

Dissertation zur Erlangung des akademischen Grades
Doctor rerum naturalium

**At the energy frontier
in particle and astroparticle physics:
Measurements of nuclear effects
in light-hadron production
with the LHCb experiment**

Julian Alexander Boelhauve

Mai 2025

Fakultät Physik
Technische Universität Dortmund

Der Fakultät Physik der Technischen Universität Dortmund vorgelegte Dissertation von Julian Alexander Boelhauve, geboren in Datteln, zur Erlangung des akademischen Grades eines Doktors der Naturwissenschaften.

Erstgutachter

Prof. Dr. Johannes Albrecht

Zweitgutachter

Priv.-Doz. Dr. Dominik Elsässer

Vorsitzender der Prüfungskommission

Prof. Dr. Heinrich Päs

Vertreter der wissenschaftlichen Mitarbeiter*innen

Dr. Gerald Schmidt

Datum des Einreichens der Arbeit

30. Mai 2025

Datum der mündlichen Prüfung

3. September 2025

Abstract

The impact of precision measurements in high-energy particle physics on astroparticle physics is illuminated in this thesis. Two measurements of light-hadron production performed with the LHCb experiment at CERN are presented. In the first measurement, the differential cross-section of prompt production of long-lived charged particles in proton–proton collisions at a centre-of-mass energy of 13 TeV is determined. The analysis is conducted in intervals of particle kinematics, and the two particle charges are studied separately. In the second measurement, the differential cross-sections of prompt production of charged pions, charged kaons and (anti)protons are analogously determined in two systems: collisions of protons with lead ions at a nucleon–nucleon centre-of-mass energy of 8.16 TeV and proton–proton collisions at 13 TeV. Moreover, hadron-production ratios and nuclear-modification factors are measured. Comparisons with contemporary hadronic-interaction models reveal that none of the models is capable of describing all the observables within the measurement uncertainties. The light-hadron composition measured in proton–proton collisions resembles the decrease in charged-pion production needed in the models to explain the Muon Puzzle in high-energy air showers.

Kurzfassung

Die Bedeutung von Präzisionsmessungen in der Hochenergie-Teilchenphysik für die Astroteilchenphysik wird in dieser Arbeit dargestellt. Zwei Messungen der Produktion leichter Hadronen, die mit dem LHCb-Experiment am CERN durchgeführt wurden, werden präsentiert. In der ersten Messung wird der differenzielle Wirkungsquerschnitt der prompten Produktion langlebiger geladener Teilchen in Proton-Proton-Kollisionen bei einer Schwerpunktsenergie von 13 TeV bestimmt. Die Analyse wird in Intervallen der Teilchenkinematik vollzogen und die beiden Teilchenladungen werden getrennt betrachtet. Analog werden in der zweiten Messung die differenziellen Wirkungsquerschnitte der prompten Produktion geladener Pionen, geladener Kaonen und von (Anti-)Protonen in zwei Systemen bestimmt: Kollisionen von Protonen mit Blei-Ionen bei einer Nukleon-Nukleon-Schwerpunktsenergie von 8,16 TeV und Proton-Proton-Kollisionen bei 13 TeV. Des Weiteren werden Hadronproduktionsverhältnisse und Kernmodifizierungsfaktoren gemessen. Vergleiche mit gegenwärtigen Modellen hadronischer Wechselwirkungen zeigen, dass keines der Modelle alle Observablen innerhalb der Messunsicherheiten zu beschreiben vermag. Die in Proton-Proton-Kollisionen gemessene Zusammensetzung leichter Hadronen ähnelt der Verringerung der Produktion geladener Pionen, die in den Modellen benötigt wird, um das Myonrätsel in hochenergetischen Luftschauern zu erklären.

Contents

1	Introduction	1
I	Foundations	5
2	Theoretical framework	7
2.1	The Standard Model of particle physics	7
2.2	Hadron production	8
3	Connection of particle and astroparticle physics	13
3.1	Cosmic rays and air showers	13
3.2	The Muon Puzzle in astroparticle physics	16
4	Experimental setup	25
4.1	The Large Hadron Collider	25
4.2	The LHCb detector	26
II	Measurements	33
5	Prompt charged-particle production	35
5.1	Data and simulation samples	35
5.2	Strategy	36
5.3	Selection	38
5.4	Background contributions	38
5.5	Efficiency	50
5.6	Propagation of uncertainty	57
5.7	Results	61
6	Prompt identified-hadron production and nuclear modification	67
6.1	Data and simulation samples	68
6.2	Strategy	68
6.3	Selection	70
6.4	Background contributions	71
6.5	Signal composition	78

Contents

6.6	Efficiency	89
6.7	Acceptance correction	94
6.8	Propagation of uncertainty	95
6.9	Results	100
III	Implications	109
7	Validation of hadronic-interaction models	111
7.1	Comparisons with measurements	111
7.2	Towards a global picture	113
8	Conclusion and outlook	117
A	Basis vectors for two-body decays	121
	Bibliography	125
	Acknowledgements	141

Abbreviations and variables

The abbreviations employed in the main text are defined below. Their places of first occurrence are given as well.

CERN	European Organization for Nuclear Research, originally <i>Conseil Européen pour la Recherche Nucléaire</i> [Chapter 1]
IT	Inner Tracker [Section 4.2.1]
LHC	Large Hadron Collider [Section 2.2]
LHCb	Large Hadron Collider beauty (experiment) [Chapter 1]
OT	Outer Tracker [Section 4.2.1]
PID	Particle identification [Section 4.2]
QCD	Quantum chromodynamics [Section 2.1]
RICH	Ring-imaging Cherenkov (detector) [Section 4.2]
SM	Standard Model (of particle physics) [Chapter 1]
TT	Tracker Turicensis [Section 4.2]
VELO	Vertex Locator [Section 4.2]

The variables appearing across several chapters of the main text are defined below. Their places of first occurrence are given as well.

\sqrt{s}	Centre-of-mass energy [Chapter 1]
$\sqrt{s_{\text{NN}}}$	Nucleon–nucleon centre-of-mass energy [Chapter 1]
η	Pseudorapidity (in the centre-of-mass system) [Section 2.2]
p_{T}	Transverse momentum [Section 2.2]
X_{max}	Depth of the shower maximum [Section 3.1]
N_{μ}	Number of muons produced in a shower [Section 3.1]
R	Ratio of the mean electromagnetic- and hadronic-energy pseudorapidity densities [Section 3.2]
η_{lab}	Pseudorapidity in the laboratory system [Section 4.2]
P_{fake}	Fake-track probability [Section 5.3]
χ^2	Sum of the squared normalised residuals [Section 5.4.3]
n_{DOF}	Number of degrees of freedom [Section 5.4.3]

1 Introduction

Cosmic rays, travelling to Earth from the vastness of the universe, reach the highest particle energies ever observed. In the middle of the 20th century, studies of cosmic rays were performed by placing photographic emulsions at high altitude, which led to the discovery of the charged pion [1]. This constitutes an early example of an astroparticle experiment contributing to research in particle physics. Due to its lightness, the pion is frequently produced in experiments at particle colliders, built in the subsequent decades. Although the colliders continued to become more powerful, the highest centre-of-mass energies nowadays reachable still lie an order of magnitude below that for cosmic rays [2], which can exceed 100 TeV. Collider experiments enabled the discovery of further particles. At the same time, theories were developed, predicting or accommodating these particles.

