout showing elements of the ITS2 inner (IB) and outer (OB) barrel staves. Taken from [39, 40].

The HIC includes the sensors themselves and a flexible printed circuit (FPC) used for the distribution of data and power. These HICs are glued to a coldplate, which is essentially a carbon fibre laminate equipped with cooling pipes. This arrangement ensures efficient thermal contact with the pixel chips, facilitating the dissipation of heat generated during their operation. The entire setup is mechanically supported by a space frame, a lightweight support structure made from filament-wound carbon fibre. The design of the IB and OB staves in ITS2 is largely similar, though there are some differences for the OB on account of the greater stave length. Specifically, the OB staves for layers 4 and 5 measure 844 mm, and for layers 6 and 7 they extend to 1478 mm, as opposed to the shorter 271 mm length of the IB staves [39].

This design is remarkably lightweight, achieving an average material budget of just $0.36\% X_0$ for the IB layers and $1.1\% X_0$ for the OB staves, where the material budget constraints are more relaxed (cf. Table 1.2). The azimuthal distribution of the material budget in the first layer of ITS2 is depicted in Fig. 3.2, providing a visual representation of the different components.

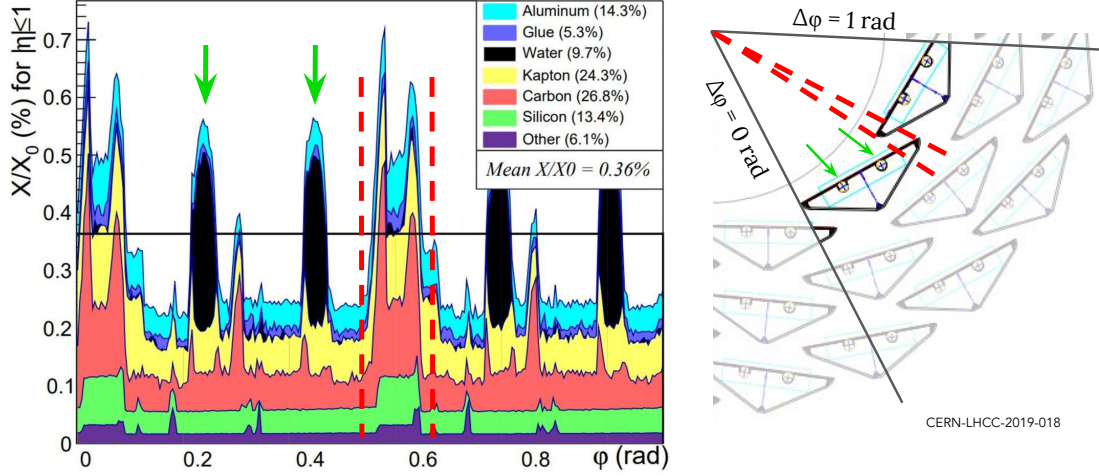


Figure 3.2 – (left) Azimuthal cumulative distribution of individual contributions to the material budget of an IB stave, expressed in units of radiation length. The relative contribution of each component is highlighted, with the average value represented by a black horizontal line. Notable peaks in the distribution are attributed to the overlap of staves (essential for ensuring the hermiticity of the detector), and to the presence of water cooling pipes.

(right) Schematic quarter-sectional view of the IB staves, emphasizing two staves: their overlapping area (delimited by the two red dashed lines) and the cooling pipes (indicated by the green arrows), which are also featured in the material budget distribution on the left. Figures adapted from [39, 40].

In a few years, the ITS2 IB will be replaced with a new detector design featuring wafer-scale silicon sensors, bent into truly cylindrical layers. The new ITS3 will approach a nearly-massless detector limit by reducing the material budget of the ITS2 IB layers by 80%. The design of the ITS3 half-layers focuses on minimizing any contribution to the material budget in the active area of the sensor, while ensuring the positional stability of the sensors over time.

Considering the ITS2 material budget composition, as illustrated in Fig. 3.2, it can be seen that silicon accounts for one seventh of the total material. Achieving the material budget constraints of the ITS3 ($0.07\% X_0$ per layer) necessitates a significant detector redesign. This involves either eliminating most components other than the silicon sensor from the active area or substantially reengineering them. Moreover, it dictates limits on the power density on the sensor, affecting in turn the mechanical and cooling system.

In the ITS2, the FPC and/or the power bus were required in order to arrange the sensors on the long staves, and were essential for distributing the supply, ground and bias voltages, along with data and control signals to the pixel sensors. The sensors have close to the maximum size allowed by the photoretic in their respective technology node ($< \mathcal{O}(20 \times 20 \text{ mm}^2)$) so, to span the length of a stave, several sensors needed to be tiled. However, with the exploration of stitching, it is possible to fabricate a single, large sensor that covers the length of a stave and is wider than a single sensor. The data and power transmission lines are to be embedded onto the silicon die itself, eliminating the need for dedicated FPC cables that would introduce extra material in the acceptance of the detector.

The wafer-scale sensors will have the periphery located at one endcap (A-side) where services are possible in ALICE (see Fig. 1.21). Here, most of the power dissipation circuits will be concentrated, resulting in significant localized heat generation. The power density in this region is expected to be about 800 mW/cm^2 [55]. In contrast, the active area is expected to have a power dissipation per unit area of 28 mW/cm^2 . Both these values are based on nominal operating condition with 1.2 V supply voltage at 25°C ambient temperature.

3.1.2 Carbon-foam design

In the design of the ITS3 cooling system, open-cell carbon foams are considered for their dual functionality. They not only play a major role in the cooling process, but ensure the mechanical stability and help maintain the curvature of the sensors. The structure of the foam, with large cells and linear ligaments, allows for effective air flow, aiding in heat dissipation.

The cooling design incorporates thermally-conductive open-cell carbon foam half rings (H-rings), strategically placed at the endcaps and in direct contact with the periphery of the sensors where the heat flux is maximal (see Fig. 3.3). These half rings are designed with a series of holes and slots to regulate airflow and function as heat exchangers. The heat from the sensors is conducted to the foam and subsequently removed from the active detector volume by forced air convection.

To maintain the cylindrical shape of the silicon sensors, the carbon foam rings in the endcap region alone are not sufficient, as the sensors would buckle. Therefore the design includes longerons extending along the beam direction, serving as mechanical supports. Two are positioned along the long edges of each sensor, and provide crucial support in the z direction. They are made from a different foam type, with lower density.

The choice of carbon foam materials is tailored to their specific roles. A high thermal conductivity foam is chosen for the high-power dissipation region at the periphery of the chips on the A-side, the same side that allows servicing in ALICE, due to the high surface-to-volume




Figure 3.3 – (top) Exploded and (bottom) assembled schematic view of the half-layer 0 of the ITS3. Taken from [55]. The A-side contains the cooling circuits, power supply and data readout units, while the C-side serves only as a secondary voltage supply for the long sensors (cf. Fig. 1.21)

ratio of the foam. The heat will be dissipated through direct forced convection, as air flows through the holes and slots of the half-ring. For the C-side half-rings and the two longerons, a lower density carbon foam with slightly inferior thermal properties is employed, since this ensures a low enough material budget.

For the ITS3, one critical requirement is the short-term sensor stability of $\leq 2 \mu\text{m}$. This necessitates a controlled, low-speed airflow to keep induced vibrations within acceptable limits. Thermal validation tests have demonstrated that an average airflow velocity of around 8 m/s between layers is sufficient to dissipate a heat flux of $40 \text{ mW}/\text{cm}^2$. This will ensure an effective detector operation while maintaining short-term position stability below $1.3 \mu\text{m}$ [55], therefore satisfying the ITS3 requirements.

A variety of foams meeting the various specifications in terms of density, thermal conductivity, mechanical stability, machinability, particle emission, and long-term behavior are available. Two types of open-cell Reticulated Vitreous Carbon (RVC) foams have been selected for their distinct mechanical and thermal characteristics:

- Carbon (RVC) Duocel[®], manufactured by ERG Aerospace, chosen for the C-side ring and the longerons due to its low density ($\rho = 45\text{--}50 \text{ kg}/\text{m}^3$)
- Allcomp K9 standard density (SD), manufactured by Lockheed Martin, with a higher density ($\rho = 200\text{--}260 \text{ kg}/\text{m}^3$), but good thermal conductivity of about $25 \text{ W}/(\text{m} \cdot \text{K})$ and higher specific surface area for better heat transfer due to added graphite

Given the high porosity (see Fig. 3.4) and low densities of these foams, they are expected to be almost transparent to particles. However, the influence on the material budget was thoroughly investigated in a specialized analysis campaign, which makes the purpose of this work.




Figure 3.4 – Optical microscope pictures of the two foams chosen for ITS3. Note the different scales between the two pictures. Taken from [55].

Two engineering models (EM) were built to address manufacturing and handling challenges and refine the sensor installation techniques for the final design. Such prototypes aim to simulate the final design as close as possible. The first model, EM1 (shown in Fig. 3.5), is equipped with dummy silicon wafers, cut to match the dimensions of the final sensors for ITS3 and held in place using carbon foam wedges, a design that was initially explored before the adoption of the longerons.




Figure 3.5 – Engineering Model 1 (EM1)

This specific design was ultimately rejected. A laser profile analysis of this prototype revealed slight deviations of the half cylinders from the targeted radii [91]. The outward

pressure on the foam wedges from the silicon trying to relax results in small bends (see Fig. 3.6a) at the edge of the chip and causes local footprints (see Fig. 3.6b) at the silicon-foam interface. Moreover, as detailed in [55, 92], due to capillarity effects on the very porous foam, glue seeps into the foam by some unknown amount. These factors combined posed a risk of compromising measurement accuracy by introducing uncontrolled deformations and material inconsistencies that could affect the precision of spatial resolution and track reconstruction.

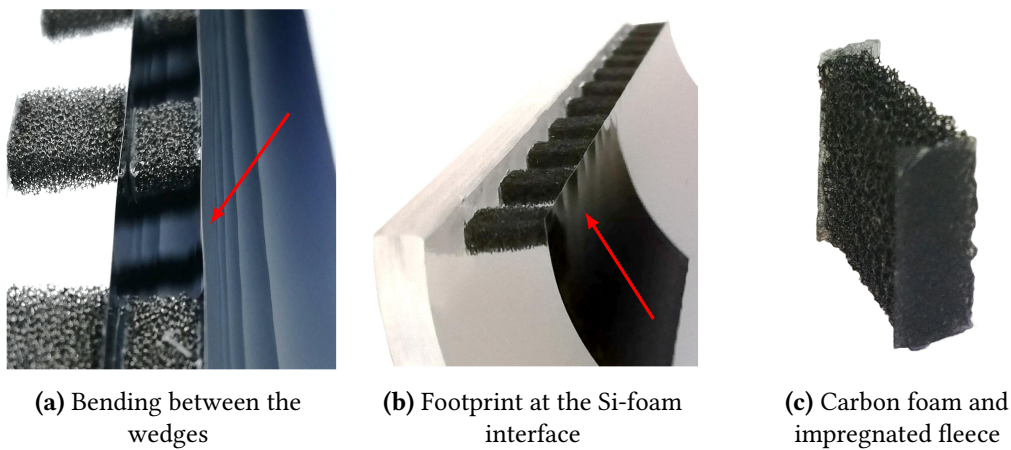


Figure 3.6 – First engineering model for the ITS3 containing a mixture of final grade components.

To mitigate these effects, it was necessary to maximize the contact surface. This led to the design of EM2 (Fig. 3.7), which replaces the wedges with a continuous foam rib called a longeron. Additionally, the footprint was minimized by adding a glue-impregnated carbon fleece, having a thickness of $120\ \mu\text{m}$ and an areal density of $8\ \text{g/m}^2$, at the interface between the foam and sensor (see Fig. 3.6c). The fleece offers a more uniform contact area and provides a better control of the glue layer that penetrates into the foam.

In addition to the carbon fleece at the foam-sensor interface, the EM2 features All-comp K9 SD half-rings and Carbon (RVC) Duocel® foam longerons for improved support and precise radii definition.



Figure 3.7 – Engineering Model 2 (EM2)

Various glues were considered for the ITS3 sensor-foam interfaces, with high thermal conductivity and radiation resistance [55]. ARALDITE® 2011 was selected for the longerons, while Epoxies® 50-3150 FR was chosen for the endcaps.

At the time this analysis started (Q2 2021) and data was first taken, EM1 was nearing completion, while the design phase for the second model had just begun. EM2 aimed to address and improve upon the weaknesses identified in the research and development phase of EM1. It is important to note that the original design for ITS3, as proposed in the Expression of Interest [52] and Letter of Intent [53], planned for sensor radii of 18, 24 and 30 mm. These dimensions slightly differ from those later targeted for the final ITS3, as detailed in the Technical Design Report [55].

The two chosen foam types, Carbon (RVC) Duocel® and Allcomp K9 SD – throughout the remainder of this chapter denoted as *ERG* and *ALLCOMP*, respectively – were subject to radiation length testing. Assemblies of foam, silicon sensors, glue, and fleece were produced and tested in a particle telescope to evaluate the amount of scattering they introduced.

3.2 Measurement idea

3.2.1 Carbon foam samples

For the in-beam characterization, two distinct samples were prepared, each incorporating a different type of foam: one contained the *ERG* foam and the other the *ALLCOMP* foam. Each sample was designed as a composite structure, comprising of two ALPIDE sensors (each measuring about $30 \times 15 \text{ mm}^2$) with a foam wedge ($10 \times 3 \text{ mm}^2$) sandwiched centrally between them, thus mimicking the structure between two ITS3 layers. Of these sensors, one is operational and bonded to a carrier card, whereas the second serves as a non-functional dummy.

These samples, provided by collaborators at CERN from the ITS3 mechanics and cooling team, are produced following methods similar to those used for building the EM1. Nevertheless, the samples required different tools than the ones used in the mechanical engineering models of the ITS3, and implied a more manual approach. Due to the unavailability of comprehensive details on the precise manufacturing processes and materials used in the samples, as well as the lack of optical measurement data, this material budget analysis proceeds on the basis of several assumptions, which are explained in the following paragraphs.




Figure 3.8 – Sample components at the sensor interface

An educated guess of the sample composition, broadly confirmed by the CERN team [93], is illustrated in a Fig. 3.8 and detailed below:

- A 50 μm thin functioning ALPIDE sensor, bonded to a carrier card
- On the substrate side of the active ALPIDE sensor, a $(75 \pm 25) \mu\text{m}$ layer of ARALDITE® 2011 epoxy is applied
- A carbon fleece (120 ± 10) μm thick, impregnated with the same type of resin, is attached next. It ensures better contact

with the carbon foam and controls the glue that can seep into the porous foam due to capillarity forces, as previously discussed

- The carbon foam wedge is then placed in contact with the glue-soaked fleece. A region of foam, $(150 \pm 25) \mu\text{m}$ thick, where glue partially seeps in develops. Previous measurements [91, 92] using comparable samples showed a similar glue penetration depth.
- The 6 mm thick carbon foam wedge with an area of $10 \text{ mm} \times 3 \text{ mm}$

The assembly procedure is symmetrically replicated at the opposite end of the foam, where the dummy ALPIDE sensor is then affixed. Both the ERG and ALLCOMP foams were subject to the same assembly process, resulting in two distinct samples for evaluation within the telescope setup.

Given the lack of specific data on the radiation length and other characteristics of the foam composites, estimates were made based on available manufacturer data and internal ITS3 reports. These are conservative estimates, to account for potential uncertainties in the design. These calculations are necessary in order to estimate the total radiation length of each sample, needed for theoretically predicting the width of the angular distribution according to the Highland formula.

ALPIDE sensor

The ALPIDE sensor is assumed to be a 50 μm thick homogeneous block of Si. In reality, the first few μm of the top part of the sensor contain aluminum metal lines. Aluminum has a radiation length similar to that of Si, so the approximation is justified.

The radiation length of Si is 9.37 cm [94]. The thickness of each sensor is known to a precision of below 1 μm . In total, for each ALPIDE sensor, the contribution to the material budget is about $x/X_0 = (0.053 \pm 0.001) \%$.

Glue layer

For gluing the two samples, ARALDITE® 2011 was used, which is a two-component (hardener + resin) epoxy. The combined density of the hardener and resin mixture is approximately 1.05 g/cm³, as specified by the manufacturer [95].

To estimate the radiation length of the ARALDITE® 2011 adhesive, it is presumed to be a phenol epoxy with the chemical composition C₆H₆O. Using the chemical formula, the relative weights for each element are calculated in Table 3.1.

Table 3.1 – Estimation of radiation length for ARALDITE® 2011 glue from basic principles.

Element		Atomic weight (g/mol)	Rel weight (%)	X_0 (g/cm ²)
C	6	12.01	76.57	42.7
H	6	1.008	6.43	63.05
O	1	16	17	34.24
Total				41.81

The table outlines the volume percentage, relative weight percentage, and radiation length, calculated based on their presence in the chemical structure of the glue.

$$\frac{1}{X_0} = \sum_i \frac{m_i}{M} \frac{1}{X_i} \quad (3.1)$$

Eq. 3.1 shows how to combine the radiation length of a material from its constituents. m_i and X_i represent the molar mass (contribution of the atoms in the molecule) and the radiation length for each individual atom type. M represents the molar mass of the compound and X_0 denotes its combined radiation length. The combined radiation length for the ARALDITE® 2011 adhesive is calculated to be $X_{0,\text{glue}} = 41.81 \text{ g/cm}^2$, equivalent to about 39.82 cm. This estimation is in close agreement with similar values reported in other studies of groups from the LHC experiments [96]. Per sensor interface, a glue layer thickness of $(75 \pm 25) \mu\text{m}$ is assumed.

Given that the assembly includes a glue layer at the interfaces with both the operational and dummy ALPIDE sensors, the total glue thickness within the sample amounts to $(150 \pm 35) \mu\text{m}$, corresponding to a total material budget contribution per sample of $x/X_0 = (0.038 \pm 0.009) \%$.

Nonetheless, the exact thickness of the glue layers carries some degree of uncertainty due to the manual nature of the sample assembly process.

Carbon foam

Both foams are morphological structures of carbon. As such, the radiation length of the carbon foam is derived by scaling the known radiation length of the carbon (graphite) according to the density of the foam (the only information available from the manufacturers):

$$X_{0,\text{foam}} = X_{0,\text{C}} \frac{\rho_{\text{C}}}{\rho_{\text{foam}}} \longrightarrow \left(\frac{x}{X_0} \right)_{\text{foam}} = \frac{t \cdot \rho_{\text{foam}}}{X_{0,\text{C}}} \quad (3.2)$$

For the foam samples a thickness of $t = (5.7 \pm 0.1)$ mm is considered for the region where the foam alone is present. The remaining foam volume that makes up the 6 mm sample is assumed to consist of a mixture of foam and glue that has seeped into the foam, a detail that will be elaborated on in subsequent sections. Air, due to its very large radiation length (303.9 m [94]), is neglected from calculations. Table 3.2 contains the relevant values derived from this assumption, alongside the measured density ranges and the resulting total radiation length of the two foams within the samples.

Table 3.2 – Quantities pertaining to the material budget of the two tested carbon foams.

	ERG	ALLCOMP
Thickness x (mm)	5.7 ± 0.1	5.7 ± 0.1
Density (g/cm ³)	0.047 ± 0.001	0.23 ± 0.03
X_0 (cm)	909 ± 19	185 ± 24
x/X_0 (%)	0.063 ± 0.002	0.307 ± 0.041

The density of the two foams was measured [93] and the ERG foam showed a smaller variability than the ALLCOMP foam. As expected from the very low density, the ERG foam has a very large radiation length and is perfectly suited for use in the active area of the ITS3.

Glue-impregnated carbon fleece

For the carbon fleece, the only known parameters are the surface density of 8 g/m², the total thickness of 120 μm and the fiber volume fraction (50% vf) [97]. Given that the fleece is impregnated with glue, a 50% volume fraction (foam-glue) is assumed.

Considering that the carbon fleece is a carbon-derived compound, the previously defined radiation length of C and the surface density of the fleece can be used. By applying the same scaling as in Eq. 3.2, a radiation length of 0.002% is calculated for the carbon component alone.

The thickness of the whole impregnated fleece is assumed to be $(120 \pm 10) \mu\text{m}$. Considering half of it glue (50 % vf), yields a total material budget contribution due to the fleece of $(0.017 \pm 0.001) \%$ calculated as:

$$\left(\frac{x}{X_0} \right)_{\text{impregnated fleece}} = \frac{0.5 \cdot (120 \pm 10) \mu\text{m}}{39.82 \text{ cm}} + \frac{8 \text{ g/m}^2}{42.7 \text{ g/cm}^2} = (0.017 \pm 0.001) \%$$

With an impregnated fleece layer present at both ends of the sample, the total material budget contribution per sample amounts to $x/X_0 = (0.034 \pm 0.002) \%$.

Foam-fleece interface

The foam-fleece interface is prone to uncertainties with respect to the depth of the glue seepage, which is sample dependent. A conservative estimate of $(150 \pm 25) \mu\text{m}$ glue penetration per interface for both foams is assumed, considering similar values measured in [92].

The porosity of both the ERG and ALLCOMP foam samples was measured by means of micro-CT scans [98]. For the ERG, a high porosity of 97% was measured, indicating that only 3% of the volume is occupied by the carbon foam itself. Conversely, the ALLCOMP showed a more sample-dependent porosity, with an average measured at 88%.

The radiation length contributions of the glue-penetrated part of the foams are calculated by considering various scenarios of glue penetration.

In the absence of any glue penetration, the radiation length would correspond to that of the respective foam. For the ERG sample, this would be $(0.003 \pm 0.001) \%$ and for ALLCOMP $(0.016 \pm 0.003) \%$.

The second scenario assumes the complete penetration of the void fraction. Considering the glue density mentioned previously, the thickness of the glue penetration layer $\delta x = (300 \pm 35) \mu\text{m}$ (accounting for penetration on both sides of the foam), and the void fraction for both samples, the material budget contribution is calculated to be $(0.073 \pm 0.009) \%$ for the ERG sample and $(0.068 \pm 0.008) \%$ for the ALLCOMP sample.

The third and most probable scenario assumes partial penetration, with 50% of the void fraction occupied by glue, as not all foam pores are expected to be filled. Following a similar calculation as above, while considering this reduced void fraction, the material budget contribution due to glue penetration is estimated to be $(0.037 \pm 0.006) \%$ for the ERG sample and $(0.035 \pm 0.006) \%$ for the ALLCOMP sample. This last scenario is considered in further calculations.

Total contribution

The contributions of the various components to the total material budget of the carbon foam samples are summarized in Table 3.3, highlighting the impact of each component based on conservative assumptions for the two samples under test in this work, values that may vary with the final sensor design.


Table 3.3 – Contribution of individual components to the material budget of the carbon foam samples. The variation in radiation length is attributed to the range of densities, measured for each foam: ERG (0.046–0.048 g/cm³) and ALLCOMP (0.2–0.26 g/cm³). The thickness refers to the total amount of material in the sample.

Sample	Material	Thickness (μm)	$x/X_0(\%)$
ERG	ALPIDE sensor	2×50	0.107 ± 0.002
	ARALDITE® 2011 glue layer	$2 \times (75 \pm 25)$	0.038 ± 0.012
	Glue-impregnated carbon fleece	$2 \times (120 \pm 10)$	0.034 ± 0.002
	Glue-infused foam interface	$2 \times (150 \pm 25)$	0.037 ± 0.006
	Carbon foam	5700 ± 100	0.063 ± 0.002
			$\Sigma = 0.279 \pm 0.011$
ALLCOMP	ALPIDE sensor	2×50	0.107 ± 0.002
	ARALDITE® 2011 glue layer	$2 \times (75 \pm 25)$	0.038 ± 0.012
	Glue-impregnated carbon fleece	$2 \times (120 \pm 10)$	0.034 ± 0.002
	Glue-infused foam interface	$2 \times (150 \pm 25)$	0.035 ± 0.006
	Carbon foam	5700 ± 100	0.307 ± 0.041
			$\Sigma = 0.521 \pm 0.043$

The total material budget of the ERG sample amounts to $(0.279 \pm 0.011) \%$, while for the ALLCOMP sample it is $(0.521 \pm 0.043) \%$. The larger uncertainty on the ALLCOMP sample is due to the wider range of measured densities. However, as it will be demonstrated in subsequent sections, the material budget analysis focuses solely on the contributions from the glue, carbon fleece, and foam, excluding the two ALPIDE sensors. Therefore, the total material budget quantities for the theoretical expectation omit the contribution of the sensors and are $(0.172 \pm 0.011) \%$ for the ERG sample and $(0.414 \pm 0.042) \%$ for the ALLCOMP sample.

Two pie charts derived from this table and shown below illustrate that the most significant material budget contribution in the ERG foam sample, which is to be used in the active area of the ITS3 detector, is the sensor itself. The foam represents only 23% of the total material budget, as expected from its low density. The various glue contributions amounts to about a quarter of the ERG sample. Minimizing the amount of glue that seeps into the foam as a

result of capillary forces is paramount to ensuring the strict material budget requirements of the ITS3, as well as minimal systematic uncertainties.



In the case of the ALLCOMP sample, the material budget contribution of the foam itself is dominant. However, this is less critical since this foam is placed within the ITS3 at the periphery of the sensors, outside of the active area of both the ITS3 and the ITS2 OB.

The different thicknesses and contributions shown in Table 3.3 are conservative estimates and subject to modifications as the final design of the sensor evolves. For instance, the contribution of the sensor to the material budget is anticipated already to increase, considering the need for extra thick metal lines needed to compensate for the voltage drops (as explained in Sec. 3.1). The final ITS3 sensors are expected to incorporate extra metal lines (of which one extra thick) and additional copper routing layers, equating to the radiation length of a 16 μm thick silicon layer [55, 99]. Consequently, after thinning, the 50 μm thick chips will effectively have the equivalent radiation length of a 66 μm thick layer of pure silicon.

Moreover, the glueing interfaces for the final sensor assembly are projected to be more precisely controlled than the manual production of the samples. In order to achieve a more controlled production environment, specialized tools were developed in house for sensor handling, bending, and assembly.

The estimates on the thickness of the individual layers are subject to uncertainties, as previously discussed. While some quantities are known with greater precision than others, the largest uncertainties arise from the amount of glue present. The previous calculation assumed a uniform layer of glue $(150 \pm 35) \mu\text{m}$ thick present at both interfaces with the sensors, as well as $(300 \pm 35) \mu\text{m}$ of glue seepage into the foam. Although these are reasonable estimates (as shown earlier), the variability observed between samples and even across the surface of a single sample can be significant. Different regions can have a different amount of glue that seeps in, or the seepage is not uniform across the surface locally. This was observed already by X-ray computed tomography images of the EM2 (see for example Fig. 4.40 in [55]) and emphasizes the need to account for these uncertainties in the final material budget estimate.