Numerous measurements in particle physics have consolidated the present theoretical description of this field, the Standard Model (SM) [3], since many years. There are nevertheless various observations that hint at physics beyond the SM. In searches for deviations from the SM conducted at high-energy colliders, a detailed understanding of the existing background processes is indispensable. The majority of these processes cannot be calculated from first principles in the SM. Therefore, phenomenological models of the interactions causing the background are needed.

Measurements in astroparticle physics allow knowledge of remote places in the universe to be acquired. Modern studies of cosmic rays are one possibility of accessing the physics prevailing in those places [4]. They can help to identify sources and acceleration mechanisms of high-energy cosmic rays, which are unknown to date. In the high-energy regime, cosmic rays can only be measured indirectly via the extensive particle cascades that they induce in the atmosphere of Earth, referred to as air showers. Inferring the primary particle or ion from the imprint of an air shower at ground requires a model of the interactions occurring in the atmosphere.

The preceding two paragraphs demonstrate that models of interactions are relevant both at the smallest scale in particle physics and at the largest scale in astroparticle physics. Since most of the particles involved in such interactions are hadrons, the models describing them are termed hadronic-interaction models. Each of those models possesses a set of parameters that depend on input from experiments [5]. Precision measurements of the production of SM particles at colliders, *e.g.*, contribute to this input.

Employing hadronic-interaction models to interpret air-shower measurements at high energies reveals a discrepancy. The mass composition of showers can be estimated in two ways. Cross-checking them indicates that the number of muons produced in air showers is too small in the models compared to observations. This discrepancy, which is already long-standing, is referred to as the Muon Puzzle [2]. One reason for shortcomings of the models is a lack of experimental input on the production of light hadrons at the TeV energy scale. In particular, the production in the forward region is important, *i.e.* under a small angle relative to the direction of the primary particle. Most interaction products in showers are emitted into this region, which many detectors of collider experiments do not cover. Moreover, the experimental input needs to focus on prompt production, signifying that the particles are produced in temporal and spatial vicinity to the first interaction, where most energy is transferred.

The Large Hadron Collider beauty (LHCb) experiment [6, 7], located at the European Organization for Nuclear Research (CERN, originally *Conseil Européen pour la Recherche Nucléaire*), is especially suited to deliver that input as the detector is a fully instrumented spectrometer covering the forward region. An experimentally accessible proxy for light hadrons is long-lived charged particles, with long-livedness denoting a timescale of picoseconds. Prompt production of such particles in proton-proton (p - p) collisions at a centre-of-mass energy of $\sqrt{s} = 13$ TeV has been studied in the first measurement performed within this thesis. The measurement consists of a series of analysis steps, which are required to translate the detector-level information into a physics quantity, the differential production cross-section. Additional insight can be gained from an extension of that measurement. Instead of a combined analysis of long-lived charged particles, the three most abundant hadrons thereof can be measured separately. Furthermore, collisions of protons with heavier ions can be investigated. This has been done in the second measurement performed in the course of this thesis, in which prompt production is analysed of charged pions, charged kaons and (anti)protons in collisions of protons with lead ions (p -Pb collisions) at a nucleon-nucleon centre-of-mass energy of $\sqrt{s_{NN}} = 8.16$ TeV and in p - p collisions at $\sqrt{s} = 13$ TeV. On a mass scale, the two collision systems enclose the dominant type of first interaction in the atmosphere, *i.e.* protons colliding with nitrogen or oxygen. Hence, the two measurements conducted are capable of constraining those processes and essential for the validation and improvement of hadronic-interaction models.

The author's contribution The main work of the author of this thesis was dedicated to data analysis and resulted in the two measurements presented in Chapters 5 and 6. These measurements were performed within the LHCb collaboration, a large international group of primarily physicists, engineers and technicians who operate the LHCb experiment. It would be impossible to conduct a measurement without

the effort of many collaborators that precedes the data analysis, for example the detector construction and commissioning as well as the development and maintenance of the software framework. During the data analysis, the author worked closely together with Dr. habil. Hans Dembinski and Prof. Dr. Michael Schmelling. In the first measurement, that of charged-particle production, the author was responsible for several analysis steps, wrote the paper describing the analysis and guided the paper through the review process. The paper was published in *Journal of High Energy Physics* [8], which is a peer-reviewed journal. The list of authors contains almost a thousand people from the LHCb collaboration, but only Hans Dembinski and the author of this thesis will serve as contacts in case of future enquiries. The author presented the measurement on behalf of the LHCb collaboration at two international conferences: the *16th International Workshop on Meson Physics* and the *12th International Workshop on Multiple Partonic Interactions at the LHC*. In the second measurement, that of identified-hadron production, the author was responsible for the entire analysis, was the sole developer of the analysis code and wrote comprehensive LHCb-internal documentation. At the time of finalising this thesis, the measurement is approved by the corresponding working group within the LHCb collaboration and is therefore in an advanced state of internal review. Michael Schmelling supported the two measurements by his scientific advice. The internal documentation of the two measurements is available to the LHCb collaboration in the LHCb-ANA-2020-007 and LHCb-ANA-2024-022 analysis notes. The analysis code, automated via a workflow engine, is centrally preserved. In addition, the author interpreted the second measurement with regard to the Muon Puzzle and contributed to the preservation of other LHCb measurements in the HEPData and RIVET systems, which is detailed in Chapter 7.

Outline This thesis is structured as follows. The theoretical and experimental foundations are presented in Part I. Firstly, the SM of particle physics and theory concepts of hadron production are described in Chapter 2. Properties and measurements of cosmic-ray-induced air showers as well as the Muon Puzzle in astroparticle physics are then introduced in Chapter 3; moreover, the connection to particle physics is established. Afterwards, in Chapter 4, the LHCb experiment is detailed. In Part II, which contains Chapters 5 and 6, the two measurements performed within this thesis are expounded. Implications of the measurements are lastly explained in Part III. This comprises an interpretation concerning the validation of hadronic-interaction models, discussed in Chapter 7, as well as a conclusion and an outlook, provided in Chapter 8.

Part I

Foundations

2 Theoretical framework

In this chapter, the theoretical basis of the work presented is laid. An overview of the elementary particles and forces in the SM is given first. The focus is then placed on the production of hadrons at small momentum transfer, which is relevant both at colliders and in air showers, and on the implementation of this production in various hadronic-interaction models.

2.1 The Standard Model of particle physics

The most comprehensive description of the elementary particles and three of the four fundamental forces known to date is the SM of particle physics [3]. The development of its current form was already finalised in the 1970s [9]. Since then, the discoveries of all elementary particles predicted by the SM, with the last one reported in 2012 [10, 11], have supported this theory.

The SM comprises 17 elementary particle species [3]. These can be separated into twelve fermions, which possess half-integer spin of $1/2$, and five bosons, possessing integer spin of 0 or 1. The fermions are divided into six quarks, which are affected by the strong force, and six leptons, which are not. Both quarks and leptons are ordered in three generations according to their mass. The *up* and *down* quarks as well as the electron neutrino and the electron form the first generation, that of the lightest particles. The second generation contains the *charm* and *strange* quarks as well as the muon neutrino and the muon, whereas the third generation consists of the *top* and *bottom* (or *beauty*) quarks as well as the tau neutrino and the tau. Neutrinos are assumed to be massless in the SM, but it was observed that the masses of at least two of the three neutrino states differ from zero [12]. The four spin-1 bosons mediate the three fundamental forces; the massless gluon is responsible for the strong force, the massless photon mediates the electromagnetic force, and the massive *W* and *Z* bosons are responsible for the weak force. The only spin-0 boson in the SM is the massive Higgs boson. It corresponds to the excitation of the Higgs field, which enables the generation of the masses of the *W* and *Z* bosons and ensures that the photon remains massless. This mechanism of electroweak symmetry breaking also gives rise to the fermion masses via a Yukawa interaction with the Higgs field.

The various particle species have different generalised charges, to which the force mediators couple [3]. The gluon couples to the colour charge, carried by the quarks

and the gluon itself. The photon couples to the electric charge, which the W boson and the fermions except for the neutrinos possess. The W boson couples to the weak isospin, carried by left-handed fermions and the W boson itself. The Z -boson coupling is determined by a combination of the weak isospin and the electric charge, which is nonzero for the W boson and both left- and right-handed fermions. Right-handed neutrinos do not exist in the SM. The concept of left- and right-handed fermions signifies that the relative configuration of the momentum and spin vectors of a fermion affects its behaviour under the weak force, which is thus a chiral theory.

Each of the fermions has a corresponding antiparticle with the same mass and spin, but opposite-sign charges [3]. The W boson occurs as particle–antiparticle pair of W^+ and W^- as well. The gluon, the photon and the Z and Higgs bosons are their own antiparticles. Due to the structure of quantum chromodynamics (QCD), the field theory of the strong force, eight types of gluons exist with different combinations of one colour and one anticolour. The chiral nature of the weak force implies that the W boson also couples to right-handed antifermions, but not to left-handed antifermions or right-handed fermions.

Quarks only occur in colour-neutral bound states unless the energy is sufficiently high [3]. This is referred to as confinement, and the states formed are termed hadrons. Mesons are hadrons containing one valence quark and one valence antiquark, while baryons consist of three valence quarks or antiquarks. Examples of mesons are charged pions (π^\pm) and charged kaons (K^\pm), which are the lightest electrically charged mesons; (anti)protons (p and \bar{p}) are the lightest baryons. Light hadrons are those that only have up, down and/or strange quarks in their valence-quark content. The valence quarks determine the charges of a hadron, but they are not its sole constituents. Besides the gluons, confining the valence quarks to the hadron, virtual quark–antiquark pairs are created from the vacuum according to the uncertainty principle and also annihilated on short scales of space and time. These virtual quarks contribute significantly to the mass of light hadrons. For a hadron to be classified as heavy, by contrast, at least one of its valence quarks needs to be a charm or beauty quark. Various exotic bound states with more than three valence quarks and/or antiquarks were discovered in recent years [13–15].

The SM, despite its success, is not able to describe several phenomena observed. One of them is neutrino oscillations, which forbids that all neutrino masses are zero as mentioned above. Moreover, the fourth fundamental force, gravitation, is not included. An overview of the shortcomings of the SM can be found in Reference [16].

2.2 Hadron production

The production of light hadrons in cosmic-ray-induced air showers proceeds via inelastic interactions of the arriving primary nucleus (which can be a proton) with

air nuclei in the atmosphere. At the TeV energy scale, hadron production in these collisions is dominated by processes with small momentum transfer between partons [17], which is the collective term for quarks and gluons. In this soft regime of QCD, perturbative calculations, *e.g.* of scattering cross-sections, cannot be performed since the coupling constant of the strong force is large. However, such predictions are needed to deduce the primary-nucleus type from the interaction products of a shower as there is no source of cosmic rays with a known mass composition, which could be used for calibration. The predictions are obtained from hadronic-interaction models, relying on effective field theories and phenomenology. Each model comprises a set of parameters, which control the physics processes implemented, determine the predictions and have to be constrained with the help of experimental results [5].

One such theory is Gribov–Regge theory [18, 19]. Elementary interactions are represented by the pomeron, which is a colour-neutral object that can be exchanged between partons. The pomeron was introduced to explain the approximately constant total cross-section in hadron–hadron collisions at high energies [20]. A major drawback of Gribov–Regge theory is that the calculations of particle production and cross-sections are inconsistent because energy conservation is not respected for cross-sections [21]. The only hadronic-interaction model treating this consistently [2] is the EPOS model [22]. Besides EPOS, the QGSJET [23] and SIBYLL [24] models are also based on Gribov–Regge theory. Not all of them use the same pomeron approach, which is detailed in the following three paragraphs. The models and their versions given below were validated and improved with early data from the Large Hadron Collider (LHC) at CERN [2], and they constitute the most recent versions widely employed nowadays by astroparticle-physics experiments to interpret their measurements [25].

In EPOS-LHC [26], a semihard pomeron is implemented. It unifies the original soft pomeron with aspects of perturbative QCD, describing hard processes between partons. These have large momentum transfer and precede the soft hadronisation process in a collision. In the semihard-pomeron approach, the colliding nuclei are represented by sets of various partons that each carry a large fraction of the initial momentum [27]. Since a large number of initial partons is involved in the subsequent process of multiple parton scattering, the production of hadrons under a small angle relative to the collision axis, *i.e.* in the forward pseudorapidity region, is strongly affected. Simultaneously, more hadrons are produced under a large angle relative to the collision axis, *i.e.* in the central pseudorapidity region, due to the large initial-parton sets, implying an increased amount of multiple scattering. Both processes cause long-range correlations in pseudorapidity, η , and the resulting mean density of light hadrons is rather wide in η . This is in agreement with measurements from LHCb and other LHC experiments [2, 28]. A comparison is shown in Figure 2.1. The edge around $|\eta| \approx 3$ in the densities stems from the kinematic requirements

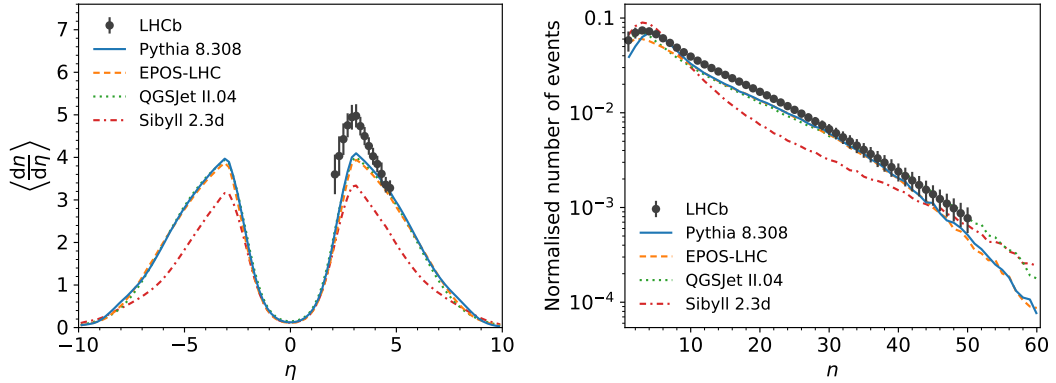


Figure 2.1: Impact of different approaches in the modelling of elementary interactions on properties of hadron production. The properties shown are the mean pseudorapidity density of prompt long-lived charged particles (left) and the distribution of their number per event (right). The predictions from the various hadronic-interaction models are generated for p - p collisions at $\sqrt{s} = 7$ TeV and compared with the data points [31] of a corresponding measurement from the LHCb experiment. The bars represent the sum in quadrature of statistical and systematic uncertainties. In the measurement, prompt long-lived charged particles with momentum above $2 \text{ GeV}/c$, transverse momentum above $0.2 \text{ GeV}/c$ and pseudorapidity in the $2.0 < \eta < 4.8$ range are studied, and an event is only retained if it produces at least one such particle. These requirements are also applied to the predictions, with that on η omitted for the left plot. The predicted distributions in the right plot are normalised to the number of generated inelastic events [32], which amounts to 10^6 .

applied. In EPOS, higher-order QCD effects are taken into account by an effective correction [22]. Furthermore, the model uses a core–corona approach to describe nucleus–nucleus collisions; the core is a state of dense matter undergoing collective effects, whereas the corona has lower energy density [29]. The goal of EPOS is a detailed understanding of particle production, which is by far not limited to the modelling of air-shower observables [30]. The EPOS-LHC version does not include heavy hadrons. EPOS can be employed to generate high-energy p - p , proton–nucleus and nucleus–nucleus collisions with nuclei up to lead ions.