In the following chapters, calculations will be based on the material budget from the two samples, excluding the contribution of the ALPIDE sensors. The results presented in Table 3.4, which include the previously discussed material composition uncertainties, represent the final theoretical expectation for the two foam samples.

Table 3.4 – Theoretical expectation of the material budget contributions of the two carbon foam samples measured in the testbeam campaign (modulo the contribution of the ALPIDE sensors).

Composite sandwich	Theoretical expectation x/X_0 (%)
ERG	0.172 ± 0.011
ALLCOMP	0.414 ± 0.042

3.2.2 Procedure

Ideally, a detector should be ultra-lightweight, minimizing the material in the path of particles coming from the interaction point. This is crucial because as particles traverse matter, they experience multiple Coulomb scattering with atomic nuclei, which can deflect them from their original trajectory. Such scattering primarily results in a change in direction, with positional offsets becoming negligible for thin layers like those in silicon sensors¹. This angular deflection leads to positional inaccuracies that can deteriorate the spatial resolution of the tracking detector, affecting track reconstruction accuracy.

Since this deterioration is measurable, it can be quantified and related to the properties of the material that was traversed. Consequently, the deflection angle from the initial trajectory needs to be measured in order to do so. For a particle, this would mean reconstructing the trajectory before entering the scattering medium and after it and measuring the deflection or kink angle between the two. Considering a beam of particles, the individual measured kink angles will lead to a distribution, the shape of which will be influenced by the amount and properties of the material that is traversed.

As detailed in Sec. 1.1, a charged particle that traverses a layer of material will be subject to multiple Coulomb scattering, getting deflected multiple times in the presence of the electric fields of the nuclei. The sum of those deflections will lead to a statistical distribution of the scattering angle, symmetrically distributed relative to the incident direction of motion. The end result is a net effective deflection angle θ , as described by the Molière theory [9].

¹For example, for the densest material studied (Ni) at a thickness approximately three times that of the thickest material under consideration (500 μm) and at the lowest momenta (1 GeV), the scattering angle is on the order of 2 mrad. The resulting offset, calculated as the product of the scattering angle and the thickness [2], is about 1 μm . For silicon sensors, this would be in the order of a few nm.

The core of the distribution, which in general represents the cases with a large number of scatterings, can be approximated by a Gaussian, as a result of the central limit theorem.

The distribution, shown exemplarily in Fig. 1.9 exhibits the Gaussian core resulting from the many Coulomb scatterings, and large non-Gaussian tails predominantly coming from the non-negligible large multiple scattering angles from the underlying Rutherford scattering processes. In a first order approximation, the total distribution can be specified by the standard deviation θ_0 of the angle projected onto the plane perpendicular to the direction of initial motion.

The transition from multiple to single scattering occurs at roughly $\theta_0 \approx 2.5$ when the Gaussian has below a percent of its peak amplitude. Therefore, it is safe to consider the approximation given by Highland and detailed in Sec. 1.1.

In this approximation, the inner 98% of the distribution is considered and fitted with a simple Gaussian, the width of which will then be θ_0 . The other 2% of the distribution, representing the tails are neglected, where rare, large scattering angles occur. This Highland parametrization, considering only the effect of multiple scattering, can be used to infer properties of the materials, and is accurate within uncertainties up to 11% or better for thicknesses between 10^{-3} and 100 radiation lengths [94]. This 11% deviation is with respect to a fit to the Molière scattering distribution for all Z, for the 98% inner core [17].

Understanding the amount of scattering present in the future ITS3 detector is crucial since this dictates the deterioration of the resolution of the tracking reconstruction. The magnitude of the multiple scattering depends on the material budget, measured in units of the radiation length (X_0).


A technique that makes use of the scattering angle to infer the radiation length has been introduced in the past years by several research groups [100, 101] and it is referred to as material budget imaging (MBI). Alongside the previously mentioned studies, the measurements in this work follow closely the procedures presented in [102–104].

This imaging procedure consists in recording the passage of particles as they traverse a scatterer, by reconstructing tracklets both before and after it and then measuring the resulting deflection angles. The angles will form a distribution, the width of which depends on the amount of material the particles encounter along their trajectory.


For an accurate measurement of the scattering and track reconstruction, very thin detectors with good spatial resolution are needed. To this end, a telescope made of ALPIDE sensors (each 50 μm thick) is used for the MBI measurement of the carbon foam samples (see Figs. 3.9a and 3.9b). Four reference ALPIDE sensors before – the last of which is at the interface with

the scatterer – and three after the target are used to reconstruct upstream and downstream tracks, from which information about the scattering angle is extracted at the position of the scatterer.


With a detection efficiency exceeding 99.9% at the nominal threshold settings and a bias voltage of -3 V, each traversing particle is expected to produce seven hits across the ALPIDE telescope. A hit is reconstructed as the centroid (geometric mean) of a cluster of pixels, which forms due to the diffusion of charge from the ionization generated by incoming electrons from an accelerator that pass through the sensor.




(a) Schematic of the ALPIDE telescope with 7 active chips



(b) Picture of the setup with the foam sample in the middle



(c) ERG and ALLCOMP carbon foam samples



(d) The DUT foam sandwich

Figure 3.9 – The carbon foam DUT is placed in the middle of an ALPIDE telescope featuring seven active sensors. The sample is a sandwich of carbon foam and other structural materials, including a non-functioning sensor, closely mimicking an ITS3 layer. An example of a reconstructed upstream and downstream track and the scattering angle is shown.

The data has been collected with the help of the EUDAQ2 software [82], while the subsequent analysis was performed using the Corryvreckan testbeam data reconstruction framework [79].

Intermediate summary

The goal of the analysis is to provide a robust method for measuring the material budget of the two types of carbon foam support structures (ERG and ALLCOMP), thereby providing a direct measurement of the scattering introduced by these structures in the future ITS3 layers.

Traditional methods exist for measuring the material budget of a detector assembly during operation at an accelerator facility. For example, probing the detector material is possible via the conversion of photons into electron-positron pairs, offering a form of γ -ray tomography of the detector. This approach has been proven for the ALICE experiment to be precise down to a 2.5% material budget systematic uncertainty [105].

However, in scenarios where the amount is minimal, as is the case for the ITS3 detector, the conversion probability is significantly reduced, requiring many photons for an accurate measurement. Moreover, such methods are used in MC studies where the material budget is a direct input, allowing conclusions to be drawn about the implementation of the detector material in the simulation. It has been observed in some cases [106] that the description of the detectors in the simulation is incomplete and a better understanding is needed.

In contrast, the imaging method discussed in this chapter offers a direct measurement of the material budget of samples with unknown radiation length. This method is advantageous over photon conversion techniques, particularly for low mass detectors, as it is less dependent on large statistical samples. Additionally, it provides a direct input on the radiation length of the structures, which can be used to refine the simulation tools. Moreover, it will be shown that this material budget imaging technique can achieve a similar precision as the photon conversion method.

This imaging analysis builds upon previous knowledge of the method and improves the understanding of a number of systematic effects present in these measurements. Various estimators of the width of the scattering angle distribution, as well as different percentages of the inner core of the distribution, are studied, and their robustness and stability is assessed. While some distribution width estimators exhibited significant variability with different percentages of the distribution under study, most identified estimators remained stable. By combining the estimators of the width of the scattering angle distribution, a better precision can be achieved.

Measurements for the carbon foam samples were performed at the DESY testbeam facility in 2021, using electron beams at beamline 24. The author was solely involved in all preparations and measurement steps at the testbeam, leading to the final results presented here.

The energy of the beam was varied to study the impact of scattering at various momenta. Datasets at 1.0, 2.4 and 5.4 GeV/ c for both foam samples were collected. This approach was

particularly relevant due to the explicit dependence of the Highland formula on the momentum of the beam particles.

Nickel and aluminum targets with well-known thicknesses were used to calibrate the material budget estimates and compare them to model predictions from the Highland parameterization. This calibration corrects a large part of the residual telescope effects convoluted into the measurement and will be later discussed in Section 3.3.6. This calibration data was taken later (April 2024), with a telescope of a similar configuration and sensor placement, operated at the same beam line at DESY, but for a very limited time (six hours).

Then, corrections for various factors are discussed in depth and the material budget estimates are presented. By selecting an estimator that is robust and stable, or by combining different such estimators, a material budget estimate consistent with the theoretical expectations and having a small systematic uncertainty is achieved. Finally, the material budget contribution of the sandwich of components (glue, carbon fleece, foam) that mimic the material separating two sensor layers for the future ITS3 sensor is quoted.

3.2.3 Experimental setup

Each of the seven ALPIDE sensors from the experimental setup (see Fig. 3.9b) is capable of measuring the passage of electrons with an intrinsic resolution of about $5\text{ }\mu\text{m}$. The tracking resolution at the position of the DUT depends on the chosen momentum, but is better than $4\text{ }\mu\text{m}$ regardless of the settings, as verified by two independent GBL track resolution calculators [65, 66]. These tools take into account the intrinsic resolution of the sensors, the beam momentum, and the material budget present in the beam path. They demonstrate that the tracking resolution is a setup-specific quantity and that an a priori calculation can be made for a chosen telescope geometry.

The remarkable position resolution of the ALPIDE telescope enables a precise reconstruction of the particle trajectories. Nonetheless, to evaluate how the plane positioning within the telescope influences the angular and tracking resolution of the setup, the same GBL track resolution calculators are used. They account for the uncertainties arising from scattering effects within the telescope planes, surrounding air, as well as the foam samples (here for simplicity modeled as a uniform scatterer of fixed material budget, as big as the sensors).

Figure 3.10 shows how variations in the distance between ALPIDE planes impact both resolution metrics for particle trajectories. For simplicity, the plots show the changes for configurations in which the distance between consecutive ALPIDE planes is varied between 10 and 100 mm, while keeping the distance between the middlemost two planes (sensors 3 and 4, the ones closest to the foam sample) fixed at 20 mm (see Fig. 3.9a).

Although the positioning of the sensor planes within the telescope impacts both position and angular resolution, the geometry of the setup was fixed by the mechanical enclosure that was available during the testbeam (see Fig. 3.9b) and could not be altered for this study. At the time of the testbeam, three telescope setups were operated in parallel. For the setup used in this work, the decision to use a readily available compact telescope with fixed sensor positions was made.

The results from the simulation show that as the distance between planes increases, the tracking resolution of the telescope worsens, highlighting a preference for minimal spacing. Conversely, a longer lever arm between the planes enhances angular resolution.

In the configuration used for this work the spatial resolution is very close to its minimal value, for each of the three momenta under study (see the left panel of Fig. 3.10) and ensured an accurate measurement point both upstream and downstream, as close as possible to the foam sample. Constraints imposed by electronic components on the readout boards and carrier cards set the minimum feasible inter-plane spacing at 15 mm, so a conservative value of 20 mm between adjacent sensor planes was chosen for the compact telescope mechanics in the first place.

However, as shown in Fig. 3.10 (right), a small inter-plane distance results in a scattering angle resolution which is larger than the theoretical minimum. The larger angular resolution leads to an increase of the minimal scattering angle that can be measured.

More detailed studies on the angular resolution and the spatial resolution were done by varying all the inter-plane distances. They are shown in Appendix B. The results presented there show what the ideal setup should be, either for obtaining the best angular resolution or the best spatial resolution using the example of a six plane setup. Nevertheless, with the 20 mm inter-plane distance that the setup used in this work had, the angular resolution was found to be close to its optimal value across the various datasets.

Tracks are formed by requiring a hit in each of the seven ALPIDE sensors. Hits are associated to a track if they are within a circle of 100 μm radius from the track intercept per plane. This starting value of ~ 3 pixel pitches for associating hits to tracks is chosen in order to have a larger sample of tracks for software-based alignment, since the telescope is not yet aligned.

The effects of scattering within the ALPIDE planes, the foam sample and the surrounding air volume, as well as the measurement uncertainties stemming from the finite size of the pixels, are accounted for through the application of a General Broken Line (GBL) track fitting algorithm, as referenced in Section 1.5.2. The difference between a simple straight line fit and

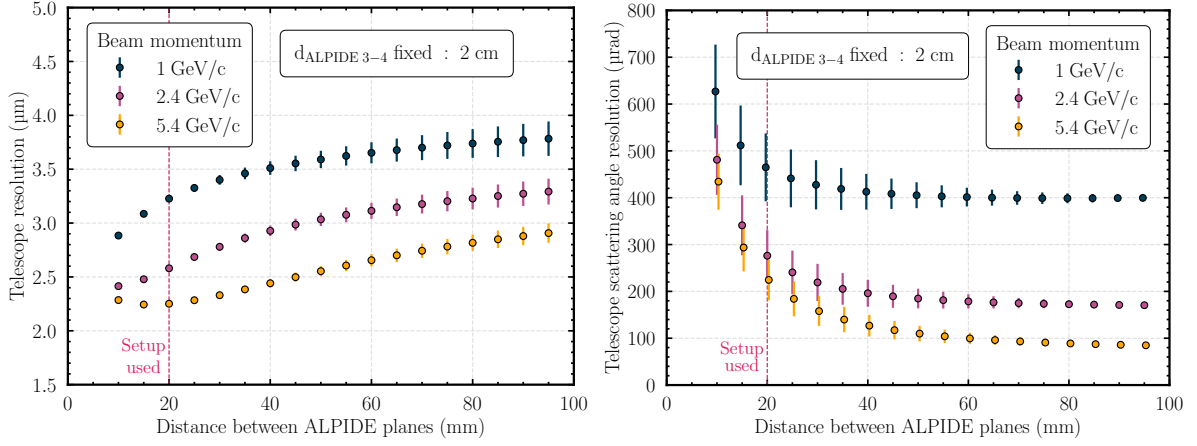


Figure 3.10 – (left) Spatial resolution and (right) angle resolution at the position of the foam sample calculated using [65, 66] as a function of the distance between ALPIDE planes, for the three momenta under study in this work. The errors account for the momentum uncertainty of about 150 MeV/c measured at the DESY II beam lines [86].

a GBL fit is illustrated in Figure 3.9a. By taking into account the effects of multiple scattering, GBL offers a better approximation of the trajectory of the particle.

The software alignment of the sensors based on the reconstructed GBL tracks is performed by means of the Millepede-II algorithm [107]. This method enables a simultaneous least-squares fit of both local parameters, which are specific to individual tracks and influence only the local track model fit, and global parameters, which include alignment constants and the degrees of freedom of the detectors.




Figure 3.11 – χ^2_{ndf} for tracks fitted with a straight line or with GBL, for the ALLCOMP foam sample

The figure of merit for the software alignment is the χ^2/ndf value. It is shown exemplarily in Fig. 3.11 for the case of a 1 GeV electron beam reconstructed with the telescope from Fig. 3.9. A straight-line fit has $(2 \cdot n_{\text{planes}} - 4)$ degrees of freedom. They are derived from the two local x and y measurements per sensor plane, with the z position determined by the position in space of the plane, and four constraints on a straight track in 3D space (two slopes and two offsets for the x and y directions). Although a GBL fit requires more parameters to define the tra-

jectory, it is constraint by the kink angles measured at each scatterer [108]. Ultimately, it results in the same number of degrees of freedom: $2 \cdot n_{\text{planes}} - 4$.

The analysis reveals a clear difference in tracking performance when accounting for the effects of multiple scattering, particularly at the lowest electron energy evaluated. The GBL algorithm demonstrates superior accuracy over conventional straight-line fitting methods.

For the middle ALPIDE sensor, which is directly attached to the carbon foam, the biased residuals are analyzed to assess the tracking precision. The usage of biased residuals is motivated by the fact that the detector is used in track fitting. The biased residuals contain an intrinsic resolution component, as well as a tracking resolution component:

$$\sigma_{\text{biased}}^2 = \sigma_{\text{intrinsic}}^2 + \sigma_{\text{track}}^2 \quad (3.3)$$

Here σ_{biased} represents the biased residual value in a spatial direction obtained for example from the fits to data, as shown in Fig. 3.12. These residuals are calculated by taking the distance between the track intercept on the sensor and the nearest associated hit for all the tracks in data. The intrinsic value of the ALPIDE sensor (typically about $5 \mu\text{m}$) is denoted by $\sigma_{\text{intrinsic}}$.




Figure 3.12 – Biased residuals in the x direction. Similar values are obtained for the y direction. A Gaussian fit is performed on the points.




Figure 3.13 – Track resolution at the position of ALPIDE 3 and at the position of the carbon foam. The lines are not a fit to the data, but are intended to aid in visual interpretation. The error bars account for the momentum divergence of about $150 \text{ MeV}/c$ measured at the DESY II beamlines.

Finally, σ_{track} is the track resolution at the position of the respective ALPIDE sensor. This tracking resolution is convoluted with the intrinsic resolution of the sensor into the measured residual distribution. In order to obtain the intrinsic resolution of the sensor, the track resolution has to be known and added quadratically from the measured biased residual.

The telescope tracking resolution (σ_{track}) is calculated using the information from the GBL track resolution calculators [65, 66]. Its magnitude is shown in Fig. 3.13. Here, the

tracking error is depicted as a function of the beam energy at the position of the active sensor attached to the foam sample (ALPIDE 3), as well as at the middle position of the foam itself (+3 mm downstream with respect to ALPIDE 3).

The measured biased residuals (see Fig. 3.12) align with the theoretical expectation outlined in Eq. 3.3, when taking into account the tracking error shown in Fig. 3.13. An intrinsic sensor resolution of $\sqrt{\sigma_{biased}^2 + \sigma_{track}^2} = \sqrt{3.5^2 + 3.7^2} = (5.1 \pm 0.1) \mu\text{m}$ is found, consistent with the known performance characteristics of the ALPIDE sensors.

With an aligned telescope, the data is separately analyzed in the upstream (first 4 ALPIDE) and downstream (last 3 ALPIDE) planes. This division allows for the formation and matching of tracklets at the position of the carbon foam, where an arbitrary kink is allowed. A maximum permitted distance between extrapolated upstream and downstream tracklets of $30 \mu\text{m}$ (one pixel pitch) is allowed, to account for scattering. The choice of this value is motivated by the fact that most (99%) of the tracks have a hit closer than $30 \mu\text{m}$ with respect to the track intercept on the ALPIDE plane (see Fig. 3.14; the red arrow indicates the point after which the number of entries is 1% of the total).




Figure 3.14 – Distance between track intercept and hit position after alignment.

Although GBL fitting typically aims to minimize or adjust for significant kink angles at scatterers in order to enhance track accuracy, this process could bias kink angle estimations, especially if the exact material budget (sample) is unknown. However, the fit is unconstrained at the position of the active ALPIDE which hosts the foam sample. This means that no constraint is applied to the track after it reaches the ALPIDE sensor, as this is the last measurement point.

Because no constraint is applied, the track is simply extended towards the foam after the ALPIDE sensor, meaning that the measurement takes place before the track experiences any scattering from the foam. In simpler terms, the track is free to follow its natural trajectory before it is deflected by the foam material, and the fit does not attempt to “predict” or compensate for any scattering beyond the sensor.

This approach allows the local derivatives (used in the GBL calculation) to remain free from constraints, allowing for an accurate measurement of the scattering angle, without influencing the χ^2 value of the track fit. In other words, the scattering that happens in the foam does not affect the trajectory data recorded at the sensor, preserving the accuracy of the measurement. Consequently, an unbiased estimation of electron scattering angles is achieved by comparing the slope differences of the upstream and downstream tracklets in both the x and y directions.

3.3 Image reconstruction

3.3.1 Kink angle reconstruction

The matching of the two tracklets allows to project their direction vectors and enables the derivation of two measurements of the deflection angle of the particles, in both the horizontal (k_x) and vertical (k_y) directions, on a virtual reconstruction plane, perpendicular to the beam direction. Given the azimuthal isotropy of the scattering phenomena and the orthogonality of the two measurement axes, these kink angle values are presumed to be uncorrelated.




Figure 3.15 – Correlation between the kink angles in the x and y directions from reconstructed tracks

This expectation is confirmed by the analysis of the correlation between k_x and k_y , as depicted in Fig. 3.15, where a negligible Pearson correlation coefficient substantiates the lack of correlation between the two measurements. Consequently, the independent assessment of the kink angles in the x and y planes enhances the robustness of the scattering angle width by offering two independent estimators.

The reconstruction plane is divided into image cells, thus offering a position-resolved scattering angle measurement. For each cell, scattering angle distributions are compiled from the trajectories of electrons that pass through, allowing for a precise calculation of the material budget in a specific area on the measurement plane.

Since the two separate angles are uncorrelated, they can be combined: the values of k_x and k_y are filled in the same histogram, denoted with k_{x+y} which is used further for analysis. A small Allpix² simulation was conducted to investigate the effects of the asymmetric pixel pitch

of the ALPIDE sensor on the kink angle distribution. The results are detailed in Appendix G. In summary, the slight asymmetry in the pixel pitch has a minor, yet negligible, impact on the kink angle distribution, slightly broadening it in the direction of the longer pitch.

To show this also for the data, a direct comparison between the kink angle distributions in the x and y directions is performed, taking the example of the ALLCOMP sample at a momentum of 1 GeV/ c in Fig. 3.16.




Figure 3.16 – The kink angle distribution in the x and y directions for the ALLCOMP sample at 1 GeV.

The two distributions are nearly indistinguishable, with only a small difference of about $10 \mu\text{rad}$ in width. This results in a systematic difference of $5 \mu\text{rad}$ when combining the k_{x+y} distributions, which is carried through the analysis.

The other distributions for the ALLCOMP sample, as well as for the ERG sample are investigated. With increasing momentum and lower material budget, this difference gets smaller, as the total width of the distribution gets narrower.

However, as will be shown later, other systematic uncertainties dominate the measurements. Therefore, introducing this small bias in the analysis is deemed acceptable in exchange for doubling the size of the available statistics sample.

After having shown that the x and y angle distributions can be combined, for the 2D reconstruction plane, the two kink angles are both filled into the bin corresponding to the matched position of the two tracklets at the reconstruction plane.

Optimizing the imaging resolution for assessing the material budget involves minimizing the size of the image cells in order to resolve finer details. Ultimately, the smallest image cell

size is inherently limited by the tracking resolution (and the corresponding spatial window for track joining), as illustrated in Fig. 3.13. Nonetheless, enough entries per cell are needed to extract a width.

In this thesis no fine details are expected or studied. The foams themselves have around 100 pores per inch which would translate into foam cells in the order of $250\ \mu\text{m}$ with filaments that can be down to $100\ \mu\text{m}$ thick. The sample is irradiated from a single, frontal angle and the information of all pores is therefore overlapped along the beam direction. The size of the image cells is chosen in order to have at least 500 entries for each bin in the two dimensional histograms for the cases where statistics are limited.

Typically, beam time is restricted to slots of a week, and sometimes the beam line must be shared with other users. Additionally, in that week various settings of the sensor are studied, as well as different momenta. Moreover, given the maximum achievable particle rate at the beam line, the available statistics is limited and will influence the number of entries per image cell. This limitation directly impacts the statistical robustness of the scattering angle distributions obtained for each cell. Large number of tracks are preferred to ensure a more reliable estimation of scattering angles and, by extension, a clearer imaging precision for the material budget analysis.

A dataset from a different beam test, used for training a group of students [109], which had higher statistics and slightly different sensors, was investigated to illustrate how varying cell sizes affect the resolution power for distinguishing different features. A custom made 3D-printed scatterer with fine features was produced and placed in the middle of the setup. The results, presented in Appendix A, explore cell sizes ranging from $100\ \mu\text{m} \times 100\ \mu\text{m}$ down to $10\ \mu\text{m} \times 10\ \mu\text{m}$.

For thin scatterers, such as the foam samples examined in this thesis, the mean of the scattering angle distribution is expected to be centered around zero and the offset negligible, offering no significant insight into the material budget. Instead, the width of this distribution becomes a critical measure, as it directly correlates with the amount of material the particles have traversed. An example of such a kink angle distribution is shown in Fig. 3.17 for the full surface of the reconstruction plane of the ERG sample for the 1 GeV dataset.




Figure 3.17 – Kink angle distribution (k_{x+y}) for the ERG carbon foam sample at 1 GeV

As explained in the introductory chapter, the scattering angle distribution exhibits a Gaussian core, given by the many small-angle independent scattering events, and long tails, from the rare, large-angle Coulomb scatterings (rather than to the cumulative effect of many small angle scatterings). As it can be seen, there are rare events where the particle is deflected by about 20 mrad or more (for the denser foam and smallest momenta dataset). This would correspond in all cases to less than a 1° deflection. For an inter-plane distance of 2 cm, these extreme cases would correspond to at most 0.4 mm deflection from the original trajectory.

Nevertheless, the vast majority of the scatterings are concentrated in the core of the distribution, within a range of 2.5σ on either side of the mean value. This corresponds to an average deviation of less than 100 μm for most incoming particles between the two planes spaced 2 cm apart.

3.3.2 Estimators of the width of the distribution

The kink angle distribution contains a convolution of effects: the intrinsic angular resolution of the telescope, the contribution of the material budget due to the foam scatterer (which is the focus of this chapter), the misalignment of the telescope planes, contributions from the non-square pixels and the beam momentum spread. While some of these factors can be approximated with Gaussian distributions, their convolution is non-trivial. Therefore, to obtain a single, concise value that effectively summarizes the larger dataset and serves as an estimator of the dispersion of the data (i.e., the width of the distribution), some careful considerations are needed.

Various methods for estimating the width of the scattering angle distributions have been explored in previous studies [103, 104], incorporating both statistical and fitting techniques. Statistical methods might include calculating for example the RMS of the kink angle distribution. Alternatively, fitting approaches may involve for example applying a Gaussian model to the distribution. In the cited works, to account for the influence of distribution tails, which can significantly affect the width measurement, these estimators have been applied to both the complete datasets and specific inner parts (quantiles) of the distribution.

The choice between these estimators depends on the specific characteristics of the distribution being analyzed and the objectives of the study. For instance, while RMS provides a broad measure of distribution width, including outliers, estimators like the average absolute deviation offer a more robust estimate to extreme values (i.e the large angles in the tails of the distribution). Gaussian fitting, on the other hand, assumes a symmetrical distribution (expected from a scattering angle distribution) and focuses on the central peak, omitting parts of the tail information.

For this work, a systematic study of the influence of different estimators on the width of the distribution is performed, addressing a gap not explored in depth in previous works. A set of estimators used to quantify the variability or dispersion within the dataset is defined, based on both statistical and fitting methods.