The QGSJET II.04 model [33] also has a semihard pomeron implemented, inducing long-range correlations in η similar to EPOS. In contrast to EPOS, QGSJET captures higher-order QCD effects by an all-order resummation of pomeron–pomeron interactions [34]. As QGSJET is mainly designed to model air-shower physics, the particle production is not as comprehensive as in EPOS [28]; none of the light strange baryons is contained except for Λ and $\bar{\Lambda}$ baryons, nor are heavy hadrons. QGSJET also supports the collision systems given for EPOS.

In SIBYLL 2.3d [35], the soft and hard regimes are decoupled and respectively described by pomeron exchanges and minijets calculated with perturbative QCD. The minimum transverse momentum required for a minijet is below that of typical jets reconstructed at colliders [36]. The transverse momentum, p_T , is the component of the momentum perpendicular to the collision axis. SIBYLL can be regarded as a simplified version of the DPMJET model [28, 37]. In the soft-plus-hard-pomeron approach implemented, multiple parton scattering starts from a universal set of partons with a large initial-momentum fraction [27]. Hadrons produced in the forward η region originate from the beginning of the parton cascade and have little spread in η due to the parton-set universality. The remaining cascade causes hadrons to be mostly produced in the central η region. These two processes only weakly depend on each other; consequently, solely short-range correlations in η are present. The resulting mean density of light hadrons is rather narrow in η , too narrow compared to measurements [2, 28], which is also exemplified in Figure 2.1. The small width is related to the underestimation of the middle part of the distribution of the number of light hadrons, shown in the same figure. Like QGSJET, SIBYLL is designed to interpret air-shower data [35]. Although the model is less complex than EPOS and QGSJET, the predictions from SIBYLL reasonably reproduce many aspects of hadronic interactions [36]. A great advantage is its speed; SIBYLL can generate collisions much faster than QGSJET and especially EPOS. It comprises more particle species than QGSJET and EPOS-LHC, namely light strange baryons and charm hadrons. SIBYLL can be used to generate p - p and proton–nucleus collisions with nuclei up to around oxygen ions.

PYTHIA 8.308 [38], by contrast, is a model extensively used in high-energy particle physics. It is not based on Gribov–Regge theory. Instead, perturbative QCD is employed, not only for the hard parton processes, but also to compute the soft multiple parton scattering. Therefore, PYTHIA contains parameters to prevent the total cross-section at small momentum transfer from diverging. The PYTHIA model includes initial- and final-state radiation, through which additional soft gluons are created. They can split into quark–antiquark pairs and then hadronise into further particles, extending the range of possible η values. This implies longer-range correlations in η and a mean density of light hadrons that is wider in η compared to the soft-plus-hard-pomeron approach [2]. Several types of collective effects, altering the particle production in dense-matter environments, are part of PYTHIA [38]. The model allows the user to easily access and modify the values of its parameters. For this reason, PYTHIA does not only occur in various versions, like the other models, but also in numerous tunes, which are sets of values for the parameters. A tune can be optimised, *e.g.*, for the η coverage of an experiment [39]. The default tune of PYTHIA 8.308 is the Monash 2013 tune [40]. Since PYTHIA is meant to be an accurate general-purpose model [38], it can generate all known particle species, for example also beauty hadrons, unlike SIBYLL. PYTHIA supports p - p collisions among other small

collision systems, and the possibility of generating proton–nucleus and nucleus–nucleus collisions with nuclei up to lead ions was recently added [41].

Model predictions are obtained from collections of events, where an event refers to the ensemble of particles created in one collision. In nature, the outcome of collisions for given initial nuclei and a given energy differs from collision to collision owing to quantum fluctuations. For the models, this behaviour is mimicked by including them in event generators [38], which use random numbers to create distinct events. Simulation denotes the general process of generating events, while experimental results are referred to as data.

The authors of the models usually determine an optimal set of parameter values by modifying them until the predictions are sufficiently similar to various measurements from collider experiments. Those data comprise properties of hadronic interactions, such as the energy dependence of the inelastic cross-section and the distribution of the number of prompt long-lived charged particles [28, 42]. For the models dedicated to air-shower simulation, QGSJET and SIBYLL, the number of measurements considered is relatively small since the number of parameters is small, in the order of ten [2]. For the general-purpose models, EPOS and PYTHIA, a wider variety of data are taken into account to constrain the parameters, of which there are hundreds [38]. Although not originally designed to simulate air showers, PYTHIA was demonstrated to generate predictions comparable to those of the other three models [43]. Cross-checking predictions from the models with data not used to set the parameters reveals that there is disagreement in numerous places. This is partially because the experimental input does not cover the highest energies occurring in air showers, meaning that the models have to extrapolate to centre-of-mass energies that are one order of magnitude above those reachable at the LHC [2]. In addition, collider experiments often cannot measure in the forward η region, which would be important for air showers, or have not yet studied key processes, for example light-hadron production. Two observables in which the models differ have already been described above, the width of the mean density of light hadrons in η and the distribution of their number per event. Another more eminent discrepancy is the Muon Puzzle in high-energy air showers. As the hadronic-interaction models constitute the largest source of uncertainty in air-shower simulation [2], continuously validating and improving them is indispensable. The benefit of a better understanding of hadronic interactions is not restricted to astroparticle physics. In measurements performed at colliders, an accurate simulation of the underlying event is desirable, which refers to the particles stemming from the soft component of a collision. These particles form the background for many analyses of heavy hadrons and are particularly relevant when rare processes are studied.

3 Connection of particle and astroparticle physics

The main properties of cosmic rays and air showers are presented in this chapter. Modern air-shower detection is described by taking the Pierre Auger Observatory as an example. Moreover, a long-standing discrepancy with regard to high-energy air showers, the Muon Puzzle, is introduced. Its connection to the LHC is emphasised, and potential explanations are given.

3.1 Cosmic rays and air showers

High-energy nuclei originating from distant places in our galaxy or beyond regularly hit Earth. These fully ionised atoms range from protons to nuclei of elements enclosing iron and are referred to as cosmic rays [4]. The flux of cosmic rays approximately spans eleven orders of magnitude in kinetic energy, from below 1 GeV up to around 10^{11} GeV. It is remarkable that, despite this wide range, the spectrum can be described by a simple power law in each of only a few large energy intervals. There are two transition regions that are known since several decades [44, 45]. Above 3×10^6 GeV, the energy spectrum steepens, forming the *knee*. This is related to the transition from the probably dominant galactic mechanism, which is shock acceleration in supernova remnants, to another one. Above 3×10^9 GeV, the *ankle* marks a flattening in the spectrum, which is usually interpreted as the transition from galactic to extragalactic sources of these ultrahigh-energy cosmic rays. The third well-known feature of the spectrum is its suppression above 4×10^{10} GeV [46, 47], caused by energy loss of the cosmic rays during propagation through the cosmic microwave background. Neither the sources nor the acceleration mechanisms capable of producing ultrahigh-energy cosmic rays have been identified so far [4]. More recently, some subtle structures between the knee and the ankle were reported as well [48].