1. RMS (root mean square)

For a set of n values of a discrete distribution k_i with $i \in (1, n)$, where k_i represents the kink angles k_{x+y} , the RMS is the square root of the mean of the values k_i^2 with \bar{k} the sample mean:

$$\text{RMS} = \sqrt{\frac{1}{n} \sum_{i=1}^n |(k_i - \bar{k})^2|}$$

Although the standard deviation is one of the most common ways of measuring the spread of data, due to the fact that it squares its differences, the RMS tends to give more weight to larger differences and less weight to smaller differences compared to other measures of dispersion.

2. MAD (median absolute deviation)

This effect can be minimized by choosing a more robust metric, like the median absolute deviation. The measure of central tendency ² is chosen to be the median.

$$\text{MAD} = \zeta \cdot \text{median} (|k_i - \text{median}(k)|)$$

Like the standard deviation, it is a measure of dispersion, but more robust to outliers (e.g. the entries in the tail of the kink angle distributions). This estimator employs the sample median twice. Initially, the sample median provides an estimate of the center of the data, which is used to form absolute residuals around this median. Subsequently, the sample median of these residuals is computed.

To make the MAD comparable to the standard deviation, a normalized version of the MAD is used throughout this work. The normalization constant ζ is about 1.4826 ($\approx 1/\Phi^{-1}(75\%)$ ³ = 1/0.6745 [110, 111]). This scale factor for MAD is used for non-normal distributions in order to make it an unbiased Fisher-consistent estimator [112].

²A statistical metric used to describe a whole set of data with a single value that represents the middle or center of its distribution. The most common are the mean, median and mode. The central tendency is often contrasted with the dispersion, which describes the spread of the data around it.

³ Φ^{-1} is the quantile function of the standard normal deviation (or the inverse of the cumulative distribution function)

The reason for this definition is that 0.6745 is the MAD of a standard normal random variable, thus aligning the MAD with the standard deviation scale for large samples. Essentially, $\mathbb{E}(\text{RMS}) \approx \sigma$ and $\mathbb{E}(\zeta \cdot \text{MAD}) \approx \sigma$, implying that the factor scales the MAD to be comparable not to the data, but to the standard deviation estimator of a normal distribution. In other words, a normally distributed variable with mean $\mu = 0$ and variance σ^2 has a normalized MAD of σ . Here, the normal distribution is used not as an assumption, but as a calibration tool; the MAD is multiplied by a factor that, under normality, aligns it with the standard deviation estimator, at least asymptotically. Consequently, the size of the robust z-scores is also comparable to standard z-scores⁴, allowing for the use of normal distribution quantiles for outlier detection.

Given the explanation above, one might be inclined to always use the median and the MAD as robust alternatives to the mean and standard deviation. However, they tend to exhibit poorer statistical performance. The optimal solution combines the best attributes of both approaches: estimators that perform like classical ones when the data lacks outliers, yet remain insensitive to outliers otherwise. Typically, the classical estimates are optimal when the data is distributed according to an idealized parametric model, such as the Gaussian distribution, but become suboptimal when the data distribution deviates from this model. Conversely, robust estimators maintain nearly optimal performance both under an assumed model and with slight deviations from it.

Considering the unimodal distributions with pronounced tails of the kink angles (with the bulk of approximately 98% of the distribution being Gaussian), MAD serves as a robust alternative to the RMS.

3. AAD (Average Absolute Deviation)

Similar to the MAD, the average absolute deviation can be calculated as a robust estimator. In this work, the measure of central tendency is chosen to be the mean, such that the formula used to calculate the quantity is:

$$\text{AAD} = \zeta \cdot \frac{1}{n} \sum_{i=1}^n |(k_i - \bar{k})|$$

⁴The z-score measures how many standard deviations a data point is from the mean of that dataset. Since the mean is sensitive to outliers, the standard z-score also inherits this sensitivity. A robust z-score mitigates this issue by using the median instead of the mean.

The AAD is a robust estimator, providing a good fit to the bulk of the data when the data contains outliers, like the large kink angles that form the tail of the distribution.

Similarly to the MAD, a scaling factor is needed for the AAD to have comparable results with the standard deviation. This factor is $\sqrt{\pi/2} \approx 1.2533$. In other words, for a normal distribution, the AAD is about $1/1.2533 \approx 0.8$ times the standard deviation.

4. BMWV (Biweight-midvariance)

Yet another possible estimator robust to violations of normality is called the biweight-midvariance [113]. This estimator empirically determines whether a value is unusually large or small and discards them, while the variation among the remaining values is computed.

If one defines the quantity:

$$u_i = \frac{k_i - \text{median}(k)}{c \cdot \text{MAD}(k)}$$

where c is a tuning constant (larger c indicates more data values are included in the computation of the statistic; the usual value in literature, also used in this work, is $c = 9$), then the biweight-midvariance can be defined as:

$$\xi_{bmwv} = n \frac{\sum_{i=1}^n (k_i - \text{median}(k))^2 (1 - u_i^2)^4}{\sum_{i=1}^n (1 - u_i^2) (1 - 5u_i^2)^2}, |u_i| < 1$$

Its square root $\sqrt{\xi_{bmwv}}$ is a robust estimator that is also investigated in this work.

5. σ_{Gauss} estimator

In addition to the statistical estimators previously mentioned, a fitting procedure assuming an underlying distribution can be employed. The simplest approach is to fit a Gaussian distribution to the data. This is motivated by the observation that the core of the kink angle distribution follows a normal distribution, in accordance with the central limit theorem for the large number of scatterings. Moreover, the Highland formula also describes the inner core of the scattering distribution using a Gaussian approximation, with its width given by 1.7.

The fit function follows the following formula:

$$f_{\text{Gauss}}(k_{x,y}) = \frac{N}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right)$$

The standard deviation of this distribution, further denoted σ_{Gauss} is used as an estimator.

6. $\sigma_{\text{G+St}}$ – Convolution of Gauss and Student's t estimator

As demonstrated in [114], the sum of a Gaussian and a Student's t distribution can effectively describe the full kink angle distribution. The Gaussian component represents the multiple scattering core of the distribution, while the Student's t distribution accounts for the distribution tails.

The Student's t distribution is a continuous probability distribution that generalizes the standard normal distribution. It is symmetric around the zero and bell-shaped, but with heavier tails. The parameter ν controls how much of the total probability (or likelihood) is found in the tails of the distribution. For $\nu = 1$ the Student's t distribution becomes the standard Cauchy distribution, characterized by very “fat” tails. Conversely, as $\nu \rightarrow \infty$, it approximates the standard normal distribution $\mathcal{N}(0, 1)$ which has very “thin” tails.

The convolution of both distributions is given by:

$$f_{\text{G+St}}(k_{x,y}) = N \cdot \left((1-a) \cdot \frac{1}{\sigma_G\sqrt{2\pi}} e^{-\frac{(k_{x,y}-\bar{k})^2}{2\sigma_G^2}} + a \cdot \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\sigma_S\Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{(k_{x,y}-\bar{k})^2}{\nu\sigma_S^2}\right)^{-\frac{\nu+1}{2}} \right) \quad (3.4)$$

The kink angle distributions are fitted using this convolution that has six free parameters: an overall normalization N , the relative fraction a of the Student's t distribution, a common mean μ , the width of the Gaussian (σ_G) and that of the Student's t distribution (σ_S), as well as the tail parameter ν .

The estimator is denoted as $\sigma_{\text{G+St}}$ and is calculated as a weighted sum of the standard deviations of the two separate parts:

$$\sigma_{\text{G+St}} = \frac{\sigma_G A_G + \sigma_S A_S}{A_G + A_S} \quad (3.5)$$

where A_G and A_S represent the integrals of the bin contents for the Gauss and for the Student's t distributions, respectively, within the range of definition of the Gaussian fit function (i.e. excluding the contribution from the tails).

To illustrate this, an exemplarily dataset was fitted using a convolution of Gaussian and Student's t distributions using eq. 3.4. This is shown in Fig. 3.18. The top part shows the kink angle distribution k_{x+y} in blue, alongside the convolution fit in red. Additionally, the individual contributions of the Gaussian and Student's t distributions to the convoluted fit are displayed, using the parameters derived from the combined fit.




Figure 3.18 – Kink angle distribution k_{x+y} for the ERG foam sample measured at $1 \text{ GeV } c^{-1}$, fitted with a convolution of Gaussian and Student's t distributions. The top plot shows the data and the convolution fit, with individual contributions from the Gaussian (yellow) and Student's t (green) distributions. The inset displays the fit parameters of the convolution and their values. The middle plot shows the ratio of the fit to the data, and the bottom plot shows the residuals (pulls) of the fit.

It can be seen that the convolution of the two distributions describes well the full kink angle dataset. A χ^2/ndf of 1 is achieved which is consistent with the results for other datasets examined. The values of the six free parameters are also displayed using the same color as the fit. For this specific dataset, the contribution of the Student's t distribution to the convoluted fit is approximately half that of the Gaussian distribution. Furthermore, the separate standard deviation value for the two convolution components are shown, as well as the combined one calculated according to eq. 3.5.

In order to check the robustness against outliers of the aforementioned estimators, a quick check is performed on all estimators defined above, with the exception of the last one.

A clean data sample of 1M events is generated, following a normal distribution with a mean of 0 and a standard deviation of 1. To simulate scenarios which deviate from normality, contamination is introduced into the clean dataset. Specifically, a fraction of up to 10% of the data points is replaced with values drawn from a Student's t distribution with a tail parameter $\nu = 3$, which introduces heavier tails compared to the normal distribution.

As will be shown later, it is expected that the tails take at most 3-4% of the whole distribution, so these large fractions are not expected. This exercise is exclusively used to illustrate the behavior of the estimators to outliers.

To evaluate the robustness of these estimators, their sensitivity to different levels of contamination is measured by varying the number of outliers. Contaminated datasets are generated, and the value of each estimator is calculated at different contamination levels. The performance of each estimator is presented in Fig. 3.19. Sensitivity curves are plotted to illustrate how the value of each estimator changes with the percentage of outliers.




Figure 3.19 – (simulation) Evaluation of the robustness of various width estimators against outliers. **(left)** Sensitivity curves of different width estimators as a function of the percentage of outliers. **(right)** Histogram comparison of clean data ($\mathcal{N}(\mu = 0, \sigma = 1)$) and contaminated data (Student's t distribution ($\mu = 0, \nu = 3$)). The filled area shows the inner 90% of the contaminated data, the maximum percentage of outliers under study.

For the Gauss + Student's t distribution, the nature of the convolution and the free parameters do not easily allow such a test. The contaminated samples will be accommodated by changes of the tail parameter and interplay of the two sigma of the distributions and ultimately might not represent the true behavior of the convolution.

Inspecting the results, some conclusions can be drawn about the robustness of different width estimators in the presence of outliers, at least in this simplified example.

The RMS estimator is highly sensitive to outliers. As the percentage of outliers increases, the RMS value rises fast. This behavior reflects the lack of robustness of this estimator in the presence of large values, as it gives equal weight to all data points, making it particularly susceptible to deviations from normality.

The AAD estimator also shows sensitivity to outliers, but to a lesser extent compared to the RMS. It increases steadily as the percentage of outliers increases, suggesting that while it is more robust than RMS, it is still affected by the values in the tails of the distribution.

The other estimators (MAD, BMWV and σ_{Gauss}) demonstrate considerable robustness, exhibiting only a very slight increase as the fraction of large values grows, making them a reliable choice for datasets such as those found for the kink angle distributions.

This small example provides only a preliminary check. Since the “contamination” in the actual foam data can differ, a comprehensive evaluation is necessary to accurately measure the robustness of these estimators.

3.3.3 Quantiles

The aforementioned estimators can be utilized on the entire distribution or confined to a specific inner segment. One effective method to partition a dataset is by employing quantiles.

Quantiles are cut points that divide the range of a distribution into continuous intervals with equal probabilities. Specifically, q -quantiles split a finite set of values into q subsets of (nearly) equal size, with $q - 1$ partition points corresponding to each integer k that satisfies $0 < k < q$. This approach can also be extended to continuous distributions, allowing for the generalization of rank statistics⁵ to continuous variables.

In both discrete and continuous population densities, the k -th q -quantile is the data value where the cumulative distribution function intersects k/q . In other words, x is a k -th q -quantile for a variable X if

$$Pr[X < x] \leq k/q \quad \text{or, equivalently,} \quad Pr[X \geq x] \geq 1 - k/q$$

⁵Ranking is the data transformation in which numerical or ordinal values are replaced by their rank when the data are sorted

To investigate the effects of the tail of the distribution and its influence on the estimators, quantiles are used in the thesis to restrict the estimators of width on a subset of samples from the core of the kink angle distributions. The n -th percentile divides a dataset such that $n\%$ of the values lie below and $(100 - n)\%$ lie above. With this definition, a set of quantiles is defined in this work, starting from the full distribution denoted as 100% and restricting to the inner 80% of the kink angle distribution. For example, Gauss_{95} is the σ_{Gauss} estimator applied to the inner 95% of the kink angle distribution.

In this work, the following quantiles are defined and further used: 100, 99, 98, 97, 95, 93, 90, 85, 80%.

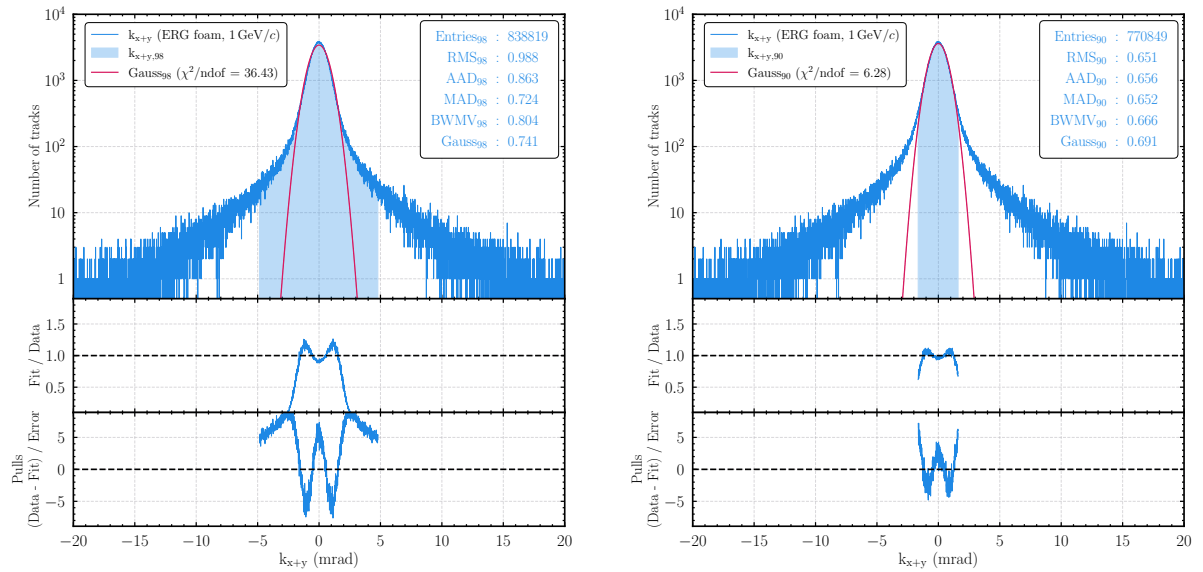


Figure 3.20 – Kink angle distributions fitted with Gaussian models for the ERG foam at 1 GeV/c considering (left) 98% and (right) 90% of the inner distribution. Both plots showcase the calculation of the width estimators for the two cases. The ratios of fit to data and the pulls are also shown.

Figure 3.20 illustrates an example of the kink angle distribution for the ERG foam sample taken at a momentum of 1 GeV/c. The left panel shows the calculation of estimators restricted to the inner 98% of the distribution, while the right panel further restricts it to the inner 90%, highlighted by the blue shaded region.

It is immediately apparent that the value of the width estimators vary between the two quantiles, decreasing as the core of the distribution becomes smaller. Notably, while there is considerable variation among the estimators at the 98% quantile (where the influence of the tails can still be seen and which is the typical quantile used in conjunction with the Highland formula), they converge to similar values at the 90% quantile, irrespective of the chosen estimator.

The ratio of fit to data, together with the pulls, pertain to the Gaussian fit, which, alone, fails to properly describe the data, particularly towards the tails of the distribution. At the 90% quantile the influence of the tails is diminished. Still, deviations from normality persist, suggesting that the core of the distribution may not be entirely Gaussian.

In the following subchapters, the combination of quantiles and estimators is explored to examine their impact on describing the bulk of the distribution and addressing the tails.

Some of the following figures in the rest of this chapter will show various quantities using different estimators and quantiles. This choice does not impact the results, but serves to show that any of them can be used effectively.

3.3.4 Position resolved scattering angle distribution

To differentiate between tracks traversing areas with and without foam, a two-dimensional, position-resolved approach for analyzing the kink angle distributions is utilized. As mentioned in Sec. 3.3.1, this method involves mapping the combined kink angles k_{x+y} onto corresponding image cells on a virtual reconstruction plane defined at the middle position of the scatterer.

The results of the 2D position-resolved analysis are shown in Fig. 3.21, which depicts the width of the kink angle distributions confined to the area where tracks are reconstructable, a fiducial region slightly smaller than the ALPIDE sensors themselves due to slight sensor misalignment. The measurement reveals a pronounced region of larger scattering angles at the center of the reconstruction plane, matching the dimensions ($10 \times 3 \text{ mm}^2$) of the carbon foam sample.




Figure 3.22 – Close-up of an ALPIDE on a carrier card. The bonds are visible on the top side of the figure, close to the digital periphery of the chip. The red dashed line indicates the opening in the PCB, a region with minimal scattering.

Additional areas of increased material budget are identified, correlating with the structural design of the carrier card on which the ALPIDE sensor is bonded. Figure 3.22 shows the ALPIDE on the carrier card, alongside a red region where the PCB has a cutout to allow in-beam measurements of the sensors. Three specific regions needed to ensure stable sensor contact with the carrier can be seen outside this red area: two corners and a longitudinal strip near the edge of the sensor that is bonded. These areas contain on top a part of the active matrix, resulting in an increased material budget that can be measured, as evident in the 2D histogram.




Figure 3.21 – Superposition of kink angles (k_{x+y}) per $100\,\mu\text{m} \times 100\,\mu\text{m}$ image cell, representing a measure of the material budget distribution at the position of the sample. The values in each cell are calculated using the RMS_{100} estimator. Regions of interest used to select only kink angles corresponding to the foam (green) or only to no-foam (yellow or orange) regions are shown.

Sensor misalignment within the telescope array further impacts reconstruction accuracy, reducing the effective area for track reconstruction. This misalignment is noticeable along the edges of the 2D histogram. Ideally, a perfect alignment would make the full 30 mm length of the ALPIDE sensor accessible in the x direction, but only about 26 cm are usable. Moreover, in the y direction, part of the sensor has the PCB behind, near the periphery of the chip, which leads to increased scattering.

Despite these limitations, a significant portion of the sensor area remains suitable for accurate kink angle measurements, allowing for comprehensive analysis with minimal measurement bias.

To accurately evaluate the material budget and isolate specific contributions from the foam and silicon sensors only, regions of interest (ROIs) are defined where exclusively foam or silicon sensor material is present. The extent of these ROIs is important in order to minimize effects stemming from the presence of momentum gradients along the sensor surface, which result from the beam generation at the DESY II beamlines.

Once ROIs are defined (for example containing only the foam part), the information from the kink angles for both x and y axes is integrated over the full surface of the ROI. These kink angle distributions for various ROIs are then further used.

3.3.5 Momentum divergence considerations

The kink angle distributions for the ALLCOMP sample for a foam-only ROI, observed at two different beam momenta, illustrate the effect of beam momentum on the scattering angle behavior. As depicted in Figure 3.23, the increase in beam momentum results in a narrower kink angle distribution, consistent with the $1/p$ dependency of the scattering angle width (θ_0) as described by the Highland formula (cf. eq. 1.7). The MAD_{100} estimator is used in this example as a measure of the width of the two distributions.




Figure 3.23 – Kink angle distributions measured for electrons with momenta of 1 and 5.4 $\text{GeV } c^{-1}$ scattering with the ALLCOMP sample.

At the DESY II test beam facility, the final beam generation steps involve finely tuning the particle momentum using a dipole magnet and a primary collimator, ensuring that only electrons with the selected momenta are delivered into the testbeam area. A secondary collimator, positioned by users, further refines the spatial profile of the beam depending on the experimental purposes. Despite these controls, the beam exhibits a finite momentum spread, which depends mostly on the horizontal position across the aperture of the collimators, and potentially on the vertical position as well. This variation cannot be ignored in material budget measurements, as the kink angle width can vary along the horizontal axis, given the direct proportionality with the momentum.

Previous measurements [86] on beam line 21 (another beamline than the one used to take data in this thesis, but the only one for which such a measurement is available) have quantified this momentum spread, recording an absolute momentum variation of $158 \pm 6 \text{ MeV}/c$ across the momentum range of the beam. This absolute spread results in a decreasing relative momentum uncertainty with increasing beam energy, from approximately 15% at $1 \text{ GeV}/c$ to about 2% at $5.4 \text{ GeV}/c$.

This effect was investigated in this study. Figure 3.24 shows the two-dimensional position-resolved kink angle histogram for the case where no scatterer is present.

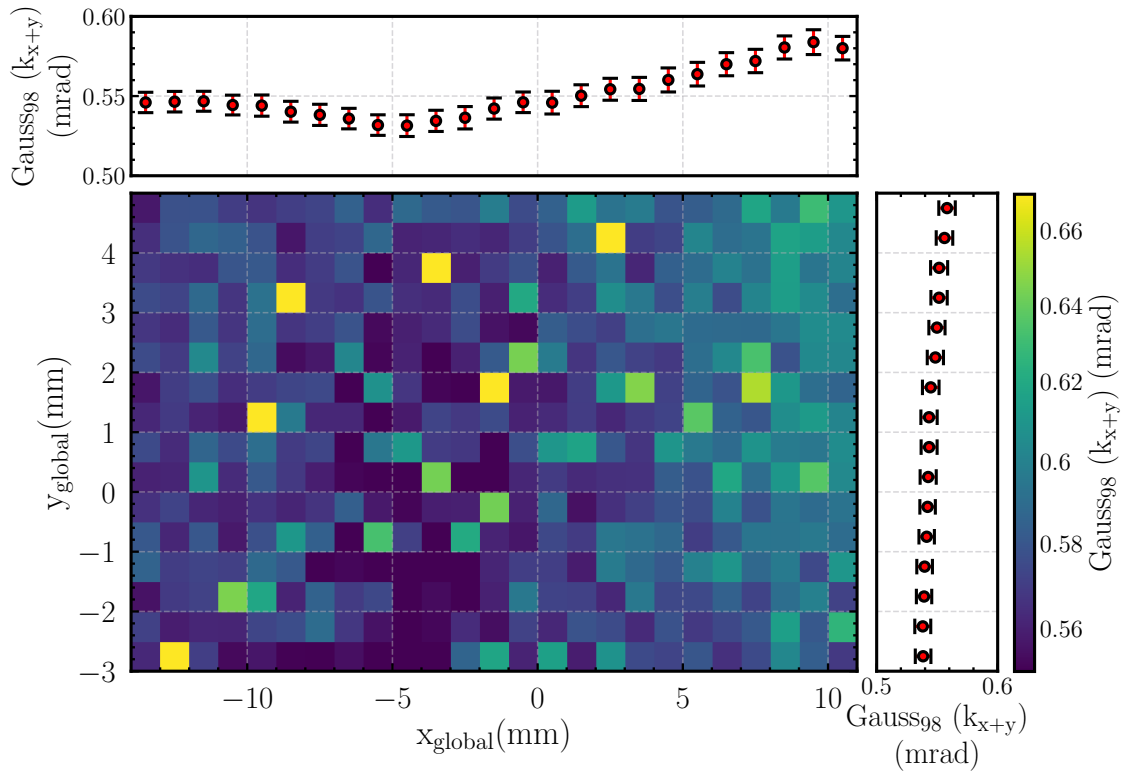


Figure 3.24 – Spatially-mapped distribution showing the Gauss98 estimator and quantile pair for the width of the kink angle distribution measured at 1 GeV in the absence of any scatterer from the telescope (i.e. an “empty” telescope). A region smaller than the sensor surface is displayed, excluding areas around the edges where scattering with the PCB occurs. Projections along the x and y axes highlight contributions from the horizontal momentum gradient.

A gradient is observed, with larger kink angles towards the positive x direction. A smaller, non-zero gradient is also visible along the y axis. The bins in each direction are averaged, resulting in two projection scatter plots.

The absolute value of the beam momentum at a specific x position is not known. The fact that the horizontal gradient is not strictly monotonic is not currently understood. Nonetheless,

in the region where it is strictly increasing, a relative variation of $(0.479 \pm 0.037) \text{ \%/mm}$ in the x direction and $(0.659 \pm 0.018) \text{ \%/mm}$ in the y direction is observed, corresponding to approximately 5 MeV/mm and 7 MeV/mm, respectively. These values are consistent with observed changes of the θ_0 angle in the Highland formula for the same amount of material considering the difference in momenta ($p \pm \Delta p$). Moreover, this is consistent with previously measured values [102, 104] at DESY II. The variation shown here is the largest, since this is the smallest momentum for which data was taken. For larger momenta, a smaller relative variation is expected.

No further efforts were made by the author to investigate this effect in more detail. However, under the supervision of the author, a master thesis is being prepared, examining this phenomenon in more detail [115]. A telescope with a better kink angle resolution is used and more statistics are available. A two dimensional fit is applied over the surface of the sensor and the momentum divergence is extracted and corrected on a bin-by-bin basis.

The influence due to the momentum variation along the x axis is expected to be minimized when selecting regions that span the same x range for a particular run and studying them together. This strategy will be employed in the studies from the next subchapters. Nonetheless, a small systematic uncertainty is considered further. Since a gradient in y is also observed, a total variation of 20 μrad is further carried as a systematic uncertainty due to the momentum spread on all values for the remainder of this chapter.

As noted in the previous subchapter (see Sec. 3.3.4), the extent of the ROIs is important in order to minimize the effect of the momentum gradient. The potential bias can be reduced by equalizing the extent of the ROIs in the x -direction. Therefore, subsequent measurements are taken from regions with the same extent in the x direction.

For example, in the case of the foam measurement, two ROIs will be defined: one where there is no foam and one where only the foam is present. The extent of these ROIs along the x axis is chosen to be the same (see Fig. 3.21). To account for the no-foam contribution, an average is taken from two regions adjacent to the foam region, one located above and one located below the foam in the 2D histogram (yellow and orange ROIs in Fig. 3.21). This helps minimize the impact of momentum divergence along the y axis.