At energies below 10^5 GeV, the flux is sufficiently high to measure cosmic rays directly by experiments located at the top of the atmosphere, *e.g.* on balloons or satellites [4]. At higher energies, the flux is too low for an efficient direct detection. Instead, indirect detection of cosmic rays is performed with ground-based experiments that cover a large area and have a long exposure time. Such experiments

employ the atmosphere as calorimeter and measure products of the cosmic-ray-induced air showers.

An air shower consists of extensive hadronic and electromagnetic cascades of particles, which develop in a coupled way [4]. High-energy hadrons are produced in the first interaction, that of the primary nucleus with nitrogen or oxygen, and in subsequent interactions of the hadrons produced before. The majority of particles emerging from this prompt production is pions, but there are also other light hadrons, such as kaons and protons [2]. Nonprompt production occurs as well, for example via decays of kaons into pions, but appears less often. Charm hadrons can also be created, but their fraction is small. The energy contained in the hadronic cascade is steadily transferred to the electromagnetic cascade, approximately one third in each interaction. This amount is primarily due to neutral pions, which constitute one third of all pions produced because of isospin symmetry and almost always decay into two photons. The photons interact via pair production, giving rise to electrons (e^\pm), which can in turn create photons via bremsstrahlung. Pair production continues until the photon energy falls below the critical energy of around 1 MeV. Feedback from the electromagnetic to the hadronic cascade only happens to a small extent; one example is photonuclear interactions [2]. The hadronic cascade proceeds until the energy is small enough for decay to become more probable than reinteraction. This leads to the creation of muons (μ^\pm) and muon neutrinos (ν_μ and $\bar{\nu}_\mu$), *e.g.* in $\pi^+ \rightarrow \mu^+ \nu_\mu$ decays. A few muons decay into electrons and neutrinos. The electromagnetic contribution to the muon content of the shower is small since $\mu^+ \mu^-$ -pair production seldom occurs. Thus, the muons trace the evolution of the hadronic cascade, especially its low-energy end. The type of the primary nucleus has to be inferred from the particles produced, *i.e.* from air-shower observables. This is essential to deduce the mass composition of cosmic rays at a specific energy, enabling potential source scenarios of cosmic rays to be distinguished [49]. At energies above 10^6 GeV, the shower fluctuations caused by the processes described above can be larger between showers induced by the same primary nucleus than between showers induced by nuclei of adjacent elements [2]. This advocates using a single quantity to reflect the composition, which is usually chosen to be the mean logarithmic mass, $\langle \ln A \rangle$.

Insight into air-shower observables can be obtained based on the Heitler model of electromagnetic cascades [50], which was extended to hadronic cascades by Matthews [51]. It is assumed that the hadronic cascade only consists of pions. In each hadronic interaction with air nuclei, two thirds of the energy enter charged-pion production, whereas one third is converted into neutral pions. Let the fixed total number of pions produced in each step of the cascade be denoted by N_π . Each neutral pion decays immediately into two photons, interacting further via pair production and bremsstrahlung, with the energy equally shared among the products in each step. Charged pions reinteract after traversing a fixed amount of

slant depth until their energy falls below a critical energy, ζ_π . Slant depth, defined as the integral of the air density along the length crossed, is a measure of absorption. Lastly, each charged pion decays into one muon. The number of steps of the hadronic cascade can be calculated in the model; it is found to be 5 for a shower at 10^9 GeV if values of $N_\pi = 50$ and $\zeta_\pi = 10$ GeV are assumed [2]. Furthermore, the model allows the two main air-shower observables to be related to $\langle \ln A \rangle$: the depth, X_{\max} , of the shower maximum and the number, N_μ , of muons produced. It follows that, for a given energy, heavier-nucleus-induced showers comprise more muons than proton-induced showers, while the latter penetrate more deeply in the atmosphere. Relating N_μ to the fraction, α , of energy retained in charged pions, which amounts to exactly $2/3$ in the model, demonstrates that a change in α of only 10% already leads to a change in N_μ of 50% at 10^9 GeV [2]. This is because the effect accumulates along the steps of the cascade. It should be noted that the model is built on several approximations, discussed further in Reference [2], where refinements of the model are also outlined.

Two giant observatories exist in the world that measure air showers induced by ultrahigh-energy cosmic rays. The Pierre Auger Observatory (or simply *Auger*) [52] is the largest air-shower experiment to date. It approximately covers 3000 km^2 and is located in the Southern Hemisphere, in western Argentina. The area contains a sparse grid of 1660 water-Cherenkov-detector stations and is overlooked from four sites of six air-fluorescence telescopes each; three additional telescopes were installed after the completion of the baseline observatory [53]. This structure makes the Auger experiment a hybrid detector. The Cherenkov stations observe charged particles via the electromagnetic radiation emitted if the particles travel faster than the speed of light in water. High-energy photons can be detected as well if they produce e^+e^- pairs in the water volume. Consequently, the Cherenkov stations are sensitive to the electromagnetic cascade and the muon content of a shower. Relative timing permits the arrival direction of a shower to be determined, and integrating the lateral muon density at ground gives N_μ [2]. In addition, the energy of the primary nucleus can be reconstructed. The fluorescence telescopes observe the ultraviolet light emitted by nitrogen along the shower path. This light results from the charged particles in the shower, which excite the molecules. As the telescopes can only be operated in clear and moonless nights and when the wind conditions are safe, the duty cycle is limited to 15%. It also determines the fraction of hybrid events, measured simultaneously with both detection techniques, because the Cherenkov stations possess a duty cycle of 100%. The fluorescence telescopes measure the energy of the primary nucleus via integration of the longitudinal light profile. They are able to detect X_{\max} , which is the most direct estimator of the cosmic-ray mass composition. Hybrid detection provides measurement redundancy and therefore enables important cross-checks between techniques with different systematic uncertainties; the muon content, obtained from the Cherenkov stations, can also be used to estimate the mass

composition, whereas the fluorescence telescopes can detect the arrival direction as well. Moreover, the Auger experiment was upgraded by installing additional Cherenkov stations on a denser grid and buried plastic-scintillation-detector stations, with the underground location of the latter providing shielding against the particles from the electromagnetic cascade and hence a direct muon measurement [54]. The largest hybrid air-shower experiment in the Northern Hemisphere is Telescope Array [55, 56]. It is located in western Utah in the United States. The area covered by the sparsely spaced 507 plastic-scintillation-detector stations approximately amounts to 700 km². It is viewed from three sites of 38 air-fluorescence telescopes in total.