The same procedure is repeated for calibration measurements, as discussed in the next subchapters.

3.3.6 Corrections and calibrations

The shape of the kink angle distribution is influenced by multiple factors and can be represented by a model function $f_{\text{measured}}(\theta)$ as follows:

$$f_{\text{measured}}(k_{x+y}) = f_{\text{Highland}}(x/X_0, \theta_{\text{HL}}) \otimes f_{\text{telescope}}(\theta_{\text{telescope}}) \quad (3.6)$$

The measured angular distribution results from the convolution of two Gaussian probability density functions: one representing the Highland model and the other accounting for telescope effects.

The Highland part can be approximated with a Gaussian as a result of the core of the multiple scattering distribution. The width of this distribution is described by eq. 1.7 which depends on both the momentum and the material budget. As such, the first part of the convolution can be written as:

$$f_{\text{Highland}} = \frac{1}{\varsigma \cdot \theta_{0,\text{HL}} \sqrt{2\pi}} e^{-\frac{(k_{x+y} - \bar{k})^2}{2\varsigma \cdot \theta_{0,\text{HL}}^2}} \quad (3.7)$$

An overall scaling factor ς is introduced to account for deviations from the true angular distribution shape of this Gaussian approximation. On top, such a term is needed when using other estimators than the Gaussian fit for which the Highland formula is applicable, as well as for the different quantiles used. It compensates for any incorrect assumptions that might cause a global offset in the Highland curve. The value of ς is determined through calibration measurements, as will be later shown in this chapter.

The second part of the convolution considers effects stemming from the telescope setup and other global influences, such as the momentum divergence. This component is also expressed as a Gaussian (simplistic assumption) with a standard deviation $\sigma_{\text{telescope}}$ given by:

$$f_{\text{telescope}} = \frac{1}{\theta_{0,\text{tele}} \sqrt{2\pi}} \exp -\frac{(k_{x,y} - \bar{k})^2}{2\theta_{0,\text{tele}}^2} \quad (3.8)$$

This quantity is also determined through a dedicated set of measurements where no scatterer is present in the telescope, but only the presence of the air and the telescope planes is affecting the measurement.

For the scope of this thesis, a proper deconvolution of the two functions is not attempted. Instead, a quadratic subtraction will be employed to compensate for effects arising from the telescope setup, as will be explained below.

Considering the case of a foam sample measurement, the two contributions will be quadratically subtracted in order to estimate only the scattering angle due to the presence of the foam:

$$\theta_{\text{foam}} = \sqrt{\epsilon_{q, \text{foam}}^2 - \epsilon_{q, \text{no-foam}}^2} \quad (3.9)$$

Here, ϵ stands for the width of the scattering angle distribution given by the chosen estimator, and q for the quantile.

3.3.6.1 Calibration samples

During a testbeam campaign in the spring of 2024, new data was collected using a carbon foam sample that is closer to the final design of the ITS3. Those results are subject to a master thesis in preparation [115], carried out with the help of the author, and address several shortcomings identified prior to the testbeam and presented in this work. Notably, the telescope was configured with larger distances between planes to achieve better angular resolution, and the samples were carefully prepared with known amounts of glue.

During this measurement campaign, there was a brief opportunity to collect data using a setup similar to that used for the carbon foam measurements in this thesis. Over a period of about six hours, a setup like the one depicted in Fig. 3.9 was used. Data was gathered using calibration targets, from which the scaling factor ς (see eq. 3.7) is later calculated, as well as data with no scatterer present (i.e., an “empty” telescope), needed to estimate the contribution of the telescope.

The calibration target holder, shown in Fig. 3.25, consists of a 3D-printed holder that secures three sheets of metal (5 cm × 5 cm) of varying thicknesses. Both nickel and aluminum targets were used for calibration, but only one flavor was used per measurement. The target areas were larger than the ALPIDE sensor, and stacking the targets minimized the need for frequent target exchanges. Additionally, the targets were aligned along the long edge of the ALPIDE to allow construction of ROIs that counteract the momentum divergence along the x axis. These ROIs are designed to span the same range in x across all four




Figure 3.25 – 3D printed calibration target holder. The measurement area of the ALPIDE sensors and four sub-measurement areas are shown.

regions. The momentum spread is expected therefore to affect all four regions in the same manner along this direction. A region without a target (region 4) serves as a reference point for comparison between different runs and calibration targets.

For the nickel calibration measurement, the three sheets were 25, 50, and 90 μm thick, corresponding to regions 3, 2, and 1, respectively. Due to overlap, region 2 had a total thickness of 75 μm , and region 1 had a thickness of 165 μm . With a radiation length of $X_{0,\text{Ni}} = 1.424 \text{ cm}$ [94], the material budget ranges from 0.18% to 1.16% X_0 .

The aluminum samples were 70, 200, and 200 μm thick (without overlap), corresponding to regions 3, 2, and 1, respectively. When overlapped, region 1 has an equivalent thickness of 470 μm aluminum and region 2 270 μm . Considering for Al a radiation length of $X_{0,\text{Al}} = 8.897 \text{ cm}$, the material budget spans a range of 0.08% to 0.53% X_0 .

These target thicknesses were selected to constrain the expected values for the two foam samples, enabling interpolation rather than extrapolation and providing a wide range for a robust fit, as will be demonstrated later.

Using this stacked configuration and quadratic subtraction, six measurement points can be derived per target flavor. Taking the case of the Ni target, these are:

- 25 μm Ni : by subtracting in quadrature region 4 from region 3
- 50 μm Ni : by subtracting in quadrature region 3 from region 2
- 75 μm Ni : by subtracting in quadrature region 4 from region 2
- 90 μm Ni : by subtracting in quadrature region 2 from region 1
- 140 μm Ni : by subtracting in quadrature region 3 from region 1
- 165 μm Ni : by subtracting in quadrature region 4 from region 1

Similar calculation can be achieved for the Al targets.

There is an underlying assumption that the quadratic subtraction completely removes the contribution from the telescope effects. Careful considerations are made to ensure this, such as maintaining the same extent of the ROI in the x direction during quadratic subtraction, in order to have the same momentum divergence bias in both measurements. Nevertheless, several checks are performed to assess the influence of quadratic subtraction on the results.

3.3.6.2 Quadratic subtraction tests

A first check is performed by comparing the regions with air (region 4 in Fig. 3.25) for the cases of Ni and Al, as well as the scenario where no target is inserted. Under the

assumption that the quadratic subtraction fully cancels the telescope effects, there is no momentum divergence in the vertical direction, and the setup is similar in the three cases, the measurements should be compatible: by comparing any two of these values and subtracting one from the other, the resulting difference should be close to zero.

Figure 3.26 shows the results of the quadratic subtraction between these three different cases. The y axis represents the difference between any of the three combinations taken in pairs, while the x axis shows the different kink angle width estimators. Within each estimator region, the quantiles are depicted by the transparency of the data points: full color represents the full distribution ($q = 100\%$) and the least transparent markers represent the inner core of the distribution ($q = 80\%$). Different colors indicate the three energies at which data was taken, and symbols mark the type of quadratic subtraction combination: either between the empty telescope and the Ni or Al air region, or between the Ni and Al air regions themselves.




Figure 3.26 – Results of the quadratic subtraction between the three different cases. The y axis represents the difference between any of the three combinations taken in pairs, while the x axis shows the different estimators. Within each estimator region, the quantiles are depicted by the transparency of the markers. The $\sigma_{\text{Gauss}} + \text{Student's } t$ is only defined for $q=100\%$.

From the figure, it is evident that although values fluctuate around zero, there can be a systematic shift of up to $40 \mu\text{rad}$. Moreover, depending which two quantities are subtracted, values seem to group together, showing a specific offset from zero based on the momentum at which the data was taken. These results indicate that the quadratic subtraction does not entirely cancel out the telescope effects, leaving behind some residual systematic shift. In the following, some of the contributions (marked in bold) to this effect are discussed.

Part of this shift can be attributed to the measurement precision of the distances between the sensor planes, as well as setup changes due to ambient **temperature variations**, which are not accounted for. The alignment is performed once per run type and maintained throughout.

There are variations between day and night and over shorter periods which are then not accounted for.

Small changes of up to 1 mm in the **distances between the telescope planes** can lead to changes of up to 20 μrad in the angular resolution of the telescope, as demonstrated in the figures from Appendix B. The precision of measuring the z distances⁶ between planes at the test beam is at best on the order of 1 mm. Due to the way the target holder is mounted, the measurement of distances between planes may be slightly off between the empty telescope and the one with the target holder. The change between the Ni and Al targets was done by swapping the targets without altering the setup, whereas for the empty telescope, the telescope arms were opened, modifications were made, and then they were reassembled. This procedure can add small variations to the setup which might not have been detected and corrected during the testbeam. This assumption is validated by the results in Fig. 3.26, which shows that the quadratic differences between the air regions in the Ni and Al setups are very close to zero.

The ordering of the quadratic differences with respect to the **momentum change** suggests another effect. The Highland formula is **non-linear**, and the quadratic subtraction that is performed is only an approximation. When performing a quadratic subtraction, the logarithm term introduces a residual difference. This is discussed in detail in Appendix C.

This is discussed in detail in Appendix C, which shows that there can be a few percent difference between the quadratic subtraction term and the true material value (as given by the Highland formula). A concrete example is shown in the following paragraphs.

For the case where the quadratic subtraction is performed with a region not containing any material (“empty” ; silicon-only), it is challenging to quantify such a change in the Highland formula, as there is no absolute value to reference for the air. However, for the calibration targets this is done, as shown further on in this subchapter. Moreover, for regions where nothing but the silicon is present, the value of the measurements go below the sensitivity level (smallest thickness that can be accurately measured). Therefore, for any quadratic subtraction involving air, a value of 30 μrad is chosen as a systematic uncertainty.

A second check of the quadratic subtraction is performed on a dedicated set of Al data. As previously mentioned, two separate Al calibration targets of 200 μm and one of 70 μm are overlapped and measured at the same time. By selecting the pair of Al regions with 470 μm and 270 μm , or the 270 μm and the 70 μm equivalent thickness, the quadratic subtraction can be performed and should yield an equivalent of 200 μm of aluminum thickness in both cases.

⁶The distances between planes represent a weak mode in the alignment and therefore need to be inputted.

The resulting values can be compared with scattering angle computed with the Highland formula for a 200 μm slab of aluminum.

These two checks, along with the checks on the three air regions, help test the quadratic subtraction, extract systematic uncertainties, and identify any remaining effects.

For the Ni/Al case, the non-linearity introduced by the quadratic subtraction using the Highland formula can be corrected. Figure 3.27 shows the results of applying the corrections discussed in Appendix C. Further results will contain this correction due to the non-linear nature of the Highland formula under quadratic subtraction.

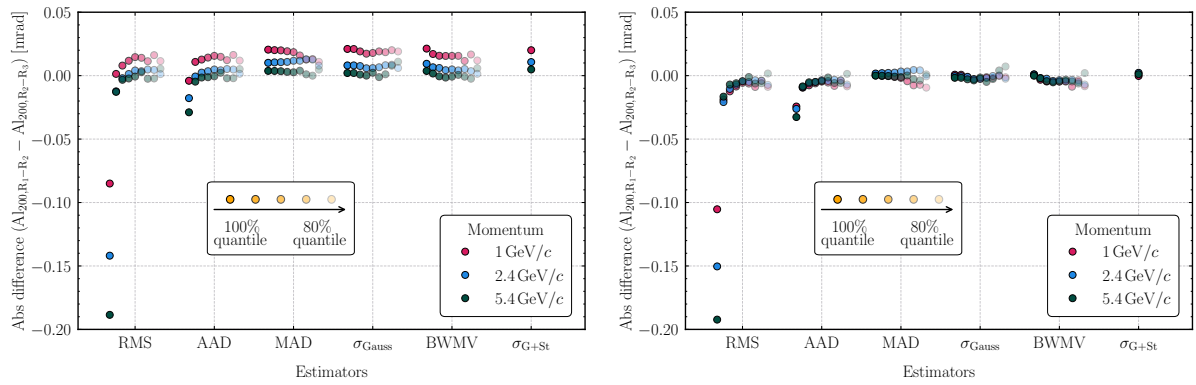


Figure 3.27 – Absolute difference between quadratic subtractions performed on two separate regions of the Al calibration targets that are expected to yield the same material contribution (200 μm Al): **(left)** before correcting the Highland formula for the remainder of the logarithm term (as discussed in Appendix C) and **(right)** after the correction is applied.

The figure illustrates the absolute difference between the two quadratic subtractions for different sets of regions across various chosen estimators and quantiles. Similar to the quadratic subtraction of the air regions previously discussed, a non-zero remainder persists. Additionally, a similar ordering with momenta is observed. After correcting the Highland formula for the non-linear part, the results are closer to zero, as expected. This correction is not applied in the case of the carbon foam samples or in calculations involving the no scatterer present in the telescope, since there no absolute value is known.

However, a small deviation remains unaccounted for, on the order of up to 20 μrad , dependent on the quantile and estimator. This deviation aligns with the systematic uncertainty observed in the previous measurements with air regions and is further considered during the quadratic subtractions involving the Ni/Al samples as a systematic uncertainty.

This non-linearity correction does not completely resolve the difference between the measured data and the Highland model expectation. To take a concrete example with the

σ_{Gauss} estimator at the 98% quantile, consider the quadratic subtraction between Ni region 1 (equivalent to 165 μm) and Ni region 3 (25 μm) data at 1 GeV, as shown in Table 3.5.

Table 3.5 – Differences between measurement values and the theoretical expectation from the Highland formula for the quadratic subtraction between two Ni regions. Beam momentum is 1 GeV, the estimator for the data values is σ_{Gauss} and the quantile is 98%.

	Measurement (mrad)	Theoretical Highland value (mrad) (no telescope effects)
θ_{165}	1.543	1.216
θ_{25}	0.725	0.432
$\sqrt{\theta_{165}^2 - \theta_{25}^2}$	1.362	1.136

Performing the quadratic subtraction for the theoretical (Highland) expectation yields:

$$\sqrt{\theta_{\text{Ni},165(\text{Highland})}^2 - \theta_{\text{Ni},25(\text{Highland})}^2} = 1.136 \neq 1.112 = \theta_{\text{Ni},140(\text{Highland})}$$

Taking the theoretical Highland values of $\theta_{\text{Ni},165(\text{Highland})}$ and $\theta_{\text{Ni},25(\text{Highland})}$ and quadratically subtracting them, yields a value (1.136) that is about 2.2% larger than expected (1.112) for this material thickness and energy. This value comes from the non-linearity of the Highland formula and can be corrected for, as explained previously and detailed in Appendix C.

Performing the same calculation with the values from the data, a value that is 20% larger is obtained:

$$\sqrt{\theta_{165,\text{data}}^2 - \theta_{25,\text{data}}^2} = 1.362$$

The measured values each contain the convoluted effect of the telescope, which cannot be corrected here. Performing the same calculation on other Ni or Al data results in values consistently about 20% larger (for this specific estimator and quantile choice). This disparity will later be assimilated to the ς scaling factor discussed earlier in this subchapter and further explained below. Other estimator and quantile choices will yield different percentages.

The stability of the estimators and quantiles after the quadratic subtractions in the case of the Al/Ni target samples, as well as for the foams, is investigated next.

The quadratic subtraction, corrected for the non-linearity in the Highland formula, is applied to various region pairs for the Al/Ni samples. The resulting value from data is then divided by the corresponding thickness value derived from the Highland formula, which has also been corrected for non-linearity. These ratios for the different region pairs are then

averaged for each estimator and quantile, and are exemplarily shown for the 1 GeV case in Fig. 3.28.




Figure 3.28 – The ratio between the value obtained from quadratic subtraction of various region pairs for the Al/Ni samples to those predicted by the Highland formula, for the case of 1 GeV momentum. The values are corrected for the non-linearity of the Highland formula after quadratic subtraction. Each point represents the average value across the various region pairs and their corresponding thickness differences, extracted for the specific estimator and quantile.

The values for Ni are systematically larger than those for Al, due to the smaller radiation length of Ni during the quadratic subtraction.

An important observation in Fig. 3.28 is the varying stability of the estimators. It can be seen that the RMS is the most unstable estimator, followed closely by the AAD. A small change in quantile results in a large change in the resulting kink angles, and therefore the ratio.

The BWMV estimator also exhibits a slight variation. In contrast, MAD and σ_{Gauss} estimators are the most robust, showing less sensitivity to tighter tail cuts. Notably, the $\sigma_{\text{G+St}}$ estimator is nearly consistent with one.

Although the initial values ($q = 100\%$) are consistent among the last six estimators, indicating robust width descriptions of the whole distribution, a good estimator in this context would maintain stability as the quantile is adjusted.

In the case of the σ_{Gauss} estimator, the value remains stable up to the removal of approximately 3% of the tails. Beyond this point, the ratio starts decreasing, indicating that the Gaussian part of the distribution is being cut. This can be a good indicator of the proportion of tails in a specific distribution.

This figure served to show the necessity of the ς correction for each quantile and estimator. This correction will be further discussed and applied at a later time.

After checking the robustness of the estimators for the nickel and aluminum samples after quadratic subtraction, this procedure is also checked in the case of the foam data. Kink angles were extracted using the aforementioned estimators and quantiles for both a region with foam and a region without.

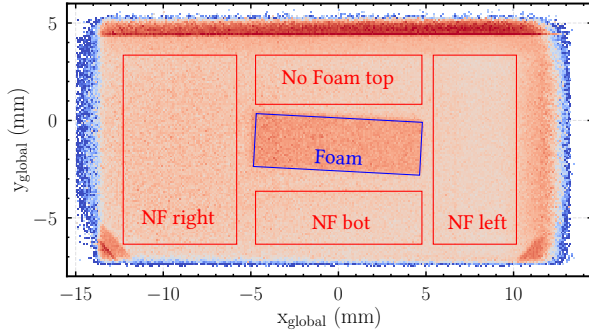


Figure 3.29 – Foam and no-foam ROI definitions.

For each foam type, several regions were defined, as shown in Figure 3.29. The region where the foam contribution is present is marked with a blue rectangle. Four regions where only the contribution of the telescope is present are marked in red. The top and bottom “no-foam” regions are chosen with the same extent as the foam region to minimize the influence of the momentum divergence during quadratic subtraction.

Minimal differences are observed between the top and bottom regions, while a mild gradient is inferred from the right-left difference. No effort is made to correct the effect of this gradient, but it has to be taken into account. In order to reduce the mistake made by not correcting, the effect can be mitigated by actively choosing the same extent in x for the ROIs (see Section 3.3.5). An average between top and bottom no-foam region is taken as the no-foam contribution.

After quadratically subtracting the no-foam region from the foam region, the results are inspected for estimator and quantile stability and robustness. This is exemplified for the ALLCOMP sample at 1 GeV in Fig. 3.30.

The results show that the RMS and AAD estimators are rather unstable until the 95% quantile, where their values approach those of the other estimators. The MAD, σ_{Gauss} and BMWV estimators initially cluster around the same value and remain stable until the 95% quantile, where they start to decrease due to the strong data cut. The σ_{Gauss} estimator remains the most stable, with only a small reduction for the full quantile range.

The estimators are consistent around the 90% quantile with the uncertainty band, which represents the manufacturing knowledge of the sample, including uncertainties in material thickness. The $\sigma_{\text{G+St}}$ estimator has the smallest value but matches the expectation from the sample material contribution well.

The results for the other foam samples and momenta are shown in Appendix D. It was observed and expected for all studied cases that cutting too much into the kink angle distribution alters the behavior of the estimators. Specifically, when more than 10% of the distribution




Figure 3.30 – Results of the quadratic subtraction for an estimator (ϵ) between foam and no-foam regions, in the case of the ALLCOMP sample, measured at 1 GeV. This provides the contribution of the foam sandwich alone (foam + carbon fleece + glue), excluding the contribution from the two ALPIDE sensors which are present in both the foam and no-foam cases.

is removed, all estimators systematically start decreasing. Thus, to preserve the integrity of the analysis, a limit is placed at the 90% quantile, and the 80% and 85% quantiles are excluded from further calculations. The remaining values, along with their associated uncertainties, will be used in the next section.

Up to this point, the primary factors influencing the results and contributing to systematic uncertainties have been discussed. The largest sources of uncertainty are the momentum spread due to beam generation (not corrected, but accounted for; bringing an uncertainty $\mathcal{O}(20 \mu\text{rad})$) and the quadratic subtraction, which does not fully cancel out the telescope effects ($\mathcal{O}(20 \mu\text{rad}$ to $30 \mu\text{rad}$), depending on the estimator and quantile). Other effects that are understood have been accounted for and corrected, such as the non-linearity in the quadratic subtraction using the Highland formula, which introduces an effect at the percent level.

3.3.7 Material budget calibration

Finally, with better control over the uncertainties that can affect the quadratic subtraction and individual measurements, the overall scaling factor ς discussed at the beginning of the subchapter can be calculated.

For a chosen estimator and quantile, the Ni and Al data points at all three beam momenta, which contain the measured width of the kink angle distribution (for this specific estimator and quantile pair), corrected for non-linearities in the Highland formula, are displayed on a scatter plot. This plot relates the θ_0 angle to the x/X_0 expected material budget. The corresponding x/X_0 entry represents the resulting material budget after subtraction. For example, in the Ni case mentioned earlier, performing a quadratic subtraction between the measurement in the region of 165 μm and the one of 25 μm should correspond to a thickness of Ni of 140 μm . A scaled Highland formula is then fit to these points. The scaled Highland formula includes the ς term and is expressed as:

$$\theta_{0,\text{scaled}} = \varsigma \cdot \frac{13.6 \text{ MeV}}{\beta c p} z \sqrt{\frac{x}{X_0}} \left[1 + 0.038 \ln \left(\frac{x z^2}{X_0 \beta^2} \right) \right] \quad (3.10)$$

The fit returns the best ς value that minimizes the χ^2/ndf . The goodness-of-fit values range from 0.7 to 2.5. Smaller values are usually associated with the 5.4 GeV curve fit due to the fact that the Highland parametrization becomes rather flat at high energies and the errors are overestimated.

These results are shown in Figure 3.31, exemplarily for the $\sigma_{\text{Gauss},98}$ estimator. A few other estimator and quantile pairs are documented in Appendix E. The scaling factor for all estimators and quantiles ranges between $0.9 \leq \varsigma \leq 1.4$. Specific quantiles from unstable estimators (e.g., RMS, AAD for quantiles 100 and 99) or estimators for quantiles that cut too much into the distribution (e.g., all estimators for quantile 80, or AAD and RMS for quantile 85) can yield values outside this range.

The data points include the systematic uncertainties related to the θ_0 angle, as discussed at the end of the previous section. The uncertainties on the horizontal axis for the values related to the two foams represent the uncertainty in the material composition of the two samples (see Table 3.4).

The plot shows both the original, unscaled Highland curves in full color, as well as the scaled-up versions for each momentum, as dotted lines. The data from the ALLCOMP and ERG foams are also plotted in green shades at the x/X_0 position calculated from first principles, considering the sample composition (see Table 3.4).

The ratio of data to original Highland fit shows that all data points are incompatible with the Highland expectation. However, after scaling up, the values cluster more closely and are compatible with the Highland expectation value within 10%.




Figure 3.31 – Calibration of the data using a scaled Highland parametrization with the $\sigma_{\text{Gauss},98}$ estimator and quantile pair. The kink angle distribution widths for nickel and aluminum, after quadratic subtraction, provide the scaling parameter at the three measured momenta. The scaling parameter ς is then applied to the ALLCOMP and ERG foam sample data points to determine the material budget.

Once the scaling factor is determined, it is applied to the measured widths (for this specific estimator and quantile pair) of the ERG and ALLCOMP foam samples. With the corrected θ_0 angles, the material budget is calculated using the Highland formula.

Since the equation cannot be solved analytically, a numerical method is employed. The `ROOT_SCALAR` function from the `SCIPY.OPTIMIZE` module is used. The Brent method is chosen for the entire data interval, as initial bracketing [116].

After repeating the procedure for all the other estimator and quantile pairs, the final results are shown in Fig. 3.32. The blue points represent the results for the ALLCOMP sample, while the green points represent the results for the ERG sample. For each quantile and estimator type, a momentum-averaged point is calculated. The separate quantiles and estimators are shown in this figure in order to assess their stability.




Figure 3.32 – Combined results for the ALLCOMP (blue points) and ERG (green points) foam samples after applying the ς scaling factor obtained from the Ni/Al data points to the measured widths of the kink angle distributions for each specific estimator and quantile pair. For each point a momentum-average value is calculated. The points are artificially slightly shifted around the respective quantile value for better visibility.

It can be seen that for both samples, the values for the different estimators and quantiles cluster around a true x/X_0 value. As previously discussed, the RMS and AAD are highly unstable, especially for the largest quantiles, resulting in values far from the expectation and with large errors.

As a result, these estimators are marked as outliers and excluded from the final calculation. As previously mentioned prior to Section 3.3.7, the most restrictive quantiles studied (85% and 80%) involve removing 15-20% of the entire distribution. Given that the tails of the scattering angle distribution only account for a few percent, removing too much of the distribution can also affect some of the estimators (see Fig. 3.30). Therefore, these quantiles were not considered.

3.3.8 Material budget measurement for the foam samples

At this point, two methods for estimating the final x/X_0 value for the two foam samples from the results of Fig. 3.32 are discussed. One method involves selecting a robust estimator for a specific quantile and discussing the choice made in this case. The second method involves calculating a combined estimator per quantile from all selected estimators and then averaging these values across all selected quantiles.

Method 1 - Estimator selection

Managing the large number of potentially competing estimators for the same quantity has been extensively researched in specialized literature over the past years (see for example [117–119]).

Estimating a parameter θ_{k_x+y} in advance from a statistical model (which in this case would be the width of the scattering angle distribution), based on a collection of available estimators ϵ_i (like RMS, AAD, MAD, σ_{Gauss}) is difficult due to the relative performance of each estimator being influenced by the size of the sample, the true parameter value and other possibly unknown factors.

In the end, if care is taken to choose a robust and stable estimator, the choice does not significantly influence the result or the measurement. Most values of the robust estimators revolve around an expected true value. From the results of Fig. 3.32 MAD is the most stable of the studied estimators across the full quantile range, and is therefore chosen for this first method. To be compatible with the multiple scattering description, the 98% quantile is selected.

Choosing the MAD yields a material budget contribution for the ALLCOMP foam sandwich (not containing the two ALPIDE sensors) of $(0.369 \pm 0.031_{\text{syst}} \pm 0.001_{\text{stat}}) \%$ and for the ERG foam sample of $(0.166 \pm 0.026_{\text{syst}} \pm 0.001_{\text{stat}}) \%$, where the systematic uncertainties are the ones discussed prior to Section 3.3.7.