3.2 The Muon Puzzle in astroparticle physics

The advantages of hybrid air-shower detection were employed by the Auger collaboration in two measurements published in 2015 [57] and 2016 [58], which rekindled interest in the muon content of ultrahigh-energy air showers due to the comparison with hadronic-interaction models based on early LHC data. In Reference [57], highly inclined showers were analysed to infer the mass composition of cosmic rays from N_μ . Most particles of the electromagnetic cascade are absorbed by the large depth of the atmosphere traversed by inclined showers, leaving a high purity of muons at ground. The cosmic-ray energy is obtained independently of the muon detection by using the fluorescence telescopes. The mean muon content as a function of the energy for measured and simulated showers is shown in Figure 3.1. The separation of the proton and iron predictions from EPOS-LHC and QGSJET II.04 demonstrates the power of the N_μ -based mass-composition estimator. However, this power is limited by the large systematic uncertainty, primarily caused by the uncertainty on the absolute energy scale of the Auger experiment. Since the measured muon content exceeds that predicted for pure iron-induced showers, the data hint at contributions from elements even heavier than iron, which grow with increasing energy. This interpretation is not in accord with X_{\max} -based mass-composition estimates, plausibly residing between proton and iron [59]. Therefore, a direct comparison of the two estimators is useful, which allows the internal consistency of the models to be investigated. That comparison is shown in the same figure. None of the models is compatible with the measurement within the sum in quadrature of statistical and systematic uncertainties, which is dominated by the systematic contribution. The discrepancy is larger for the two older versions of QGSJET, not relying on LHC data. To establish agreement, the simulated mean number of muons would have to be increased by 30 % to 80 %, depending on the model, or the experimental energy scale would have to be shifted upwards by the same factor. In Reference [58], hadronic-interaction models were tested with only minimal reliance on the Auger

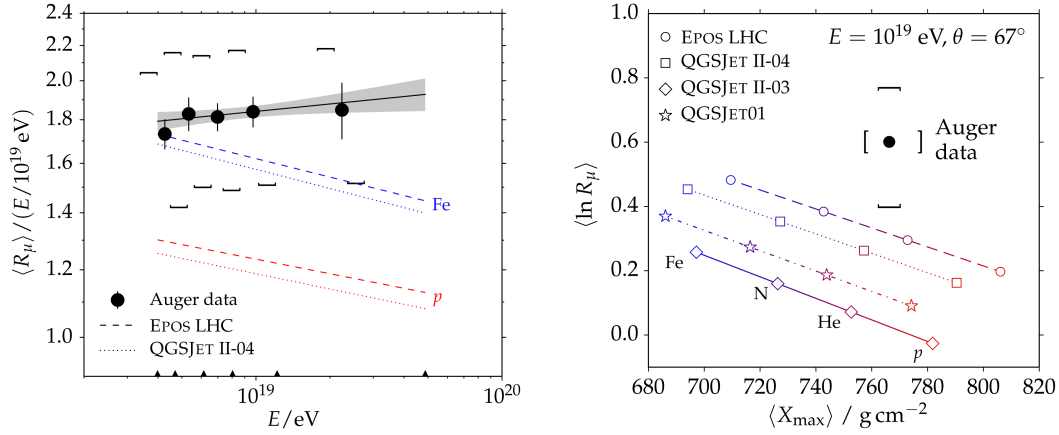


Figure 3.1: Deficit of muons in simulated ultrahigh-energy air showers compared to measured showers as reported by the Pierre Auger Observatory [57]. The mean shower muon content as a function of the shower energy (left) and the mean logarithmic shower muon content as a function of the mean depth of the shower maximum (right) are shown. The bars and brackets represent the statistical and systematic uncertainties, respectively. The black solid line and its statistical-uncertainty band in the left plot are determined by a fit to the data. The predictions from the hadronic-interaction models are obtained by simulating showers with the mean zenith angle of $\theta = 67^\circ$.

energy scale. The hadronic cascade in simulated air showers is found to be significantly smaller than in data, which corresponds to a muon deficit in simulation and matches the 2015 result. Already in 2000, the HiRes-MIA hybrid experiment had reported a similar muon-content discrepancy [60].

The findings from the Auger collaboration encouraged other air-shower experiments to perform new measurements focused on $N_{\mu'}$, which in turn necessitated a meta-analysis of all the results obtained. For this purpose, a dedicated working group with members from the corresponding experiments was formed [61, 62]. Two main analysis steps were undertaken. First, an abstract muon-content quantity, z , was defined by using simulated N_μ values for pure proton- and iron-induced showers to take into account the variety of data-recording conditions. It permits the energy dependence of an N_μ measurement to be removed, and linear biases in the measured value of $\ln N_\mu$ compared to the real value cancel. By construction, the z value depends on the hadronic-interaction model considered. In the case of no discrepancy between data and simulation, the range covered by z is between zero and unity, signifying pure proton- and iron-induced showers, respectively. Second, a cross-calibration of the experimental energy scales was applied. It is based on the cosmic-ray flux, which is isotropic up to high energies and can hence serve as a reference, as well as on the assumption that deviations in the measured flux

between different experiments are only caused by offsets in their energy scales. After cross-calibration, there can still be a global offset of all measurements relative to the real cosmic-ray energy scale, estimated to be ± 0.25 in z . The first result of the meta-analysis was published in References [61, 62]. A relatively uniform picture was obtained from the eight experiments included; almost all z values rose with energy and commenced to exceed unity around 10^8 GeV, consistently reflecting a muon deficit in simulation. The significance of this energy-dependent increase was studied by linear fits to the z values for EPOS-LHC and QGSJET II.04 after subtracting effects of the cosmic-ray mass composition, which is also energy dependent. For multiple values of correlation assumed between the systematic uncertainties on the data points from a single experiment, the increase was found to be highly significant; a constant behaviour in energy was excluded with at least eight standard deviations. This discrepancy between air-shower data and simulation was established as the Muon Puzzle, extensively reviewed in Reference [2]. An update of the meta-analysis was published in 2021 [63], which included data from one additional experiment [64] and new data from experiments already considered. The significance of the energy-dependent increase in the observed muon content was investigated in more detail, with similar conclusions reached as before. However, the picture has become partially discordant with the publication of the most recent meta-analysis in 2023 [65]. It is shown in Figure 3.2. For the first time, data from the Haverah Park experiment are included, which was operated in northern England from 1967 to 1987 [66]. Updated measurements from the Yakutsk and SUGAR experiments are used. Furthermore, applying the energy-scale cross-calibration has been made possible for the KASCADE-Grande experiment. Three groups of data can be identified. The Auger, Telescope Array, NEVOD-DECOR, SUGAR and AGASA measurements exhibit z values rising above the iron expectation at high energies. The IceCube, Yakutsk and Haverah Park measurements lie within the expected z range for the most recent hadronic-interaction models; unlike those from Yakutsk and Haverah Park, the IceCube measurement increases with energy. The KASCADE-Grande measurement decreases with energy and even tends to indicate a composition lighter than pure proton-induced showers for EPOS-LHC and Sibyll 2.3d. The EAS-MSU and HiRes-MIA measurements cannot be regarded as comparable because both results are missing information that would be required for the energy-scale cross-calibration; the HiRes-MIA result also differs systematically from all others. In summary, some experiments measure a shower muon content larger than that expected from the models, whereas others measure a muon content agreeing with the models; one measurement suggests fewer muons than the models. A conclusive interpretation is still to be provided. It is remarkable, for example, that one of the Auger results and that from Haverah Park employ the same muon detection technique, shielded scintillators, but they disagree [66]. The Haverah Park data, however, are not based on a full detector simulation. In addition, the