However, a conscious choice is made to select from any of the available estimators and quantiles, excluding those far from the true average or previously identified as outliers.

Therefore, to account for this choice, an uncertainty is calculated based on the variation in estimator values (excluding outliers) across the quantiles.

The mean (simple average) material budget value across all estimators and quantiles for both foam types is 0.368 % for ALLCOMP and 0.168 % for ERG. An uncertainty due to estimator variation is determined by calculating an interval of one standard deviation around this mean, capturing 67% of the data points. The resulting uncertainty is 0.030 % for the ALLCOMP sandwich and 0.014 % for the ERG sample.

Taking this into consideration, an estimator-(and quantile)-agnostic material budget value for the ALLCOMP foam sample is $(0.368 \pm 0.031_{\text{syst}} \pm 0.030_{\text{estim}} \pm 0.001_{\text{stat}})$ %, while for the ERG sample it is $(0.168 \pm 0.026_{\text{syst}} \pm 0.014_{\text{estim}} \pm 0.001_{\text{stat}})$ %. As shown, a significant portion of the uncertainty arises from the choice of estimator.

Method 2 - Combining estimators

Combining different estimators to obtain a more accurate estimate of a parameter is a well-established technique in statistical analysis. The idea is to leverage the strengths of various estimators, which may each perform well under different conditions, to produce a single, more reliable estimate. This process involves determining optimal weights for each estimator and using these weights to calculate a weighted average. This would better represent the mean material budget across the different quantiles and estimators than the simple average, as it takes the associated uncertainties of each estimator and quantile. Moreover, it is expected that this method reduces the uncertainty on the estimator choice. The method chosen in this work to average the estimators per quantile closely follows the approaches described in [119, 120].

If ϵ is a collection of estimators of the real parameter θ_{k_x+y} (which in this case is the width of the kink angle distribution), a combined estimator can be constructed using linear transformations:

$$\hat{\theta}_\lambda = \lambda^\top \epsilon$$

where λ^\top denotes the transpose of λ , a vector of weights which must sum up to one $\left(\sum_{i=1}^k \lambda_i = 1\right)$ to ensure the combined estimator remains unbiased.

The objective is to find a weighted average of these estimators that minimizes the mean squared error (MSE).

The optimal weights are derived by minimizing the MSE of the combined estimator. This involves estimating the MSE matrix Σ , which captures the variances (diagonal elements) and covariances (off-diagonal elements) of the estimators.

$$(\Sigma)_{ij} = \langle \sigma_i \sigma_j \rangle$$

Here, $\sigma_{i,j}$ are the uncertainties of the corresponding estimator pairs. The matrix takes into account full correlation between all quantiles and estimators for a specific foam type, since they belong to the same kink angle distribution.

The optimal weights λ^* are obtained by solving the following optimization problem:

$$\lambda^* = \arg \min \mathbb{E} \left(\lambda^\top \mathbf{e} - \theta_{k_x+y} \right)^2 = \arg \min_{\lambda} \lambda^\top \Sigma \lambda$$

using the constraint that the sum of the weights must be equal to one.

The solution is given by:

$$\lambda^* = \frac{\Sigma^{-1} \mathbf{1}}{\mathbf{1}^\top \Sigma^{-1} \mathbf{1}}$$

Here $\mathbf{1}$ is a vector of ones and the MSE matrix is inverted. By doing this, the chances of deteriorating the estimation of θ_{k_x+y} are minimized due to the smoothing effect of averaging. Although the MSE might be more complex and difficult to estimate than θ_{k_x+y} , the overall process of combining estimators is relatively robust to inaccuracies in the estimation of the MSE matrix.

When calculating the optimal weights for the combined estimator, small errors in the estimation of Σ do not significantly affect the resulting weights. This ensures that the combined estimator remains effective even if Σ might not be estimated perfectly. Moreover, errors in the individual estimators propagate through the estimation of Σ . However, because the combined estimator is a weighted average, the impact of these propagated errors is mitigated. This mitigation occurs because the weights are designed to minimize the overall MSE, balancing out the individual errors.

By combining multiple estimators in this way, their individual errors tend to cancel out to some extent, leading to a more accurate overall estimate.

A similar procedure of minimization is applied to also combine the quantiles. The final material budget results obtained using this method are $(0.365 \pm 0.031_{\text{syst}} \pm 0.011_{\text{estim}} \pm 0.001_{\text{stat}}) \%$ for the ALLCOMP sample and $(0.164 \pm 0.026_{\text{syst}} \pm 0.008_{\text{estim}} \pm 0.001_{\text{stat}}) \%$ for the ERG sample.

By combining the different estimators, the uncertainty associated with the estimator choice was reduced by a factor of about two to three. While the average value of the material budget changed ever so slightly, this change is smaller than other uncertainties.

3.4 Discussion

ALICE is set to push the frontiers of pixel vertexing detectors with an innovative concept utilizing wafer-scale silicon sensors that are bent. The upcoming ITS3 upgrade will feature an almost massless detector, achieving a mean material budget of $0.07\% X_0$ per layer. This will be accomplished by employing $50\text{ }\mu\text{m}$ thin sensors supported by open-cell carbon foam structures and cooled by forced airflow. This design eliminates the need for traditional staves, electrical connections, and cooling pipes, thus removing support structures and their associated material budget from the active area.

Given that the carbon foam will be the sole structural element within the active area of the ITS3, apart from the silicon sensors themselves, it is crucial to understand the amount of scattering it introduces. This knowledge will enhance tracking accuracy in the high multiplicity environment of ALICE and refine the dedicated simulations.

In a dedicated testbeam campaign at the DESY II facility, two carbon foam samples were evaluated: an ERG (RVC) Duocel foam, selected for the ITS3 active area due to its low density, and a denser ALLCOMP K9 SD foam, intended for the periphery of the final detector for its enhanced thermal conductivity, which is essential for efficient cooling. Both foam samples, each glued between two ALPIDE sensors, were investigated within a beam telescope to assess their material budget and scattering behavior.

This chapter introduced the concept of material budget imaging, a technique that relies on accurately measuring multiple scattering angles from reconstructed particle tracks in a reference telescope. This technique has been used in recent years by various groups [102–104], and the work presented here closely follows these results. This thesis aims to provide deeper insights into some of the considerations made in previous works (e.g., the choice of estimators and various correction factors to get a better systematic uncertainty) and offers a more detailed study of some of the effects.

Interactions between particles and the material sample cause deflections of the original track, resulting in a scattering angle distribution described by Molière theory. This distribution contains an inner Gaussian core and long tails from large-angle scatterings. The Highland parameterization can then be used to infer the radiation length from the width of such a distribution. The Highland description is said to be accurate to 11% or better for a range of x/X_0 between 10^{-3} and 100. Although corrections to the Highland formula or other better descriptions exist for the scattering angle distribution, they remove the dependence on the radiation length. Since this link is at the core of this study, no attempt is made to investigate other, potentially more accurate, formulae.

Various estimators of the width of this distribution were studied to assess their stability and robustness. Additionally, quantiles were computed where up to 20% of the tails of the kink angle distribution were excluded. This allows for a check of the robustness of the estimators against tail outliers, as the goal is to characterize the inner core created by multiple scattering, for which the Highland formula is pertinent.

The study showed that there are various contributions to the measured kink angle distribution. The distribution is primarily a convolution of the contribution from the Highland formula and a part coming from the angular resolution of the telescope and associated effects. To disentangle these effects and determine the material budget of the samples, various corrections and calibrations are necessary.

This work provides a method to deduce the material budget of unknown low material budget samples from the measured angular distribution. It also offers various insights into setup considerations, calibration samples, width estimator choices, and achievable precision.

To achieve this, systematic studies across various momenta were conducted to analyze the width of the kink angle distribution across the sensor surface. By strategically selecting regions of interest and applying quadratic subtraction, the method effectively reduces some of the systematic uncertainties.

Nonetheless, it was shown that quadratic subtraction is only an approximation and does not fully cancel the uncertainties. Due to the non-linear nature of the Highland formula, performing a quadratic subtraction between two materials of different thicknesses does not yield the material budget of the difference in thicknesses between the two cases. Correction factors are needed to account for this non-linearity. These factors are provided for a range of thicknesses and materials, along with a functional form from which other thicknesses and materials can be deduced. These corrections are most significant in the case of quadratic subtraction with samples that are very close in thickness and can be as high as 20%.

Several checks were performed with three different datasets where the telescope was “empty” (i.e., no scatterer inserted) or where two different regions of the same calibration run had the same amount of material present. This provided a good understanding of any systematic uncertainties present in the measurement after quadratic subtraction.

Using a set of nickel and aluminum targets of various known thicknesses, a calibration is performed between the expected Highland values and the measured ones. This calibration factor is then applied to the measured foam sample values, correcting the width of the kink angle distributions measured with various estimators and quantiles. Finally, this corrected angle is translated into a material budget value, per foam type and per estimator and quantile.

In this study, it is not possible to determine the material budget of individual components of the carbon foam samples. However, the contribution from the sample as a whole can be assessed, excluding the two ALPIDE sensors that sandwich it, whose influence is removed during the quadratic subtraction process.

The obtained values are then discussed. It was observed that some estimators are not robust, especially when considering the full distribution. Moreover, quantiles which cut into the kink angle distributions too much are rejected. After this initial check, a final material budget quantity for the two samples can be derived in two different ways.

One method involves selecting an estimator and quantile from the list of available ones. It was shown that the choice is arbitrary, as long as the estimator at this particular quantile is robust and stable. The uncertainty in this biased choice was estimated from the average value of the other estimators over the full range of quantiles (neglecting the outliers). An estimator-agnostic value of the material budget contribution of the ALLCOMP sample is measured at $(0.368 \pm 0.031_{\text{syst}} \pm 0.030_{\text{estim}} \pm 0.001_{\text{stat}}) \%$ and of the ERG at $(0.168 \pm 0.026_{\text{syst}} \pm 0.014_{\text{estim}} \pm 0.001_{\text{stat}}) \%$. The relative error in this case is 10-15%.

Another method involves combining estimators to obtain a single, more accurate estimate of the width of the kink angle distribution by leveraging the combined strength of various width estimators. A weighted average of the estimators is chosen, which minimizes the mean square error matrix, following closely the results in [119, 120]. Furthermore, after averaging the estimators for each quantile category, the quantiles themselves can be combined in a similar manner. For the two foam samples (modulo the two ALPIDE sensors), the final material budget values are $(0.365 \pm 0.031_{\text{syst}} \pm 0.011_{\text{estim}} \pm 0.001_{\text{stat}}) \%$ for the ALLCOMP sample and $(0.164 \pm 0.026_{\text{syst}} \pm 0.008_{\text{estim}} \pm 0.001_{\text{stat}}) \%$ for the ERG sample. The uncertainty due to the choice of estimator has been reduced.

The combined systematic uncertainty for both samples gives the precision of the method which shows the smallest amount of material that can be reliably measured. This would coincide to a single 50 μm -thick silicon ALPIDE sensor ($\sim 0.05 \% X_0$).

These values are compared to the expectations of the material contribution calculated from first principles (knowledge on the components that make up the samples), which amount to $0.414 \pm 0.042\% X_0$ for the ALLCOMP sample and $0.172 \pm 0.011\% X_0$ for the ERG sample.

It can be noted that the measurement precision is better than the material calculation, which has uncertainties primarily due to the amount of glue present (either as a uniform layer or that has seeped into the foam).

Experimental data for the ERG foam sandwich with either of the two methods is fully compatible with material predictions. For the ALLCOMP sample, the measurement values from either method are fully compatible with each other, but differ slightly from the theoretical expectation. The measurements show that the ALLCOMP samples contains less material than predicted by the material calculation done in this work. Depending on the method chosen, this discrepancy is of the order of $0.05\% X_0$. Incidentally, considering glue as the extra material, this would correspond to about $200\text{ }\mu\text{m}$. This indicates that either the amount of glue at the sensor interface was less, or(and) that the amount of glue that seeps into the foam due to capillarity forces is reduced.

This work is continued, under the guidance of the author, by a master student [115] with a set of more realistic samples. The experience and know-how gathered by the author and presented in this work is used to improve the understanding of these new measurements.

These new samples were produced with better control over the glue and thickness of materials and include final grade components (i.e. sensors produced in the 65 nm CMOS technology node with thicker metal layers). Due to their improved production process, they allow the measurement of the direct contribution and radiation length of the different components individually.

The sensors are further spaced apart in the new telescope and this enables a twofold improvement in the angular resolution, allowing for measurements of smaller angles. Additionally, more statistics were available, enhancing the robustness of the results. Various corrections not included in this work are addressed.

The momentum gradient is thoroughly investigated, a two-dimensional fit is considered across the sensor surface, and a bin-by-bin analysis of the material budget is performed. Furthermore, a proper deconvolution technique is attempted in order to better separate the contribution of the telescope from the measurements. The entire analysis is also conducted separately for the x and y kink angles to verify if the assumption made in this work about the combined kink angle holds true for pixel sensors with asymmetric pitches.

Preliminary results confirm this approach. These findings demonstrate the feasibility of the imaging technique to resolve material budgets of unknown samples. The results even provide a prediction on the possible amount of glue difference between the samples and expectations from material analysis. This will be later verified by a dedicated CT scan.

This method serves as a powerful tool in assessing the influence of new materials on the tracking efficiency of the sensors. In order to successfully implement a low-mass tracker and vertexer, understanding aspects related to the material budget distribution in the active area is mandatory.

Short summary and outlook

The Inner Tracking System 3 (ITS3) is a cutting-edge detector development within the ALICE experiment at CERN, planned to be installed during the Long Shutdown 3 (2026-2028), aimed at pushing the boundaries of vertex detection. ITS3 focuses on improving tracking precision and material budget, especially in environments with high multiplicity. This detector will replace the three inner layers of the current ITS2 system, retaining the other four outer layers, and will incorporate curved, ultra-thin, wafer-scale silicon sensors (up to $26\text{ cm} \times 10\text{ cm}$), making it the first of its kind among other physics experiments.

To access such large silicon sensors, a switch to the 65 nm CMOS technology node was necessary, replacing the 180 nm node used in the ITS2 sensors (ALPIDE). This new node allows for larger wafers (300 mm), increased integration density, and reduced power consumption.

Positioned just 19 mm from the collision point, ITS3 will feature silicon sensors as the primary material in the active area, supported by thin, ultra-low material budget carbon foam structures spanning a small region across the length of the detector. The entire assembly will be cooled by forced airflow, resulting in an expected average material budget of only $0.07\% X_0$ per layer.

The new detector will significantly enhance the tracking of low-momentum particles and improve measurements in heavy-flavor physics, among other scientific goals. The ITS3 development includes several key research and development directions, with the author contributing significantly to two: bent sensor characterization and feasibility studies and material budget assessment of the carbon foam support structures. These studies are essential for validating the use of curved sensors in ALICE and beyond.

In the bent sensor characterization, ALPIDE chips, which share many characteristics with the ITS3 final sensors, were tested both before and after bending. The results confirmed that these sensors maintained full functionality after bending, with consistent performance across metrics such as threshold distribution, detection efficiency, and position resolution. These findings represent a critical milestone for the project, proving that bent sensors perform comparably to their flat counterparts. Notably, bent ALPIDE sensors achieved over 99.9%

efficiency at thresholds around 100 electrons, a result on par with, and sometimes better than, the flat sensors due to the ability of particles to strike the chips at non-zero incident angles, resulting in increased charge deposition.

On the material budget analysis front, the author focused on evaluating the impact of two open-cell carbon foams intended to support the thin bent sensors: Carbon (RVC) Duocel used in the active area of the ITS3, and Allcomp K9 standard density used at the periphery of the sensors, outside the active area, as a thermal radiator. The goal was to assess the scattering behavior of these materials through a rather new technique called material budget imaging. Particle beams from an accelerator facility that impinge on the foam samples will scatter at an angle with respect to their original direction of motion. This angle can be ultimately related to the material budget of the samples after some calibration steps. Through this advanced method, the author helped determine the material contribution of the two foam samples and provided new insights into the systematic uncertainties that can creep into such a measurement. Ongoing studies with newer samples and sensors are being pursued by a student under the author's guidance.

Both projects have yielded promising results, demonstrating that curved sensors can perform well in the demanding conditions of the ALICE experiment and that the new materials introduce minimal additional scattering. These results play a key role in the design and successful deployment of the ITS3 in the coming years.

Looking ahead, current studies are focusing on bending 65 nm sensors to validate the promising results already obtained with the ALPIDE sensors. These investigations involve a range of sensors, including small test structures (as small as 1 mm × 1 mm), full reticle sensors (baby-MOSS), and stitched sensors (MOSS). Preliminary findings have shown that bending the 65 nm test structures, such as the APTS and DPTS, does not significantly impact their performance, confirming the robustness of this new technology under bending conditions.

The interest in bending sensors has grown considerably, with other experiments exploring similar detection systems, such as the R³B experiment at FAIR [121] and the vertex detector for the Electron-Ion Collider (EIC) [122, 123]. Additionally, ALICE 3 — the next-generation experiment set to replace ALICE after Long Shutdown 4 — plans to incorporate a vertex detector with bent silicon layers operating in a secondary vacuum [60]. This imposes stringent design requirements, ensuring that bending studies will continue to play a pivotal role in the development of future detectors.

On the material budget front, the success of the material budget imaging technique and the corrections applied to address systematic uncertainties and setup limitations have resulted in

a robust method capable of characterizing any sample with high precision. This technique will prove particularly useful for future projects like ALICE 3, which will also face strict material budget constraints. For instance, assembled modules containing sensors, electronic circuits, and cooling elements can be quickly assessed, producing a 2D map of material hot-spots. Such insights will aid in material selection and design choices, ensuring that the detector meets its performance requirements.

The world of monolithic active pixel sensors is at a pivotal moment where a paradigm shift in vertex detector design is becoming possible. With the success of bent wafer-scale sensors, many exciting opportunities are emerging for future studies. These developments open the door to innovative detector concepts that will shape the next generation of heavy-ion and high-energy physics experiments.

Appendix A

Image cell size

During a three-day testbeam training session the author had the pleasure to show how material budget imaging works to a group of students from the graduate school. For this purpose, the EUDET-type telescope at the DESY II beamline 24 was used, which was equipped with MIMOSA26 monolithic active pixel sensors. The beam consisted of electrons with an energy of 3 GeV.

The MIMOSA26 sensors have a pitch size of $18.4\text{ }\mu\text{m} \times 18.4\text{ }\mu\text{m}$, with a matrix of 1152 columns and 576 rows, resulting in an active area of approximately $21.2\text{ mm} \times 10.6\text{ mm}$. Each sensor is $50\text{ }\mu\text{m}$ thick and is shielded from ambient light by $25\text{ }\mu\text{m}$ thick Kapton foil on both sides [124].

The telescope consists of three upstream and three downstream MIMOSA26 sensors, with a 3D-printed scatterer, shown in Fig. A.1, placed in the middle.

The targets measure approximately $1\text{ mm} \times 1\text{ mm}$, are 4 mm thick, and are made of PLA. They feature small openings on the inside part of the hazard sign that are $300\text{ }\mu\text{m}$ thin, as well as grooves on the OK sign that are $200\text{ }\mu\text{m}$ thick. Some small embellishments are also visible on the top side of the samples coming from the printing technique and of a similar size as the nozzle of the 3D printer. These features are limited by the $100\text{ }\mu\text{m}$ resolution of the SLA 3D printer used.

The goal of the exercise was to demonstrate how material budget imaging works, so no additional effort was made to investigate different estimators and quantiles. The analysis plane was divided into small cells of various sizes, and the RMS of the full distribution per cell size was measured. For each of the two samples, the material budget images can be seen in the following figures.

Although the $100\text{ }\mu\text{m} \times 100\text{ }\mu\text{m}$ cell size plot is sufficient to reveal features such as the small slots in the hazard sign, other details are more difficult to discern. With sufficient statistics, smaller cell sizes reveal finer details, like the intricate design visible in the macro pictures. However, at the $10\text{ }\mu\text{m} \times 10\text{ }\mu\text{m}$ cell size not enough statistics are available and some cells are empty.




Figure A.1 – (Above) 3D render of the two scatterers used in the material budget imaging demonstration. The 3D designs were taken from [3d-scatterer]. (Below) Pictures of the 3D printed scatterers showing some finer detail.




Figure A.2 – Material budget image of the hazard sign target using 100 μm × 100 μm cell size. The small slots in the sign are visible, but finer details are harder to discern. A total of 19k cells make up the picture.




Figure A.3 – Material budget image of the hazard sign target using $50 \mu\text{m} \times 50 \mu\text{m}$ cell size.




Figure A.4 – Material budget image of the hazard sign target using $25 \mu\text{m} \times 25 \mu\text{m}$ cell size.

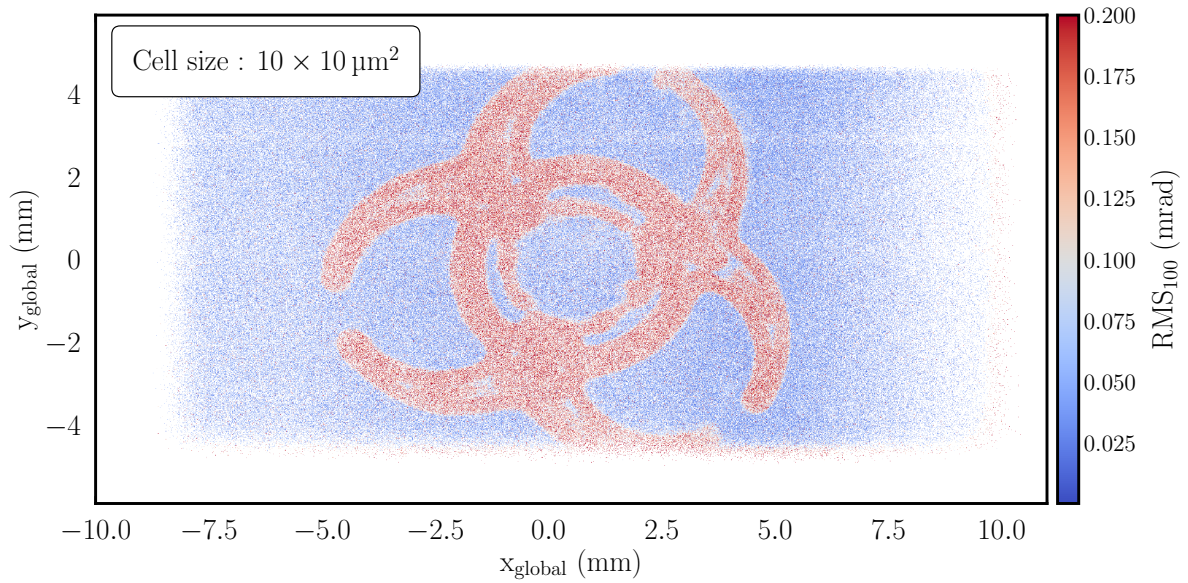


Figure A.5 – Material budget image of the hazard sign target using $10 \mu\text{m} \times 10 \mu\text{m}$ cell size. The smallest cell size shows some of the intricate details, though some cells are empty due to the limited statistics. A total of 1.9M cells are created.

A total of 45M events were recorded with this sample. For the investigated cell sizes, in decreasing order, the mean number of entries per cell were: 251, 132, 46, and 9.

For this scatterer, only 27M events were collected. In this case, the smallest cell size is inadequate for displaying the necessary details. Nevertheless, the fine details, such as the path taken by the 3D printer, can be seen in the $25 \mu\text{m} \times 25 \mu\text{m}$ cell size plot.




Figure A.6 – Material budget image of the OK sign target using $100 \mu\text{m} \times 100 \mu\text{m}$ cell size.




Figure A.7 – Material budget image of the OK sign target using $50 \mu\text{m} \times 50 \mu\text{m}$ cell size.




Figure A.8 – Material budget image of the OK sign target using $25 \mu\text{m} \times 25 \mu\text{m}$ cell size.




Figure A.9 – Material budget image of the OK sign target using $10 \mu\text{m} \times 10 \mu\text{m}$ cell size.

Appendix B

Telescope setup configurations




Figure B.1 – The simulated setup. Distances between planes are chosen such that a symmetry with respect to the middle of the setup exists.

The setup from Fig. B.1 was simulated using the resolution simulator from [66]. Electrons are passing through the setup at momenta of 1, 2.4 and 5.4 GeV. The tracking is done using the GBL algorithm and the scattering introduced by the air and the sensor planes is taken into account.

For simplicity, three distances are defined as follows:

- x – the distance between the first and second ALPIDE, as well as the distance between the last two sensors
- y – the distance between planes 1 and 2, as well as 3 and 4
- z – the distance between the two middle planes

The position resolution and the kink angle resolution is calculated in the middle of the setup, i.e. at a distance of $(x + y + z)/2$. The optimum setup configurations are expected to be symmetric with respect to this middle position, so the assumption of the distances from before is a very good approximation, which however dramatically reduces the simulation time.

The following figures show the results for the different combinations of these distances, as well as a zoomed in version, restricting to distances smaller than 10 cm.

B.1 Setup configuration for 1 GeV simulated data




Figure B.2 – Kink angle resolution




Figure B.3 – Kink angle resolution (zoomed)




Figure B.4 – Position resolution




Figure B.5 – Position resolution (zoomed)

B.2 Setup configuration for 2.4 GeV simulated data




Figure B.6 – Kink angle resolution.




Figure B.7 – Kink angle resolution (zoomed).




Figure B.8 – Position resolution.




Figure B.9 – Position resolution (zoomed).

B.3 Setup configuration for 5.4 GeV simulated data




Figure B.10 – Kink angle resolution.




Figure B.11 – Kink angle resolution (zoomed).