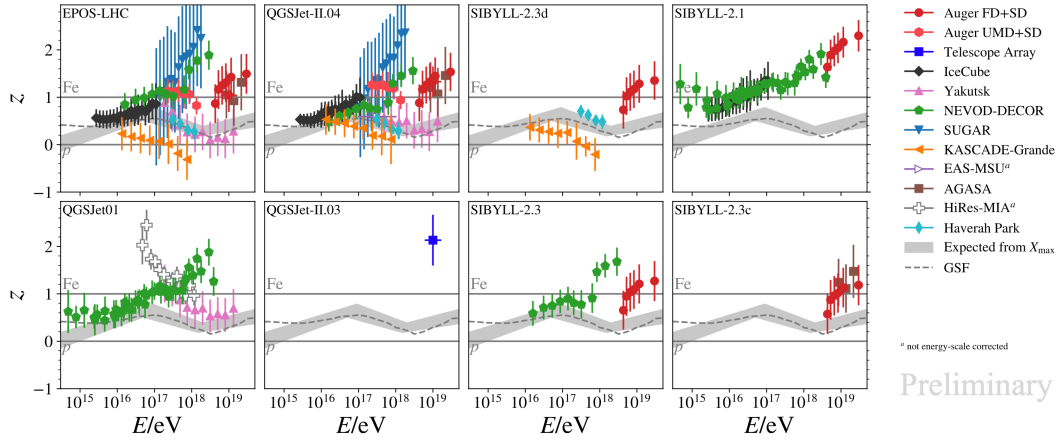


Figure 3.2: Muon content of air showers encoded in the z quantity explained in the main text as a function of the shower energy for measurements from various experiments [65]. The z values shown include the energy-scale cross-calibration. Due to the dependence of z on a hadronic-interaction model, the same data can appear several times for different models, or no data points can be present for a given model. The bars indicate the sum in quadrature of statistical and systematic uncertainties. Two results from the Pierre Auger Observatory are shown; one of them is based on the fluorescence detector and the surface detector, formed by the Cherenkov stations, while the other uses the underground muon detector and the surface detector. No cross-calibration was possible for the EAS-MSU and HiRes-MIA results because of missing information on the data, but the results are nevertheless shown for completeness. The grey band represents the z values expected from X_{\max} measurements from several experiments for a mixed cosmic-ray mass composition. The grey dashed line is obtained from a global spline fit to the cosmic-ray energy spectrum.

dissimilarity of all the measurements might be related to the fraction of muons contained in their energy estimator since most of the experiments observing the Muon Puzzle either measure the energy with only a small muon contribution or estimate it independently of muons via the cosmic-ray flux [2, 65]. Further studies are definitely needed.

Besides the current lack of a uniform picture concerning the Muon Puzzle, the large uncertainties on the measurements are unsatisfactory. In several cases, the uncertainty is close to or greater than half the difference between the proton and iron expectations in z , diluting the meaningfulness of the data. Almost all data points are dominated by the systematic uncertainty, of which a considerable fraction originates from the hadronic-interaction models used in the air-shower simulation [2] as mentioned in Section 2.2. Other parts of the simulation, such as the particle transport in the atmosphere, and experimental uncertainties contribute much less

to the systematic uncertainty. Extending the input from collider experiments to the models can aid in reducing their uncertainties and would ideally also explain the Muon Puzzle. A fortunate coincidence is that the energy at which the Muon Puzzle starts to appear, around a cosmic-ray energy of 10^8 GeV in the laboratory system, corresponds to an energy in the order of 10 TeV in the centre-of-mass system. Such energies are reachable in collisions at the LHC. The phenomenon that might clarify the Muon Puzzle is not expected to be related to physics beyond the SM, which would mainly affect the high-energy first interaction in the atmosphere. This can be concluded from an Auger measurement of the fluctuations of the shower muon content, which are primarily due to the first interaction and were observed to be consistent with the model predictions [25]. It is therefore more likely that the effect (or effects) responsible for the Muon Puzzle accumulates along the steps of the hadronic cascade of a shower. Thus, a rather small effect might be sufficient to modify the mean number of muons significantly. The most promising idea is that the effect can be attributed to the soft regime of QCD and is relevant only at high energies, where air-shower data and simulation disagree [2].

To determine which measurements at the LHC are of key importance for shedding light on the Muon Puzzle, it is instructive to study the impact of properties of hadronic interactions on the mean number of muons produced in an air shower. This was done in Reference [67]. The values of the inelastic cross-section, the secondary-particle multiplicity, the neutral-pion fraction relative to all pions and the elasticity present in the hadronic-interaction models were modified with an energy-dependent factor. Multiplicity signifies the total number of particles produced, while the fraction of energy carried by the most-energetic particle after the interaction is referred to as elasticity. The modifications were not applied based on the model theory or phenomenology, but as free variations around the default values. This enabled the dependence of the mean values of X_{\max} and N_{μ} as well as of the fluctuations of X_{\max} and N_{μ} on the four quantities modified to be investigated in detail. It was found that the mean value of N_{μ} strongly depends on the secondary-particle multiplicity and the neutral-pion fraction. Modifications of these two quantities can only be a potential solution to the Muon Puzzle if their impact on the mean value of X_{\max} and the fluctuations of X_{\max} and N_{μ} is small. As described above, the cosmic-ray mass composition derived from X_{\max} is plausible, and the measured fluctuations of N_{μ} are reproduced by the models. The fluctuations of X_{\max} measured by the Auger experiment are also mostly enclosed by the proton and iron expectations [59]. For the fluctuations of X_{\max} and N_{μ} , the impact is found to be small. The mean value of X_{\max} , however, is affected; it changes more with the secondary-particle multiplicity than with the neutral-pion fraction, favouring the latter as key to the Muon Puzzle. The Auger collaboration recently reported that, besides the production of more muons, a moderate increase in the mean values of X_{\max} in the models improves their compatibility with the data [68]. In Reference [69], the impact of modified

properties of hadronic interactions both on X_{\max} and N_{μ} was studied simultaneously. The neutral-pion fraction is not easily measurable at the LHC and was hence replaced by an analogous quantity,

$$R = \left\langle \frac{dE_{\text{EM}}}{d\eta} \right\rangle / \left\langle \frac{dE_{\text{hadr}}}{d\eta} \right\rangle, \quad (3.1)$$

defined as the ratio of the mean electromagnetic- and hadronic-energy pseudorapidity densities. These densities can be obtained by adding in each event and in intervals of η the energies of all corresponding final-state particles, photons and electrons in the former case and long-lived hadrons in the latter case, and averaging over many events. In the Matthews–Heitler model, there is a simple relation between the fraction of energy retained in charged pions as introduced in Section 3.1, the neutral-pion fraction, $1 - \alpha$, and the energy ratio: $R = (1 - \alpha)/\alpha = 1/2$. It was deduced that modifying the secondary-particle multiplicity is not able to compensate the muon deficit in simulation as that only shifts the predictions in the X_{\max} – N_{μ} plane, which has been shown in Figure 3.1, along themselves. Solely modifications of the energy ratio are capable of shifting the predictions along the muon-content axis. In particular, a decrease in R by approximately 15 % is needed; sufficient energy would then be retained in the hadronic cascade of a shower, muon production in simulation would increase, and agreement with the Auger data at 10^{10} GeV would be achieved.