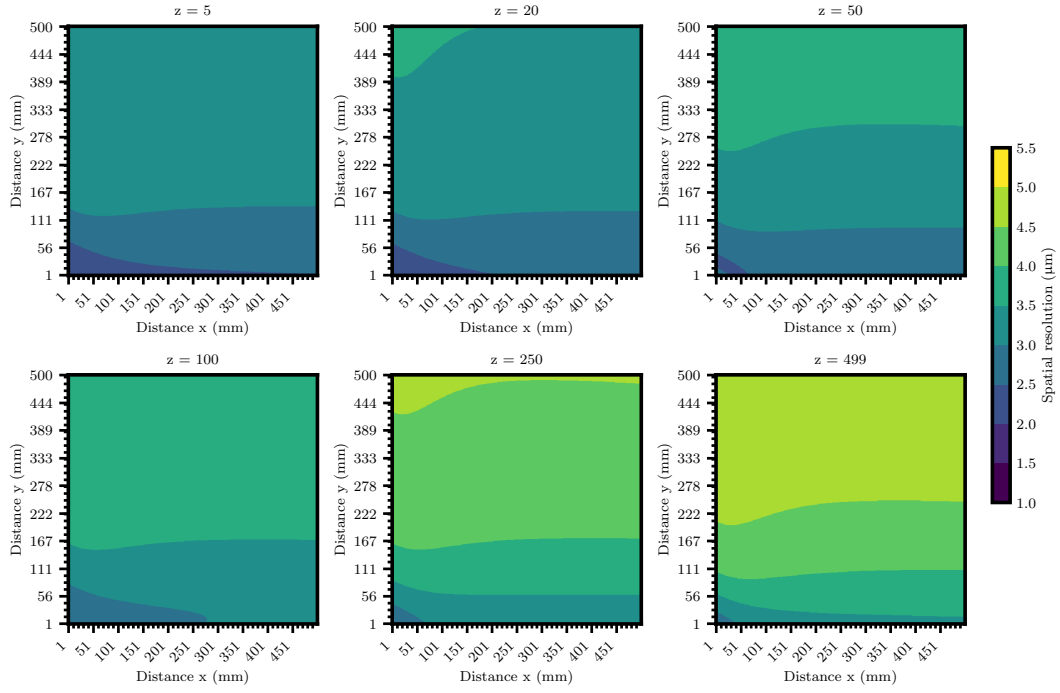


Figure B.12 – Position resolution.




Figure B.13 – Position resolution (zoomed).

Appendix C

Quadratic subtraction

Due to the non-linear nature of the Highland formula, the quadratic subtraction is not expected to fully cancel telescope effects.

To illustrate this, consider a material with a radiation length X_0 , from which two slabs of thickness $n \cdot L$ and L are constructed. When a beam of particles passes through these slabs, scattering occurs, causing the particles to deviate from their original trajectory. The resulting projected scattering angles for both slabs can be calculated using the Highland formula.

Given the Highland formula for multiple scattering:

$$\theta_0 = \frac{13.6 \text{ MeV}}{\beta p} \sqrt{\frac{x}{X_0}} \left(1 + 0.038 \ln \left(\frac{x}{X_0} \right) \right)$$

two angular widths can be calculated:

- $\theta_{0,L}$ – the scattering angle for thickness L ,
- $\theta_{0,nL}$ – the scattering angle for thickness nL .

Using the Highland formula, for the first slab:

$$\theta_{0,L} = \frac{13.6 \text{ MeV}}{\beta p} \sqrt{\frac{L}{X_0}} \left(1 + 0.038 \ln \left(\frac{L}{X_0} \right) \right)$$

For the thickness nL , the formula becomes:

$$\theta_{0,nL} = \frac{13.6 \text{ MeV}}{\beta p} \sqrt{\frac{nL}{X_0}} \left(1 + 0.038 \ln \left(\frac{nL}{X_0} \right) \right)$$

Using the property of logarithms, $\ln \left(\frac{nL}{X_0} \right) = \ln(n) + \ln \left(\frac{L}{X_0} \right)$:

$$\theta_{0,nL} = \frac{13.6 \text{ MeV}}{\beta p} \sqrt{\frac{nL}{X_0}} \left(1 + 0.038 \left(\ln(n) + \ln \left(\frac{L}{X_0} \right) \right) \right)$$

Let $A = 0.038 \ln \left(\frac{L}{X_0} \right)$, and $B = 0.038 \ln(n)$:

$$\theta_{0,nL} = \frac{13.6 \text{ MeV}}{\beta p} \sqrt{\frac{nL}{X_0}} (1 + A + B)$$

The quadratic difference between the two angular width measurements of the slabs is:

$$\theta_{\text{diff}} = \sqrt{\theta_{0,nL}^2 - \theta_{0,L}^2}$$

Substituting the expressions for $\theta_{0,nL}$ and $\theta_{0,L}$:

$$\theta_{\text{diff}} = \frac{13.6 \text{ MeV}}{\beta p} \sqrt{\frac{L}{X_0} [n(1 + A + B)^2 - (1 + A)^2]}$$

This has to be compared to the value predicted by the Highland formula for the material difference:

$$L_{\text{diff}} = nL - L = (n - 1)L$$

For this, the Highland formula is:

$$\theta_{0,(n-1)L} = \frac{13.6 \text{ MeV}}{\beta p} \sqrt{\frac{(n-1)L}{X_0}} \left(1 + 0.038 \ln \left(\frac{(n-1)L}{X_0} \right) \right)$$

The ratio of θ_{diff} to $\theta_{0,(n-1)L}$ gives the difference due to quadratic subtraction with respect to what the true value would be as given by the Highland formula for the respective L/X_0 :

$$\Delta = \frac{\theta_{\text{diff}}}{\theta_{0,(n-1)L}} = \frac{\sqrt{n(1 + A + B)^2 - (1 + A)^2}}{\sqrt{n-1}(1 + C)}$$

where $C = 0.038 \ln \left(\frac{(n-1)L}{X_0} \right) = 0.038 \ln(n-1) + A$ is chosen to make the formula simple.

Taking for example $n = 2, 3, 4$ this difference can be computed. It is shown in Fig. C.1 exemplary for the radiation length of Ni.

Extending this approach, one can compare cases for different slab thicknesses L_1 and L_2 for two materials. The resulting differences in scattering angles are illustrated in Fig. C.2.

A small difference is observed when changing the radiation length. For instance, switching from Ni to Al would result in approximately a 10% increase in the scattering angle values.

For the available Ni/Al data, which is subtracted in quadrature, the difference is calculated to be between 2-8% with respect to the Highland formula. This difference is accounted for in the analysis.




Figure C.1 – The deviation introduced by the quadratic subtraction method compared to the true value predicted by the Highland formula, plotted as a function of the radiation length per unit thickness for samples with lengths 2, 3, and 4 times greater than the original.




Figure C.2 – The deviation introduced by the quadratic subtraction method compared to the value predicted by the Highland formula for two nickel (Ni) slabs with different thicknesses, L_1 and L_2 .

Appendix D

Quadratic subtraction for the foam samples

The other results of the quadratic subtraction between the foam and no-foam (bottom) regions from Fig. [3.29](#) are shown below.

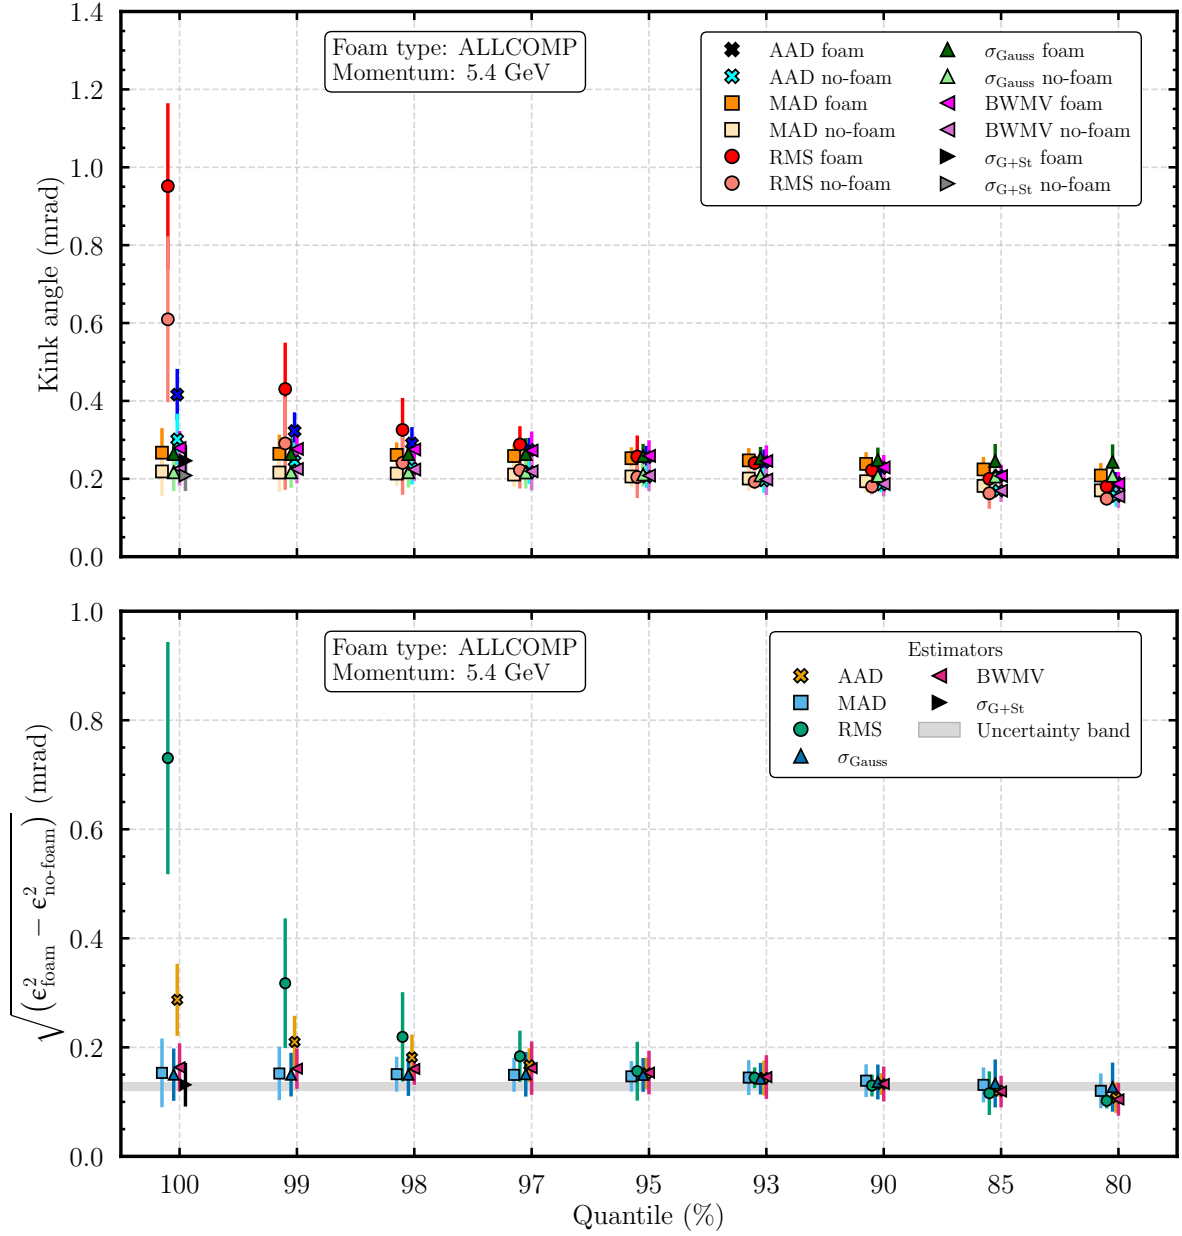


Figure D.1 – Quadratic subtraction (foam and no-foam regions) for the ALLCOMP sample at 5.4 GeV.

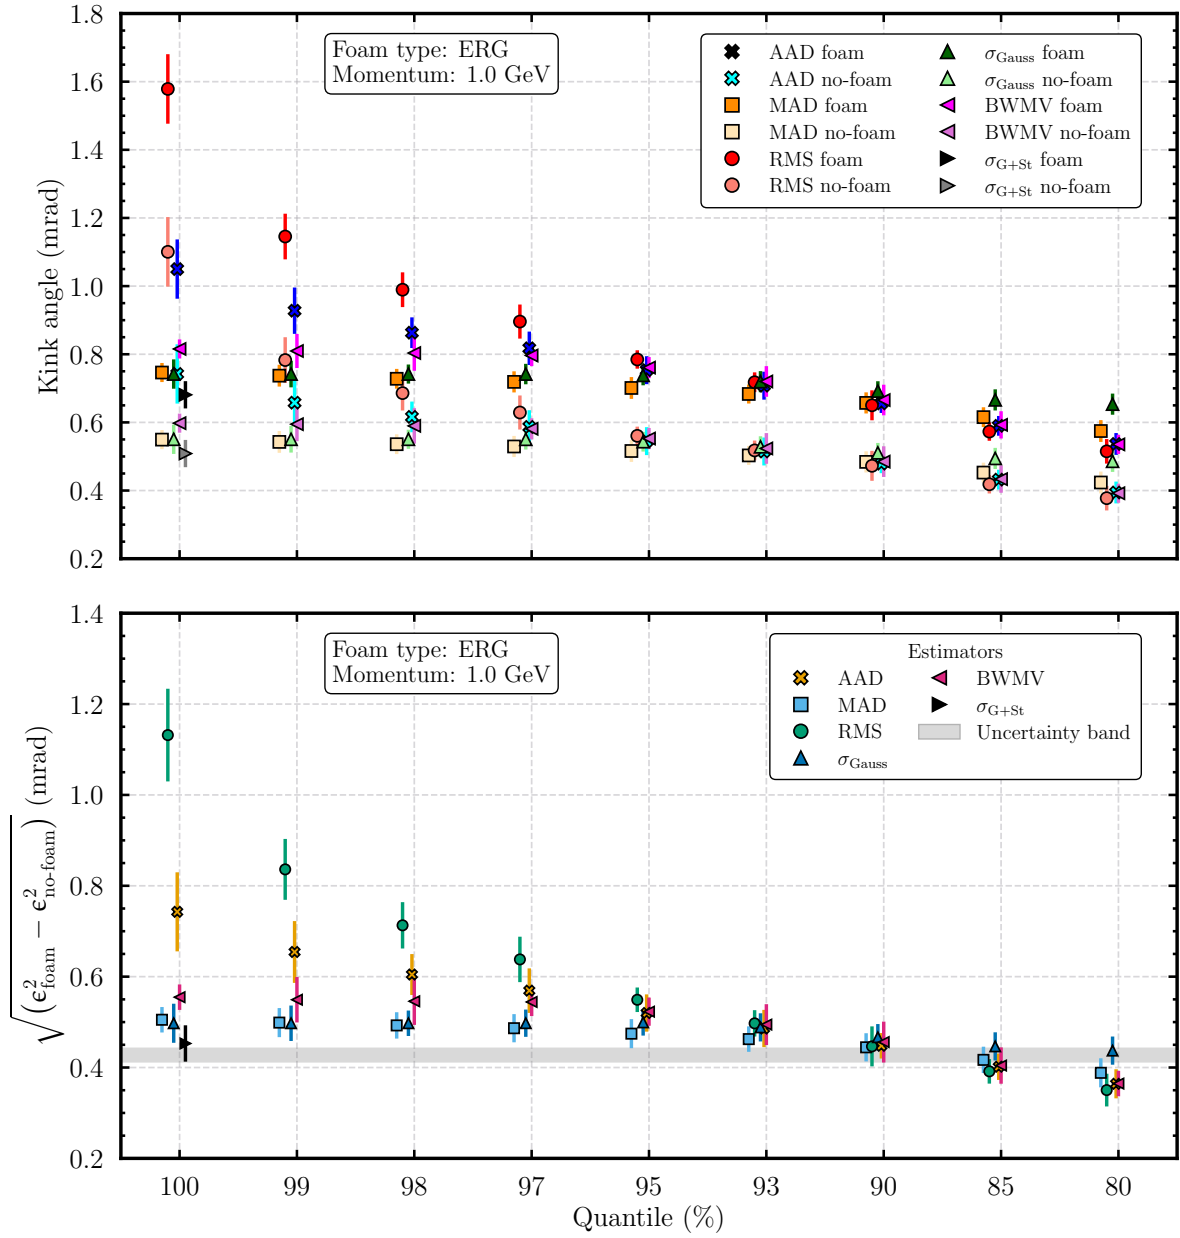


Figure D.2 – Quadratic subtraction (foam and no-foam regions) for the ERG sample at 1 GeV.

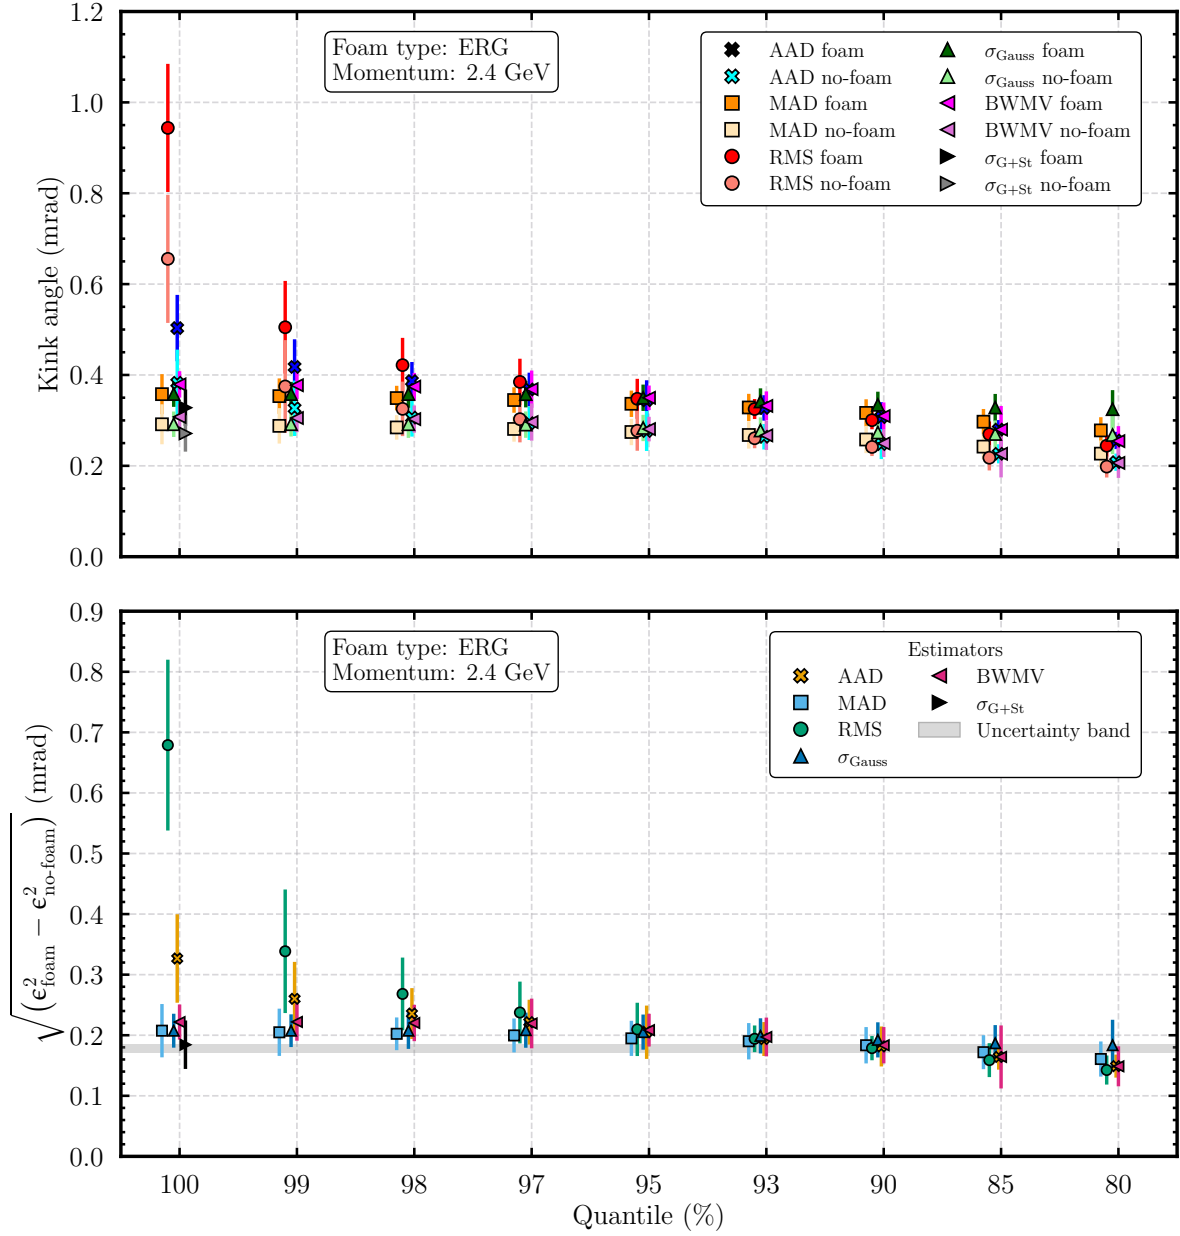


Figure D.3 – Quadratic subtraction (foam and no-foam regions) for the ERG sample at 2.4 GeV.

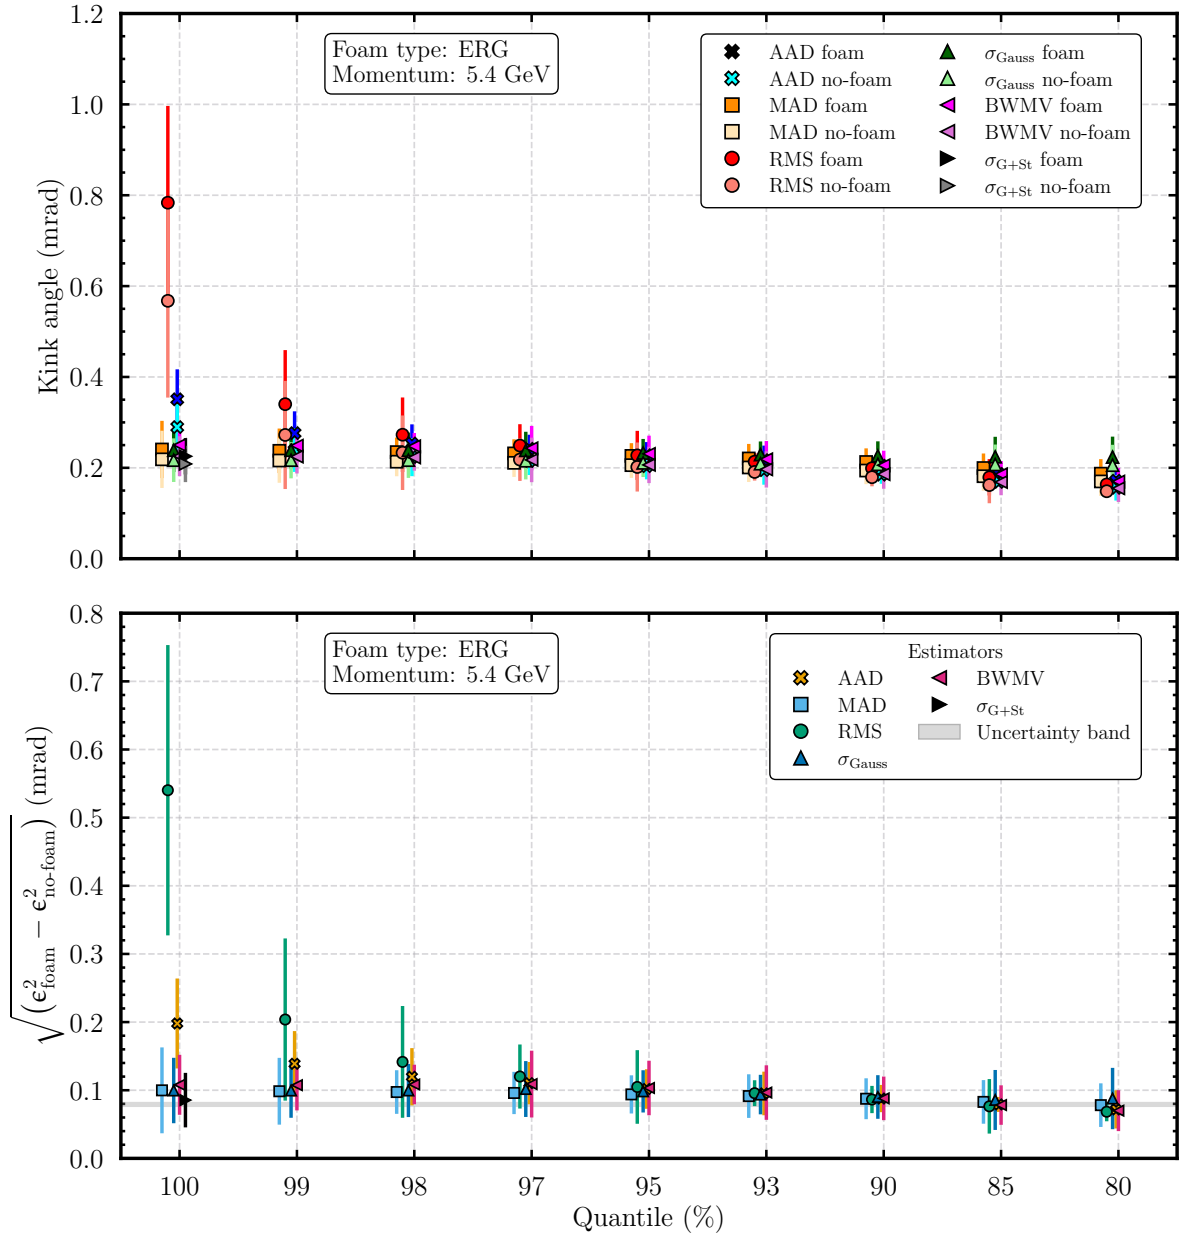


Figure D.4 – Quadratic subtraction (foam and no-foam regions) for the ERG sample at 5.4 GeV.

Appendix E

Material budget calibration

Several representative material budget calibration plots, with associated fits from which scaling factors are extracted are shown below.




Figure E.1 – Highland scaling using the AAD width estimator at the 90% quantile.




Figure E.2 – Highland scaling using the bi-weight midvariance (BWMV) width estimator at the 93% quantile.

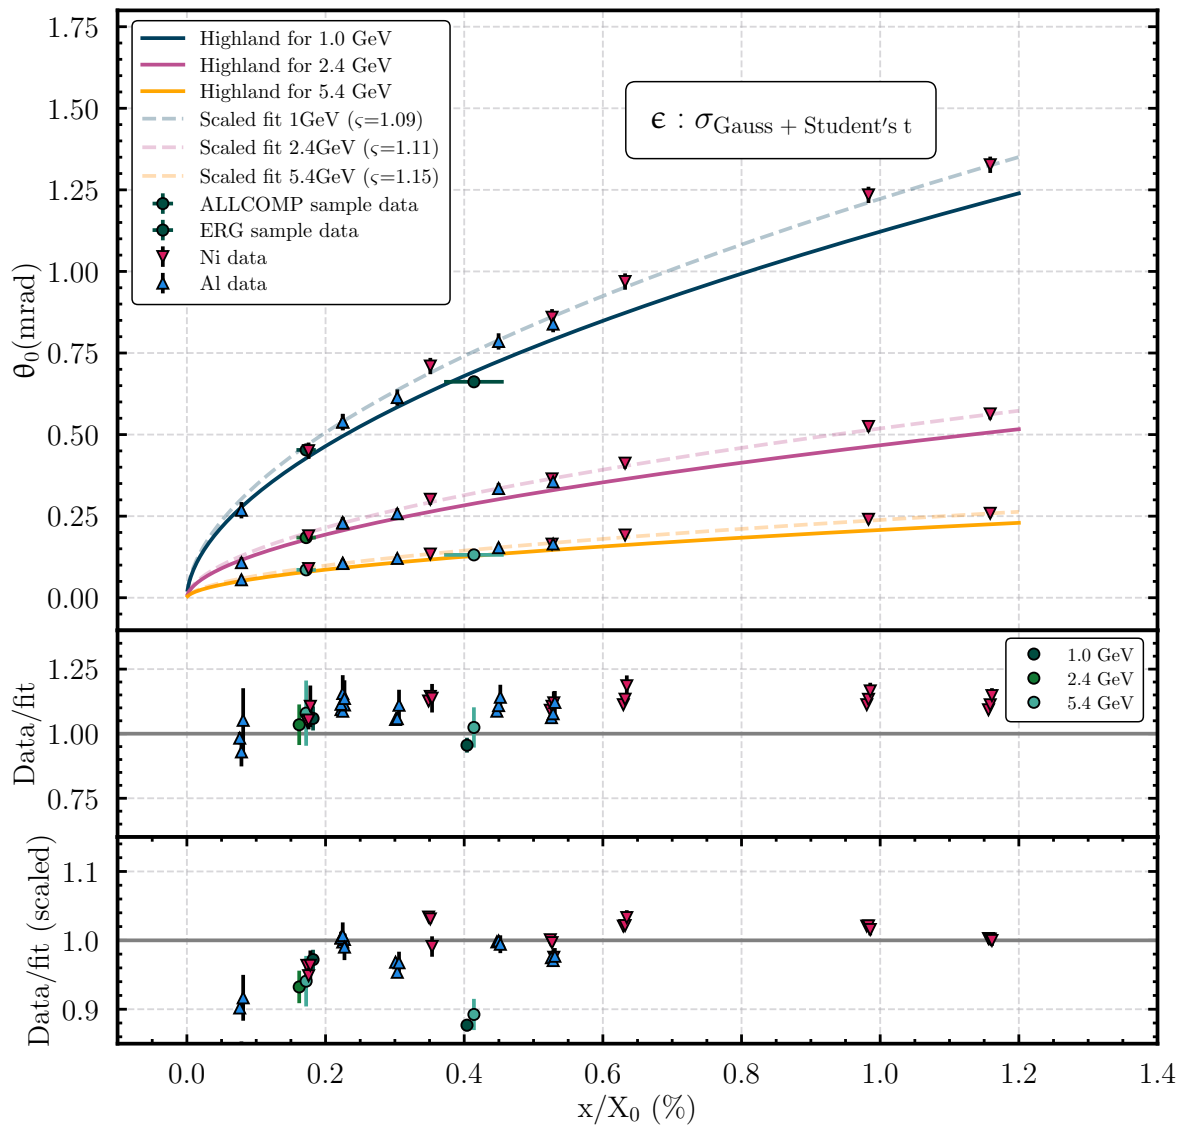


Figure E.3 – Highland scaling using the width estimator from the convolution of a Gaussian and a Student's t distribution.

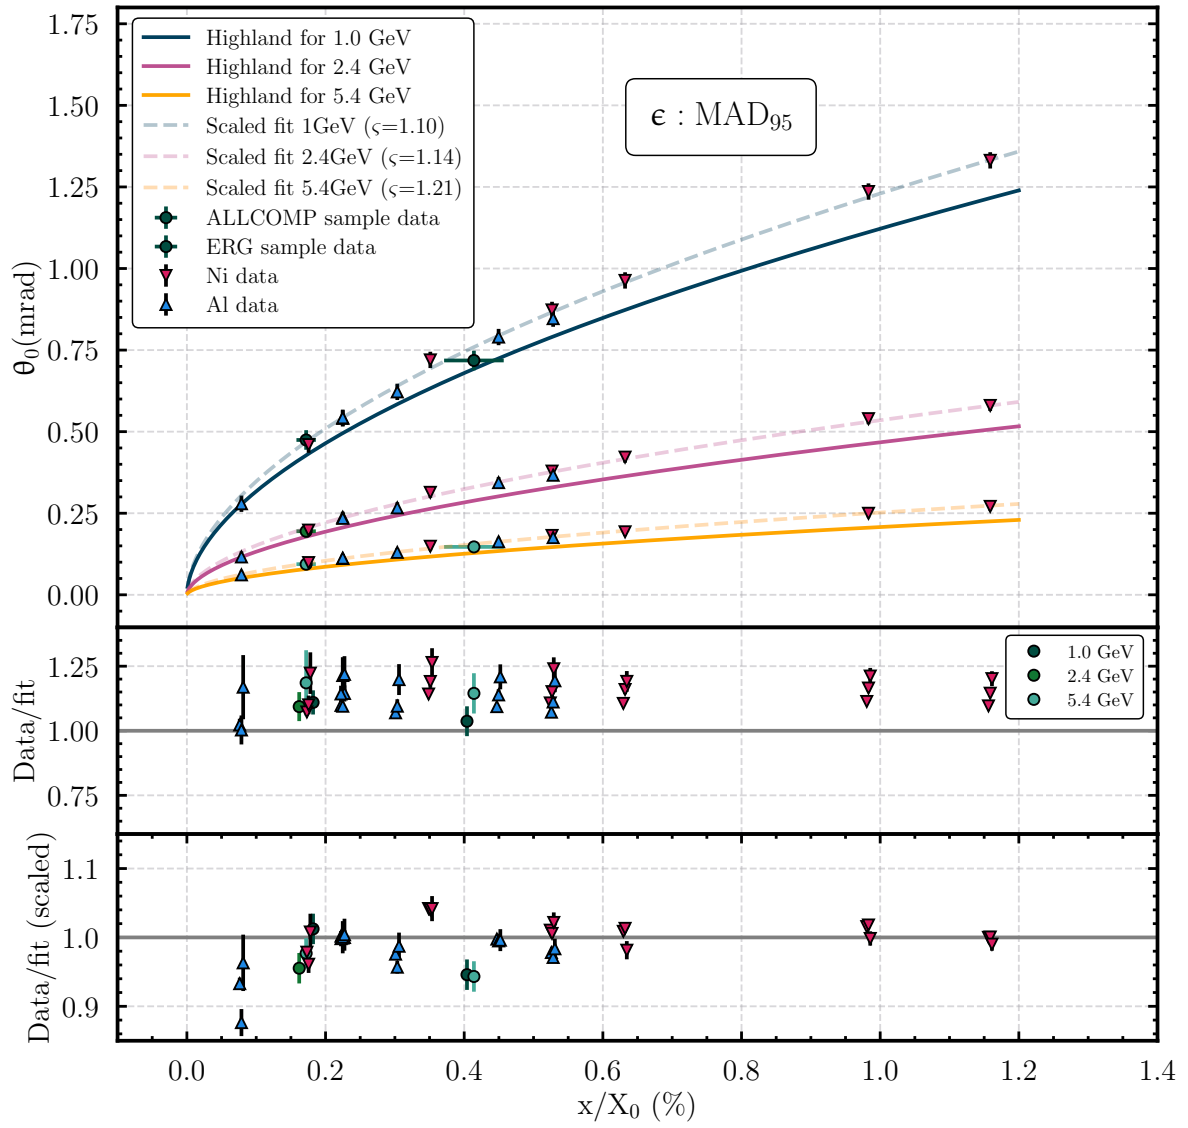


Figure E.4 – Highland scaling using the MAD width estimator at the 95% quantile.

Appendix F

3D-printed jigs cross-section




Figure F.1 – Cross section through the v1 and v2 3D-printed cylindrical jigs showing the opening of the various layers.

Appendix G

Effects of the pixel pitch asymmetry on the kink angle distributions

To ensure there is no bias introduced by the non-square pixels of the ALPIDE sensor, some checks are performed. The pixel pitch in the x direction measures $29.24\text{ }\mu\text{m}$, while in the y direction it is $26.88\text{ }\mu\text{m}$, resulting in an approximate 8.8% difference.

A toy MC simulation was performed using the Allpix² generic pixel detector simulation framework [125]. In this simulation, a telescope consisting of six $50\text{ }\mu\text{m}$ thick pixel sensors, each spaced 2 cm apart – with the exception of the middle two planes, which were separated by 4 cm in order to mimic the setup used for measuring the foam samples – was modeled. Each sensor featured a matrix of 1024×512 pixels, with the pixel pitch in the x direction fixed at $30\text{ }\mu\text{m}$, while the y pitch varied from 20 to $200\text{ }\mu\text{m}$. The exact pixel pitch of the ALPIDE sensor was also simulated. The real electric field configuration of the ALPIDE sensor is not used. Instead, a linear electric field is introduced in order to achieve a cluster size distribution with approximately 50% single-pixel clusters and 50% two-pixel clusters (shown in Fig. G.1).




Figure G.1 – The cluster size distribution for each pixel pitch configuration.

A beam of 2.4 GeV electrons was used to interact with the telescope, and for each setting about 5 million tracks were recorded. The scattering with the air volume and the sensors, as well as the deposition of energy and generation of charge carriers is done with the provided interface to Geant4. The propagation of the charge carriers in the semiconductor to the collection diodes, the amplification and digitization are computed afterwards. The resulting

data was imported into the Corryvreckan software where the same track reconstruction algorithm is applied, as in data, while keeping the kink angle distributions in the x and y directions separately.

The simulation results showed that as the asymmetry between pixel pitches increases, the kink angle distribution in the y direction (the one that varies) starts to exhibit a multi-peak structure, as illustrated in Fig. G.2. In contrast, square pixels yield a symmetric, single-peak distribution centered around zero. As the ratio of pixel pitches increases, additional peaks appear, with up to five peaks visible when the y pitch is more than three times that of x . The distance between the peaks appears to be the same for a specific setting.




Figure G.2 – Kink angle distribution in the y direction (the one where the pixel pitch varies) develops a multi-peak structure as the asymmetry between x and y pitches increases.

This multi-peak structure arises from the digital nature of pixel detectors, a phenomenon previously observed and discussed for gas detectors [126, 127]. After digitization, a pixel can only be "on" or "off" (it recorded the passage of a particle or not). If only one pixel recorded the passage of a particle, the cluster position in the global frame of reference will be in the middle of that pixel (blue point in Fig. G.3).

When charge sharing occurs between two neighboring pixels, the reconstructed cluster position for a two-pixel cluster, determined using the center of gravity method, falls between the two involved pixels. Given that each pixel has four neighboring sides, this results in four possible cluster positions. This pattern is repeated across the entire pixel matrix, creating a discrete grid-like distribution of cluster positions.




Figure G.3 – Center of gravity reconstruction for 1 and 2 px clusters.

The fact that all the other detectors also have this sort of pattern discretises the 3D space in terms of the measurements and thus the possible measured kink angles considering a decent alignment. Misalignment helps to mitigate this effect as the relative translation and rotation counteracts the discretisation. Ideally, with perfect alignment and perpendicular tracks that do not scatter, this would lead to a multi-peak structure, even for square digital pixels.

In practice, however, each sensor is slightly misaligned relative to the others, and the tracks approach the sensors at various small angles, forming a distribution centered around the perpendicular point on the sensor surface. Additionally, scattering in the air and the sensors further diffuses the track reconstruction. This quantization effect becomes more pronounced when sensors are positioned closer together and decreases as they are spaced further apart. For small differences between the pixel pitches, these fluctuations are absorbed into the main peak. For more details, refer to [102].

For the different pixel pitches studied in this toy MC simulation, the distance between the peaks was measured at the center of the telescope and is shown in Fig. G.4.




Figure G.4 – Relation between the simulated pixel pitch in the y direction and the inter-peak distance in the k_y angle distribution. The x distribution remains unaffected.

For pixel pitches below 30 μm , seeing the multiple peaks is difficult, as they are absorbed into the main peak. In these cases, the width of the distribution is used as the uncertainty. However, as the pixel pitch increases, the distance between the peaks also increases, as previously explained. The measurements indicate a linear relationship between the pixel pitch and the distance between the peaks.

This can be understood as the distance between two points on the grid. For instance, with the largest pixel pitch in the simulation, a distance between peaks of 2.5 mrad is measured.

The angular displacement is given by $\Delta\theta = s/l$ rad, where s is the displacement, l the distance between sensors. Using this formula, a displacement value of $s = \Delta\theta \cdot l = 2.5 \text{ mrad} \cdot 4 \text{ cm} = 100 \text{ }\mu\text{m}$ is obtained. This represents half the pixel pitch for this measurement, corresponding to the distance between two reconstructed cluster positions by means of the center of gravity method along the y axis.

For the ALPIDE sensor, the cluster size distribution extends to sizes of four, which are common at nominal operating settings, increasing the density of points in the grid.

The simulation was extended to study the kink angle distribution for both a perfectly aligned telescope, using the MC truth, and for standard alignment procedures that more closely resemble a realistic telescope alignment. It was observed that the multi-peak distribution is most pronounced in a perfectly aligned telescope, since the grids align perfectly. This effect is still be visible when performing a manual alignment when the difference between the alignment values and the true values is minimal. In general, detecting a multi-peak distribution can be an indicator of good alignment, as it suggests that the alignment is very close to the true position.

Publications

Some of the work presented in this thesis has been previously published or is currently in the process of publication in the following journal articles and conference proceedings. The author's contributions to these works vary in scope, ranging from major involvement in preparation, data acquisition, analysis and writing, to more supportive roles in data analysis and technical development. The publications are categorized into three groups: those where the author had a major contribution, those with a minor contribution, and other work not related to pixel sensors.

Major contribution

- 2024 Demonstration of tracking capabilities of bent MAPS mimicking a truly cylindrical barrel configuration**
In preparation. Draft ready.
- 2024 Detector efficiency and spatial resolution of Monolithic Active Pixel Sensors bent to different radii measured with a 5.4 GeV electron beam**
In preparation. Draft ready.
- 2022 Testbeam performance results of bent ALPIDE Monolithic Active Pixel Sensors in view of the ALICE Inner Tracking System 3**
B.M. Blidaru on behalf of the ALICE collaboration
JINST 17 C09006
DOI: 10.1088/1748-0221/17/09/C09006
- 2022 First demonstration of in-beam performance of bent Monolithic Active Pixel Sensors**
ALICE ITS Project
Nucl. Instrum. Methods Phys. Res. A 1028 (2022) 166280
DOI: 10.1016/j.nima.2021.166280

Minor contribution

- 2024 Characterisation of analogue Monolithic Active Pixel Sensor test structures implemented in a 65 nm CMOS imaging process**
Gianluca Aglieri Rinella et al.
Nucl. Instrum. Methods Phys. Res. A 1069 (2024) 169896
DOI: 10.1016/j.nima.2024.169896

Other work not related to silicon pixel sensors**2021 The upgrade of the ALICE TPC with GEMs and continuous readout**

ALICE TPC collaboration

JINST 16 P03022

DOI: 10.1088/1748-0221/16/03/P03022

2019 Secondary discharge studies in single- and multi-GEM structures

A. Deisting, B.M. Blidaru et al.

Nucl. Instrum. Methods Phys. Res. A 937 (2019) pg 168-180

DOI: 10.1016/j.nima.2019.05.057

Bibliography

- [1] Andrew Purcell. “Go on a particle quest at the first CERN webfest. Le premier webfest du CERN se lance à la conquête des particules”. In: 35/2012 (2012), p. 10. URL: <https://cds.cern.ch/record/1473657>.
- [2] S. Navas et al. “Review of Particle Physics”. In: *Phys. Rev. D* 110 (3 Aug. 2024), p. 030001. DOI: [10.1103/PhysRevD.110.030001](https://doi.org/10.1103/PhysRevD.110.030001). URL: <https://link.aps.org/doi/10.1103/PhysRevD.110.030001>.
- [3] Peter Braun-Munzinger and Johanna Stachel. “The quest for the quark-gluon plasma”. In: *Nature* 448 (2007), pp. 302–309. DOI: [10.1038/nature06080](https://doi.org/10.1038/nature06080).
- [4] Esma Mobs. “The CERN accelerator complex in 2019. Complexe des accélérateurs du CERN en 2019”. In: (2019). General Photo. URL: <https://cds.cern.ch/record/2684277>.
- [5] CERN. *CERN’s accelerator complex*. URL: <https://home.cern/science/accelerators/accelerator-complex>. (accessed: 15.05.2024).
- [6] Arturo Tauro. “ALICE Schematics”. General Photo. 2017. URL: <https://cds.cern.ch/record/2263642>.
- [7] K. Aamodt et al. “The ALICE experiment at the CERN LHC”. In: *JINST* 3 (2008), S08002. DOI: [10.1088/1748-0221/3/08/S08002](https://doi.org/10.1088/1748-0221/3/08/S08002).
- [8] J. Adolfsson et al. “The upgrade of the ALICE TPC with GEMs and continuous readout”. In: *JINST* 16.03 (2021), P03022. DOI: [10.1088/1748-0221/16/03/P03022](https://doi.org/10.1088/1748-0221/16/03/P03022). arXiv: [2012.09518 \[physics.ins-det\]](https://arxiv.org/abs/2012.09518).
- [9] Hermann Kolanoski and Norbert Wermes. *Particle Detectors*. Oxford University Press, June 2020. ISBN: 978-0-19-885836-2.
- [10] William R. Leo. *Techniques for Nuclear and Particle Physics Experiments*. Springer, 1994. ISBN: 978-3-540-57280-0, 978-3-642-57920-2. DOI: [10.1007/978-3-642-57920-2](https://doi.org/10.1007/978-3-642-57920-2).
- [11] Claude Leroy and Pier-Giorgio Rancoita. *Principles of radiation interaction in matter and detection*. Singapore: World Scientific, 2011. ISBN: 978-981-238-909-1. DOI: [10.1142/5578](https://doi.org/10.1142/5578).
- [12] Hans Bichsel. “Straggling in Thin Silicon Detectors”. In: *Rev. Mod. Phys.* 60 (1988), pp. 663–699. DOI: [10.1103/RevModPhys.60.663](https://doi.org/10.1103/RevModPhys.60.663).

- [13] S. Meroli, D. Passeri, and L. Servoli. “Energy loss measurement for charged particles in very thin silicon layers”. In: *JINST* 6 (2011), P06013. DOI: [10.1088/1748-0221/6/06/P06013](https://doi.org/10.1088/1748-0221/6/06/P06013).
- [14] R. Frühwirth and M. Regler. “On the quantitative modelling of core and tails of multiple scattering by Gaussian mixtures”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 456.3 (2001), pp. 369–389. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(00\)00589-1](https://doi.org/10.1016/S0168-9002(00)00589-1). URL: <https://www.sciencedirect.com/science/article/pii/S0168900200005891>.
- [15] H. A. Bethe. “Molière’s Theory of Multiple Scattering”. In: *Phys. Rev.* 89 (6 Mar. 1953), pp. 1256–1266. DOI: [10.1103/PhysRev.89.1256](https://doi.org/10.1103/PhysRev.89.1256). URL: <https://link.aps.org/doi/10.1103/PhysRev.89.1256>.
- [16] Virgil L. Highland. “Some Practical Remarks on Multiple Scattering”. In: *Nucl. Instrum. Meth.* 129 (1975), p. 497. DOI: [10.1016/0029-554X\(75\)90743-0](https://doi.org/10.1016/0029-554X(75)90743-0).
- [17] Gerald R. Lynch and Orin I. Dahl. “Approximations to multiple Coulomb scattering”. In: *Nucl. Instrum. Meth. B* 58 (1991), pp. 6–10. DOI: [10.1016/0168-583X\(91\)95671-Y](https://doi.org/10.1016/0168-583X(91)95671-Y).
- [18] Frank Hartmann. *Evolution of Silicon Sensor Technology in Particle Physics*. Vol. 275. Springer Tracts in Modern Physics. Springer, 2017. ISBN: 978-3-319-64434-9, 978-3-319-64436-3. DOI: [10.1007/978-3-319-64436-3](https://doi.org/10.1007/978-3-319-64436-3).
- [19] Chris Damerell. *Tracking the rise of pixel detectors*. 2.07.2021. URL: <https://cerncourier.com/a/tracking-the-rise-of-pixel-detectors/>.
- [20] Leonardo Rossi et al. *Pixel Detectors. From Fundamentals to Applications*. Particle Acceleration and Detection. Springer, 2006. ISBN: 978-3-540-28332-4, 978-3-642-06652-8, 978-3-540-28333-1. DOI: [10.1007/3-540-28333-1](https://doi.org/10.1007/3-540-28333-1).
- [21] William Gayde. *How CPUs are Designed, Part 3: Building the Chip*. 20.05.2019. URL: <https://www.techspot.com/article/1840-how-cpus-are-designed-and-built-part-3/>. (accessed: 22.06.2024).
- [22] C. Claeys and E. Simoen. *Radiation Effects in Advanced Semiconductor Materials and Devices*. Springer Series in Materials Science. Springer Berlin Heidelberg, 2013. ISBN: 9783662049747. URL: <https://books.google.de/books?id=AonqCAAQBAJ>.
- [23] W. Snoeys. “CMOS monolithic active pixel sensors for high energy physics”. In: *Nucl. Instrum. Meth. A* 765 (2014). Ed. by Yoshinobu Unno et al., pp. 167–171. DOI: [10.1016/j.nima.2014.07.017](https://doi.org/10.1016/j.nima.2014.07.017).

- [24] W. Snoeys. “Monolithic CMOS Sensors for high energy physics - Challenges and perspectives”. In: *Nucl. Instrum. Meth. A* 1056 (2023), p. 168678. DOI: [10.1016/j.nima.2023.168678](https://doi.org/10.1016/j.nima.2023.168678).
- [25] K. Arndt et al. “Technical design of the phase I Mu3e experiment”. In: *Nucl. Instrum. Meth. A* 1014 (2021), p. 165679. DOI: [10.1016/j.nima.2021.165679](https://doi.org/10.1016/j.nima.2021.165679). arXiv: [2009.11690](https://arxiv.org/abs/2009.11690) [[physics.ins-det](#)].
- [26] Hannah Schmitz et al. “Mighty Tracker – Performance Studies of the MightyPix for LHCb”. In: (Feb. 2024). arXiv: [2402.08428](https://arxiv.org/abs/2402.08428) [[physics.ins-det](#)].
- [27] A. Schöning et al. “MuPix and ATLASPix – Architectures and Results”. In: *PoS Vertex2019* (2020), p. 024. DOI: [10.22323/1.373.0024](https://doi.org/10.22323/1.373.0024). arXiv: [2002.07253](https://arxiv.org/abs/2002.07253) [[physics.ins-det](#)].
- [28] A. Dorokhov et al. “High resistivity CMOS pixel sensors and their application to the STAR PXL detector”. In: *Nucl. Instrum. Meth. A* 650 (2011). Ed. by Roland Horisberger, Danek Kotlinski, and Andrey Starodumov, pp. 174–177. DOI: [10.1016/j.nima.2010.12.112](https://doi.org/10.1016/j.nima.2010.12.112).
- [29] Gianluca Aglieri Rinella. “The ALPIDE pixel sensor chip for the upgrade of the ALICE Inner Tracking System”. In: *Nucl. Instrum. Meth. A* 845 (2017). Ed. by G. Badurek et al., pp. 583–587. DOI: [10.1016/j.nima.2016.05.016](https://doi.org/10.1016/j.nima.2016.05.016).
- [30] M. Deveau et al. “Observations on MIMOSIS-0, the first dedicated CPS prototype for the CBM MVD”. In: *Nucl. Instrum. Meth. A* 958 (2020). Ed. by Manfred Krammer et al., p. 162653. DOI: [10.1016/j.nima.2019.162653](https://doi.org/10.1016/j.nima.2019.162653). arXiv: [1909.05614](https://arxiv.org/abs/1909.05614) [[physics.ins-det](#)].
- [31] J. A. Ballin et al. “Design and performance of a CMOS study sensor for a binary readout electromagnetic calorimeter”. In: *JINST* 6 (2011), P05009. DOI: [10.1088/1748-0221/6/05/P05009](https://doi.org/10.1088/1748-0221/6/05/P05009). arXiv: [1103.4265](https://arxiv.org/abs/1103.4265) [[physics.ins-det](#)].
- [32] N.Wermes. *The Monopix detectors – CMOS pixel detectors with large and small electrodes*. 18.02.2020. URL: <https://indico.cern.ch/event/813597>.
- [33] W. Snoeys et al. “A process modification for CMOS monolithic active pixel sensors for enhanced depletion, timing performance and radiation tolerance”. In: *Nucl. Instrum. Meth. A* 871 (2017), pp. 90–96. DOI: [10.1016/j.nima.2017.07.046](https://doi.org/10.1016/j.nima.2017.07.046).
- [34] Gianluca Aglieri Rinella et al. “Characterisation of analogue Monolithic Active Pixel Sensor test structures implemented in a 65 nm CMOS imaging process”. In: (Mar. 2024). arXiv: [2403.08952](https://arxiv.org/abs/2403.08952) [[physics.ins-det](#)].

- [35] Gianluca Aglieri Rinella et al. “Digital pixel test structures implemented in a 65 nm CMOS process”. In: *Nucl. Instrum. Meth. A* 1056 (2023), p. 168589. DOI: [10.1016/j.nima.2023.168589](https://doi.org/10.1016/j.nima.2023.168589). arXiv: [2212.08621](https://arxiv.org/abs/2212.08621) [[physics.ins-det](#)].
- [36] Valentina Raskina and Filip Křížek. “Characterization of Highly Irradiated ALPIDE Silicon Sensors”. In: *Universe* 5.4 (2019), p. 91. DOI: [10.3390/universe5040091](https://doi.org/10.3390/universe5040091).
- [37] ALICE collaboration. “The ALICE experiment – A journey through QCD”. In: (Nov. 2022). arXiv: [2211.04384](https://arxiv.org/abs/2211.04384) [[nucl-ex](#)].
- [38] B Abelev et al. “Upgrade of the ALICE Experiment: Letter Of Intent”. In: *J. Phys. G* 41 (2014), p. 087001. DOI: [10.1088/0954-3899/41/8/087001](https://doi.org/10.1088/0954-3899/41/8/087001).
- [39] ALICE collaboration. “ALICE upgrades during the LHC Long Shutdown 2”. In: (Feb. 2023). arXiv: [2302.01238](https://arxiv.org/abs/2302.01238) [[physics.ins-det](#)].
- [40] B Abelev et al. “Technical Design Report for the Upgrade of the ALICE Inner Tracking System”. In: *J. Phys. G* 41 (2014), p. 087002. DOI: [10.1088/0954-3899/41/8/087002](https://doi.org/10.1088/0954-3899/41/8/087002).
- [41] Jochen Klein and Felix Reidt. “ALICE ITS Inner Barrel installation”. General Photo. 2021. URL: <https://cds.cern.ch/record/2766802>.
- [42] F. Reidt. “Upgrade of the ALICE ITS detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1032 (2022), p. 166632. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2022.166632>. URL: <https://www.sciencedirect.com/science/article/pii/S0168900222002042>.
- [43] L Musa. *Conceptual Design Report for the Upgrade of the ALICE ITS*. Tech. rep. Geneva: CERN, 2012. URL: <https://cds.cern.ch/record/1431539>.
- [44] S. et al. Acharya. “First measurement of Λ_c^+ production down to $p_T = 0$ in pp and p -Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV”. In: *Phys. Rev. C* 107 (6 June 2023), p. 064901. DOI: [10.1103/PhysRevC.107.064901](https://doi.org/10.1103/PhysRevC.107.064901). URL: <https://link.aps.org/doi/10.1103/PhysRevC.107.064901>.
- [45] M. Mager. “ALPIDE, the Monolithic Active Pixel Sensor for the ALICE ITS upgrade”. In: *Nucl. Instrum. Meth. A* 824 (2016). Ed. by Maria Giuseppina Bisogni et al., pp. 434–438. DOI: [10.1016/j.nima.2015.09.057](https://doi.org/10.1016/j.nima.2015.09.057).
- [46] Chenfei Yang et al. “A Prototype Readout System for the ALPIDE Pixel Sensor”. In: *IEEE Trans. Nucl. Sci.* 66.7 (2019), pp. 1088–1094. DOI: [10.1109/TNS.2019.2913335](https://doi.org/10.1109/TNS.2019.2913335).

- [47] Sergio B. Ricciarini et al. “Enabling low-power MAPS-based space trackers: a sparsified readout based on smart clock gating for the High Energy Particle Detector HEPD-02”. In: *PoS ICRC2021* (2021), p. 071. DOI: [10.22323/1.395.0071](https://doi.org/10.22323/1.395.0071).
- [48] J. Heymes et al. “Study of the depletion depth in a frontside biased CMOS pixel sensors”. In: *JINST* 14.01 (2019), P01018. DOI: [10.1088/1748-0221/14/01/P01018](https://doi.org/10.1088/1748-0221/14/01/P01018).
- [49] Håkan Wennlöf et al. *Simulating Monolithic Active Pixel Sensors: A Technology-Independent Approach Using Generic Doping Profiles*. 2024. arXiv: [2408.00027](https://arxiv.org/abs/2408.00027) [[physics.ins-det](#)]. URL: <https://arxiv.org/abs/2408.00027>.
- [50] G. Aglieri Rinella et al. “Charge collection properties of TowerJazz 180nm CMOS Pixel Sensors in dependence of pixel geometries and bias parameters, studied using a dedicated test-vehicle: The Investigator chip”. In: *Nucl. Instrum. Meth. A* 988 (2021), p. 164859. DOI: [10.1016/j.nima.2020.164859](https://doi.org/10.1016/j.nima.2020.164859). arXiv: [2009.10517](https://arxiv.org/abs/2009.10517) [[physics.ins-det](#)].
- [51] G. Aglieri Rinella et al. “Optimization of a 65 nm CMOS imaging process for monolithic CMOS sensors for high energy physics”. In: *PoS Pixel2022* (2023), p. 083. DOI: [10.22323/1.420.0083](https://doi.org/10.22323/1.420.0083).
- [52] “Expression of Interest for an ALICE ITS Upgrade in LS3”. In: (2018).
- [53] Luciano Musa. *Letter of Intent for an ALICE ITS Upgrade in LS3*. Tech. rep. Geneva: CERN, 2019. DOI: [10.17181/CERN-LHCC-2019-018](https://doi.org/10.17181/CERN-LHCC-2019-018). URL: <https://cds.cern.ch/record/2703140>.
- [54] Gianluca Aglieri Rinella. *MOSAIX as a Case Study: Overview of the Stitched Chip for ITS3*. 30.06.2024. URL: <https://indico.cern.ch/event/1417976>.
- [55] ALICE collaboration The. *Technical Design report for the ALICE Inner Tracking System 3 - ITS3 ; A bent wafer-scale monolithic pixel detector*. Tech. rep. Co-project Manager: Magnus Mager, magnus.mager@cern.ch. Geneva: CERN, 2024. URL: <https://cds.cern.ch/record/2890181>.
- [56] David Cohen et al. *Stitching design rules for forming interconnect layers*. United States Patent US 6225013 B1. 1.05.2001.
- [57] “Upgrade of the ALICE Inner Tracking System during LS3: study of physics performance”. In: (2023).
- [58] Carolina Reetz. *ALICE: strangeness tracking (short track segment from charged strange particles like Xi and Hypertriton)*. 13.10.2023. URL: <https://indico.gsi.de/event/18341/>.

- [59] Jochen Klein and Marco Van Leeuwen. “The ALICE Upgrades: Toward a Next-Generation Heavy-Ion Experiment at the LHC”. In: *Nucl. Phys. News* 31.4 (2021), pp. 30–33. DOI: [10.1080/10619127.2021.1954450](https://doi.org/10.1080/10619127.2021.1954450).
- [60] “Letter of intent for ALICE 3: A next-generation heavy-ion experiment at the LHC”. In: (Nov. 2022). arXiv: [2211.02491 \[physics.ins-det\]](https://arxiv.org/abs/2211.02491).
- [61] Nicola Nicassio. “ALICE 3: a next-generation heavy-ion detector for LHC Run 5 and beyond”. In: *PoS EPS-HEP2023* (2023). Contribution to proceedings of The European Physical Society Conference on High Energy Physics (EPS-HEP2023), 21-25 August 2023, Hamburg, Germany, 6 pages, 4 figures, p. 540. DOI: [10.22323/1.449.0540](https://doi.org/10.22323/1.449.0540). arXiv: [2311.10164](https://arxiv.org/abs/2311.10164). URL: <https://cds.cern.ch/record/2880429>.
- [62] Anton Andronic et al. “The multiple-charm hierarchy in the statistical hadronization model”. In: *JHEP* 07 (2021), p. 035. DOI: [10.1007/JHEP07\(2021\)035](https://doi.org/10.1007/JHEP07(2021)035). arXiv: [2104.12754 \[hep-ph\]](https://arxiv.org/abs/2104.12754).
- [63] Lauren Biron. *Put it to the test beam*. 11.05.2019. URL: <https://www.symmetrymagazine.org/article/put-it-to-the-test-beam>.
- [64] Lennart Huth. “A High Rate Testbeam Data Acquisition System and Characterization of High Voltage Monolithic Active Pixel Sensors”. PhD thesis. University of Heidelberg, 2019. DOI: [0.11588/heidok.00025785](https://doi.org/0.11588/heidok.00025785).
- [65] M.Mager. *The Telescope Optimizer*. 9.02.2024. URL: <https://mmager.web.cern.ch/telescope/tracking.html>. (accessed: 15.01.2024).
- [66] S. Spannagel and H. Jansen. *GBL Track Resolution Calculator V2.0*. 2 Apr. 2016. URL: [doi:10.5281/zenodo.48795](https://doi.org/10.5281/zenodo.48795). (accessed: 15.01.2024).
- [67] *TEfficiency Class*. URL: <https://root.cern.ch/doc/master/classTEfficiency.html>.
- [68] G. Aglieri Rinella et al. “First demonstration of in-beam performance of bent Monolithic Active Pixel Sensors”. In: *Nucl. Instrum. Meth. A* 1028 (2022), p. 166280. DOI: [10.1016/j.nima.2021.166280](https://doi.org/10.1016/j.nima.2021.166280). arXiv: [2105.13000 \[physics.ins-det\]](https://arxiv.org/abs/2105.13000).
- [69] B.M. Blidaru et.al. “Detector efficiency and spatial resolution of Monolithic Active Pixel Sensors bent to different radii measured with a 5.4 GeV electron beam”. first version ready, unpublished article. 2024.
- [70] B.M. Blidaru et.al. “Demonstration of tracking capabilities of bent MAPS mimicking a truly cylindrical barrel configuration”. first version ready, in preparation, unpublished article. 2024.

- [71] Bogdan-Mihail Blidaru. “Testbeam performance results of bent ALPIDE monolithic active pixel sensors in view of the ALICE Inner Tracking System 3”. In: *JINST* 17.09 (2022), p. C09006. DOI: [10.1088/1748-0221/17/09/C09006](https://doi.org/10.1088/1748-0221/17/09/C09006). arXiv: [2112.10414](https://arxiv.org/abs/2112.10414) [[physics.ins-det](#)].
- [72] D.A. van den Ende et al. “Mechanical and electrical properties of ultra-thin chips and flexible electronics assemblies during bending”. In: *Microelectronics Reliability* 54.12 (2014), pp. 2860–2870. ISSN: 0026-2714. DOI: <https://doi.org/10.1016/j.microrel.2014.07.125>. URL: <https://www.sciencedirect.com/science/article/pii/S0026271414003278>.
- [73] Magnus Mager. “The future of bent MAPS, full-wafer (stitched) design: Status and challenges”. In: *Nucl. Instrum. Meth. A* 1064 (2024), p. 169447. DOI: [10.1016/j.nima.2024.169447](https://doi.org/10.1016/j.nima.2024.169447).
- [74] M. J. Rossewij et al. “Electrical/piezoresistive effects in bent Alpide MAPS”. In: *JINST* 19.04 (2024), p. C04057. DOI: [10.1088/1748-0221/19/04/C04057](https://doi.org/10.1088/1748-0221/19/04/C04057). arXiv: [2403.12298](https://arxiv.org/abs/2403.12298) [[physics.ins-det](#)].
- [75] Alperen Yüncü. “PhD thesis in preparation”. PhD thesis. Heidelberg University.
- [76] G. Contin. *ITS3 R&D ALPIDE bending tests update*. 14.02.2020. URL: <https://indico.cern.ch/event/886294/>.
- [77] Alperen Yüncü. *Assembly and operation of bent silicon pixel sensors towards the ALICE ITS3*. 8.04.2024. URL: <https://indico.cern.ch/event/1347765/>.
- [78] Franck Agnese. *WP4 - Sensor Thinning, Bending and gluing, Interconnection : C4PI – Microtechnics contribution*. 11.11.2020. URL: <https://indico.cern.ch/event/983701>.
- [79] Dominik Dannheim et al. “Corryvreckan: A Modular 4D Track Reconstruction and Analysis Software for Test Beam Data”. In: *JINST* 16.03 (2021), P03008. DOI: [10.1088/1748-0221/16/03/P03008](https://doi.org/10.1088/1748-0221/16/03/P03008). arXiv: [2011.12730](https://arxiv.org/abs/2011.12730) [[physics.ins-det](#)].
- [80] D. Kim et al. “Front end optimization for the monolithic active pixel sensor of the ALICE Inner Tracking System upgrade”. In: *JINST* 11.02 (2016), p. C02042. DOI: [10.1088/1748-0221/11/02/C02042](https://doi.org/10.1088/1748-0221/11/02/C02042).
- [81] Pascal Becht. “Performance and characterisation of bent monolithic active pixel sensors (MAPS) for the application in future tracking detectors using the example of the ALPIDE chip”. MA thesis. Heidelberg University, 2021.
- [82] Y. Liu et al. “EUDAQ2—A flexible data acquisition software framework for common test beams”. In: *JINST* 14.10 (2019), P10033. DOI: [10.1088/1748-0221/14/10/P10033](https://doi.org/10.1088/1748-0221/14/10/P10033). arXiv: [1907.10600](https://arxiv.org/abs/1907.10600) [[physics.ins-det](#)].

- [83] Lukas Lautner. “Characterization and performance of curved chip prototypes for the ALICE ITS3 upgrade”. Ph.D. thesis in preparation. 2024.
- [84] E.M. Okkinga. “ITS3 upgrade: Toward a fully cylindrical detector and beyond. Study on bending silicon chips for the Inner Tracking System Upgrade”. Supervisors: Alessandro Grelli & Marcel Rossewijn. bachelorthesis. Fontys University of Applied Science, 2023.
- [85] P. Martinengo. “The new Inner Tracking System of the ALICE experiment”. In: *Nucl. Phys. A* 967 (2017). Ed. by Ulrich Heinz, Olga Evdokimov, and Peter Jacobs, pp. 900–903. DOI: [10.1016/j.nuclphysa.2017.05.069](https://doi.org/10.1016/j.nuclphysa.2017.05.069).
- [86] Ricardo Wolker. “Measurement of the Momentum Spread of a DESY Test-Beam Line Using CMS Pixel Modules”. In: *Summer student programme DESY 2016* (2016). URL: <https://www.desy.de/f/students/2016/reports/RicardoWoelker.pdf.gz>.
- [87] Dominik Dannheim et al. “Combining TCAD and Monte Carlo Methods to Simulate CMOS Pixel Sensors with a Small Collection Electrode using the Allpix² Squared Framework”. In: *Nucl. Instrum. Meth. A* 964 (2020), p. 163784. DOI: [10.1016/j.nima.2020.163784](https://doi.org/10.1016/j.nima.2020.163784). arXiv: [2002.12602](https://arxiv.org/abs/2002.12602) [[physics.ins-det](#)].
- [88] Zbynek Drasal and Werner Riegler. “An extension of the Gluckstern formulae for multiple scattering: Analytic expressions for track parameter resolution using optimum weights”. In: *Nucl. Instrum. Meth. A* 910 (2018), pp. 127–132. DOI: [10.1016/j.nima.2018.08.078](https://doi.org/10.1016/j.nima.2018.08.078). arXiv: [1805.12014](https://arxiv.org/abs/1805.12014) [[physics.ins-det](#)].
- [89] M. Concas et al. *Status and perspectives of the Strangeness Tracking in Run 3*. 2023. URL: https://indico.cern.ch/event/1276818/contributions/5366438/attachments/2631012/4550558/strangeness_pag_180423.pdf (visited on 04/18/2023).
- [90] Fabrizio Grosa. personal communication. May 29, 2024.
- [91] Aitor Amatriain et al. “Development of an air cooling system with low material budget for high-energy physics applications”. In: *Appl. Thermal Eng.* 236 (2024), p. 121699. DOI: [10.1016/j.applthermaleng.2023.121699](https://doi.org/10.1016/j.applthermaleng.2023.121699).
- [92] Aitor Amatriain, Corrado Gargiulo, and Gonzalo Rubio. “Numerical and experimental study of open-cell foams for the characterization of heat exchangers”. In: (June 2023). DOI: [10.1016/j.ijheatmasstransfer.2023.124701](https://doi.org/10.1016/j.ijheatmasstransfer.2023.124701). arXiv: [2306.03699](https://arxiv.org/abs/2306.03699) [[physics.flu-dyn](#)].
- [93] Pieter Ijzermans and Massimo Angeletti. personal communication. May 16, 2024.
- [94] R. L. Workman et al. “Review of Particle Physics”. In: *PTEP* 2022 (2022), p. 083C01. DOI: [10.1093/ptep/ptac097](https://doi.org/10.1093/ptep/ptac097).

- [95] *Araldite® 2011*. Technical Data Sheet. Huntsman. Aug. 2014.
- [96] Peter Kodys. *SCT end-cap module - Materials*. 1.08.2005. URL: https://hepwww.pp.rl.ac.uk/groups/atlas-sct/engineering/material_budget/models/Endcap_Module/ATLAS_ECSCCT_Materials.pdf.
- [97] Swiss Composite. *Produkteuebersicht - Faserverbundwerkstoffe*. 13.07.2023. URL: <https://www.swiss-composite.ch/pdf/Produkteuebersicht.pdf>.
- [98] C.Gargiulo. *BBM: microCT Analysis*. 24.11.2020. URL: <https://indico.cern.ch/event/967692/>.
- [99] P.V. Leita. *TPSCo-65: Setting up an access framework and platform*. 20.11.2023. URL: <https://indico.cern.ch/event/1339888/>.
- [100] U. Stolzenberg et al. “Radiation length imaging with high resolution telescopes”. In: *Nucl. Instrum. Meth. A* 845 (2017). Ed. by G. Badurek et al., pp. 173–176. DOI: [10.1016/j.nima.2016.06.086](https://doi.org/10.1016/j.nima.2016.06.086). arXiv: [1609.02402](https://arxiv.org/abs/1609.02402) [[physics.ins-det](#)].
- [101] Hendrik Jansen et al. “Scattering Studies with the DATURA Beam Telescope”. In: *Springer Proc. Phys.* 212 (2018). Ed. by Zhen-An Liu, pp. 243–250. DOI: [10.1007/978-981-13-1313-4_47](https://doi.org/10.1007/978-981-13-1313-4_47). arXiv: [1903.04169](https://arxiv.org/abs/1903.04169) [[physics.ins-det](#)].
- [102] Ulf Hagen Stolzenberg. “Radiation length measurements with high-resolution telescopes”. PhD thesis. Gottingen U., Gottingen U., 2019. DOI: [10.53846/goediss-7624](https://doi.org/10.53846/goediss-7624).
- [103] Paul Schütze. “Silicon Pixel Detectors - Performance after Irradiation and Application in Three-dimensional Imaging”. PhD thesis. Hamburg: Hamburg U., 2019. DOI: [10.3204/PUBDB-2019-04814](https://doi.org/10.3204/PUBDB-2019-04814).
- [104] Jan-Hendrik Arling. “Detection and Identification of Electrons and Photons - Applications in the ATLAS Experiment, for the ATLAS ITk Detector and at the DESY II Test Beam”. PhD thesis. Hamburg: Dortmund U., DESY, 2020. DOI: [10.3204/PUBDB-2020-03891](https://doi.org/10.3204/PUBDB-2020-03891).
- [105] Shreyasi Acharya et al. “Data-driven precision determination of the material budget in ALICE”. In: *JINST* 18.11 (2023), P11032. DOI: [10.1088/1748-0221/18/11/P11032](https://doi.org/10.1088/1748-0221/18/11/P11032). arXiv: [2303.15317](https://arxiv.org/abs/2303.15317) [[physics.ins-det](#)].
- [106] Alica Enderich. “Performance study on photon measurements using the Photon Conversion Method with ALICE in Run 3”. bachelorthesis. Heidelberg University, 2024.
- [107] V. Blobel. “Software alignment for tracking detectors”. In: *Nucl. Instrum. Meth. A* 566 (2006). Ed. by R. Bernhard et al., pp. 5–13. DOI: [10.1016/j.nima.2006.05.157](https://doi.org/10.1016/j.nima.2006.05.157).

- [108] Anouk Kaiser. “Tracking low momentum protons with the ALICE Pixel Detector”. bachelorthesis. Heidelberg University, 2023.
- [109] HighRR. *HighRR Testbeam School @ DESY 2024*. 10.04.2024. URL: <https://indico.cern.ch/event/1374460/>.
- [110] Christophe Croux and Peter J Rousseeuw. “Time-Efficient Algorithms for Two Highly Robust Estimators of Scale”. en. In: *Computational Statistics*. Ed. by Yadolah Dodge and Joe Whittaker. Heidelberg: Physica-Verlag HD, 1992, pp. 411–428. ISBN: 978-3-662-26813-1 978-3-662-26811-7. DOI: [10.1007/978-3-662-26811-7_58](https://doi.org/10.1007/978-3-662-26811-7_58). URL: http://link.springer.com/10.1007/978-3-662-26811-7_58 (visited on 08/02/2022).
- [111] Andrey Akinshin. *Finite-sample bias-correction factors for the median absolute deviation based on the Harrell-Davis quantile estimator and its trimmed modification*. 2022. arXiv: [2207.12005](https://arxiv.org/abs/2207.12005) [stat.ME]. URL: <https://arxiv.org/abs/2207.12005>.
- [112] I. Molina Peralta and E. García-Portugués. “A First Course on Statistical Inference”. In: Version 2.4.1. ISBN 978-84-09-29680-4. 2024. Chap. 3.3 - Consistent estimators. URL: <https://bookdown.org/egarpor/inference/>.
- [113] Rand R. Wilcox. “3 - SUMMARIZING DATA”. In: *Applying Contemporary Statistical Techniques*. Ed. by Rand R. Wilcox. Burlington: Academic Press, 2003, pp. 55–91. ISBN: 978-0-12-751541-0. DOI: <https://doi.org/10.1016/B978-012751541-0/50024-9>. URL: <https://www.sciencedirect.com/science/article/pii/B9780127515410500249>.
- [114] Niklaus Berger et al. “Multiple Coulomb Scattering in Thin Silicon”. In: *JINST* 9 (2014), P07007. DOI: [10.1088/1748-0221/9/07/P07007](https://doi.org/10.1088/1748-0221/9/07/P07007). arXiv: [1405.2759](https://arxiv.org/abs/1405.2759) [physics.ins-det].
- [115] Simon Groß-Bölting. “Material budget studies for carbon foam samples for the ITS3”. M.Sc. thesis in preparation. 2024.
- [116] W.H. Press. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Numerical Recipes: The Art of Scientific Computing. Cambridge University Press, 2007. ISBN: 9780521880688. URL: <https://books.google.de/books?id=1aAOdzK3FegC>.
- [117] Franklin A. Graybill and R. B. Deal. “Combining Unbiased Estimators”. In: *Biometrics* 15.4 (1959), pp. 543–550. ISSN: 0006341X, 15410420. URL: <http://www.jstor.org/stable/2527652> (visited on 07/28/2024).
- [118] Vanamamalai Seshadri. “Constructing Uniformly Better Estimators”. In: *Journal of the American Statistical Association* 58.301 (1963), pp. 172–175. ISSN: 01621459, 1537274X. URL: <http://www.jstor.org/stable/2282961> (visited on 07/28/2024).

- [119] Hening Huang. “Combining estimators in interlaboratory studies and meta-analyses”. In: *Research Synthesis Methods* 14.3 (2023), pp. 526–543. DOI: <https://doi.org/10.1002/jrsm.1633>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/jrsm.1633>.
- [120] Frédéric Lavancier and Paul Rochet. “A general procedure to combine estimators”. In: *Computational Statistics and Data Analysis* 94 (2016), pp. 175–192. URL: <https://hal.science/hal-00936024>.
- [121] José Luis Rodríguez-Sánchez. *Update of the ALPIDE Si-tracker simulations for studies with quasi-free scattering reactions*. 2023. URL: <https://indico.gsi.de/event/17415> (visited on 05/24/2023).
- [122] G. Contin. “MAPS-based tracking and vertexing for the Electron–Ion Collider”. In: *Nucl. Instrum. Meth. A* 1050 (2023), p. 168120. DOI: [10.1016/j.nima.2023.168120](https://doi.org/10.1016/j.nima.2023.168120).
- [123] Xuan Li. “Advanced silicon tracking detector developments for the future Electron–Ion Collider”. In: *Nucl. Instrum. Meth. A* 1057 (2023), p. 168687. DOI: [10.1016/j.nima.2023.168687](https://doi.org/10.1016/j.nima.2023.168687).
- [124] R. Diener et al. “The DESY II Test Beam Facility”. In: *Nucl. Instrum. Meth. A* 922 (2019), pp. 265–286. DOI: [10.1016/j.nima.2018.11.133](https://doi.org/10.1016/j.nima.2018.11.133). arXiv: [1807.09328 \[physics.ins-det\]](https://arxiv.org/abs/1807.09328).
- [125] Simon Spannagel et al. “Allpix²: A Modular Simulation Framework for Silicon Detectors”. In: *Nucl. Instrum. Meth. A* 901 (2018), pp. 164–172. DOI: [10.1016/j.nima.2018.06.020](https://doi.org/10.1016/j.nima.2018.06.020). arXiv: [1806.05813 \[physics.ins-det\]](https://arxiv.org/abs/1806.05813).
- [126] D. Drijard, T. Ekelof, and H. Grote. “On the Reduction in Space Resolution of Track Detectors Caused by Correlations in the Coordinate Quantization”. In: *Nucl. Instrum. Meth.* 176 (1980). Ed. by W. Bartl and M. Regler, p. 389. DOI: [10.1016/0029-554X\(80\)90732-6](https://doi.org/10.1016/0029-554X(80)90732-6).
- [127] I. Duerdoth. “Track fitting and resolution with digital detectors”. In: *Nucl. Instrum. Meth.* 203 (1982), pp. 291–297. DOI: [10.1016/0167-5087\(82\)90640-8](https://doi.org/10.1016/0167-5087(82)90640-8).

Document template

The template used for this thesis was sourced from the GitHub repository created by Krishna Kumar, available at <https://github.com/kks32/phd-thesis-template>, Version: v2.4, and is licensed under the MIT License.

Acknowledgment of AI Assistance

In the preparation of this thesis, I utilized the AI language model, ChatGPT, developed by OpenAI (GPT-4 and GPT-4o versions), as a tool for assistance in several areas.

Specifically, I used it for coding assistance and code generation, which helped streamline the implementation of various data analyses and offered more efficient coding solutions. To give a concrete example of this workflow: I tasked the AI model to generate a Python program that processes a CSV file with a specific structure, giving instructions on how to handle the specific fields and what quantities to plot and in what way. This was an iterative process leading to a final usable script, where I continuously intervened to refine the output at each step based on my specifications, ensuring that both input and output values were thoroughly checked.

Additionally, I used ChatGPT for support in writing this manuscript. In this context, I used it to suggest synonyms for various expressions in order to avoid word repetitions, check the grammar or provide reformulations of paragraphs originally written by me, in order to improve the clarity and readability. At no point did I use the software to generate text entirely on its own or to paraphrase content from work on which I was not the primary author.

I fully embrace this technology, as it significantly sped up my workflow and made the analysis and writing processes more efficient. However, I was fully aware of its potential flaws, such as hallucinations or inaccuracies, and I carefully reviewed all content modified by ChatGPT before incorporating it into this thesis.

All intellectual content, technical work, and the ideas expressed in this thesis remain my own, with ChatGPT being used as a tool to enhance the efficiency and/or readability of the process.

Declaration of originality

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. I confirm that:

- I have documented all methods, data and processes truthfully
- All data and findings in the work have not been falsified or embellished
- I have mentioned all persons who were significant facilitators of the work

The submitted electronic version of the thesis matches the printed version.

Heidelberg, 14.11.2024

Mihail-Bogdan Blidaru