Modifying the energy ratio implies changes in the hadron composition because producing fewer neutral pions and thus photons is coupled to producing fewer charged pions via isospin symmetry. This has to lead to an increased production of other long-lived hadrons, which are mostly strange hadrons except for two light baryons, the proton and the neutron. In 2017, a surprising measurement from the ALICE experiment at the LHC was published [70], in which exactly the effect described was discovered. The production of K_S^0 mesons as well as of Λ , $\bar{\Lambda}$, Ξ^- , $\bar{\Xi}^+$, Ω^- and $\bar{\Omega}^+$ baryons in p – p collisions was found to be significantly enhanced relative to charged pions at high multiplicity. The enhancement becomes more pronounced as the number of valence strange quarks increases from one to three and is therefore identified as strangeness enhancement. In this case, multiplicity is represented by the mean pseudorapidity density of prompt long-lived charged particles in the central η region. Strangeness enhancement is not a novel phenomenon; it was already measured in nucleus–nucleus collisions in pre-LHC times as described, *e.g.*, in References [71, 72]. Enhanced production of strange hadrons in nucleus–nucleus collisions was originally interpreted as a sign of the formation of quark–gluon plasma [73, 74]. In this state of matter, emerging at high energy density and temperature, partons are not confined to hadrons; more pairs of strange quarks and antiquarks (possessing a higher mass than up and down quarks) are expected to be formed due to the abundance of high-energy gluons. Consequently, the hadron composition is altered.

The discovery of multiplicity-dependent strangeness enhancement in p - p collisions was a surprise since a sign of quark–gluon plasma was not anticipated in such a small collision system [2]. Another important observation was that the strangeness enhancement measured in p - p collisions closely resembles that measured in p -Pb [75, 76] and Pb–Pb [77] collisions also by the ALICE collaboration. Thus, it seems to be universal across collision systems and likewise across centre-of-mass energies [78]. A dependence only on the mean pseudorapidity density would result in substantial predictive power. With regard to the Muon Puzzle, it should be emphasised that the ALICE measurements were performed in the central η region, which is not directly relevant for air showers. The LHCb experiment, by contrast, covers the forward η region and also discovered multiplicity-dependent strangeness enhancement [79]. In the measurement, the production of strange charm mesons was found to be enhanced relative to that of nonstrange charm mesons in high-multiplicity p -Pb collisions. This requires an increased presence of strange quarks, with which charm quarks then hadronise, and could also be indicative of quark–gluon plasma in a system smaller than that of nucleus–nucleus collisions. Moreover, evidence was established by the LHCb collaboration of enhanced production of strange beauty mesons compared to nonstrange beauty mesons in high-multiplicity p - p collisions [80], with similar implications. It is not definite if the origin of the strangeness enhancement measured in p - p and p -Pb collisions lies in the formation of quark–gluon plasma. Nevertheless, such a plasma is assumed to be formed in the core region of a collision in the EPOS model. Another mechanism that is able to change the hadron composition and enhance strangeness is string–string interactions [2]. A string [81, 82] represents the field of the strong force between two colour charges as a flux tube that is colour neutral and composed of virtual quarks and gluons. With increasing distance of the colour charges, the energy in the string rises, and the string can split while converting energy into quark–antiquark pairs. One type of string–string interaction is colour reconnection [83], which enhances the production of baryons at high multiplicity, but not that of strangeness [84]. Both can be achieved by assuming rope hadronisation [85]. A rope is formed if strings overlap in a dense-matter environment. This leads to a higher energy in the combined rope and increases the probability of creating strange quark–antiquark pairs and diquarks in string splits, with the latter giving rise to baryons. String shoving [86] is another type of string–string interaction, which might prevent enhanced production in high-multiplicity nucleus–nucleus collisions from becoming too large [87]. All of these three concepts are available in the PYTHIA model.

Although the measurements described above do not quantify the energy ratio, they demonstrate that changes in the hadron composition, needed to explain the Muon Puzzle, occur at LHC energies and even in p - p collisions. A direct measurement of R in the forward η region of p - p collisions was performed with the CMS experiment at the LHC [88]. The measured energy ratio was found to be larger

than that predicted by hadronic-interaction models, which seems to contradict the decrease in R that would compensate the muon deficit in the models. Furthermore, the energy ratio was measured to be almost independent of multiplicity, whereas strangeness enhancement is expected to force R to fall with multiplicity. However, the large systematic uncertainty on the measurement precludes a firm interpretation. Hence, more measurements providing insight into R have to be conducted in the forward η region at the highest energies reachable at the LHC. One possibility of accessing the hadron composition is to measure the production of charged pions, charged kaons and (anti)protons, constituting the most abundant hadrons in air showers. To investigate the potential evolution of composition-changing effects with η , it is advisable to perform an analysis in several intervals of this variable. The transverse momentum is suited as second independent variable as it reflects the urge of a particle to depart from the shower axis. Demanding the particle kinematics to be measurable in the forward η region restricts the range of eligible LHC experiments to LHCb. Still, many particles in a shower are produced even more forward. Another unique advantage of the LHCb experiment is its capability to distinguish different species of long-lived charged particles produced in the forward η region. Collisions of protons with nitrogen or oxygen ions, mimicking the dominant type of first interaction in the atmosphere, have not yet been realised at the LHC. Therefore, collisions with nuclei enclosing these ions on a mass scale, in particular p - p and p -Pb collisions, can be studied to set constraints. Interestingly, the mass numbers of a proton, nitrogen and lead reside equidistantly on a logarithmic scale, the scale over which the physics usually varies [2]. These considerations have resulted in the two measurements performed within this thesis. Their relevance is also stressed in Reference [2].

4 Experimental setup

In this chapter, the LHC is presented first, which has provided the collisions analysed. Then, the LHCb detector is described, and its subdetectors relevant for the two measurements of this thesis are introduced.

4.1 The Large Hadron Collider

The most powerful particle accelerator built to date is the LHC [89]. It was designed for centre-of-mass energies up to 14 TeV, and an energy of 13.6 TeV has been reached so far. The accelerator is a synchrotron with a circumference of 26.7 km. It is located between 45 m and 170 m underground at CERN, near the Franco-Swiss border in the Geneva area. The LHC comprises two evacuated beam pipes, in which protons or ions can travel in opposite directions. The bending of the proton or ion trajectories along the ring orbit is achieved by superconducting dipole magnets, which are operated at a temperature of 1.9 K and of which there are 1232. Further magnets are present to focus the beams and correct their orbits. The beams consist of numerous bunches, which are arranged in several groups. These groups of bunches, referred to as bunch trains, are separately injected into the LHC. The bunches in a bunch train possess a time spacing of at least 25 ns. In the case of proton acceleration, one bunch contains approximately 10^{11} particles, and the maximum number of bunches that can circulate amounts to 2808 per beam. The LHC constitutes the last stage in a chain of accelerators [90]. The protons or ions first pass through a linear accelerator before successively gaining more energy in three other synchrotrons. The bunch trains created are then injected into the LHC beam pipes via two transfer lines, with an injection energy of 450 GeV for protons. Afterwards, the energy is further increased and lastly kept constant.

At four points along the LHC, the orbits of the two beams cross, enabling the protons or ions to be collided. Each of these interaction points is located in a cavern constructed to accommodate a large particle detector measuring the products of the collisions. In the caverns, the four main experiments at the LHC are operated: ALICE, ATLAS, CMS and LHCb. The ALICE experiment [91, 92] was designed to study quark–gluon plasma in collisions of heavy ions. The ATLAS [93, 94] and CMS [95, 96] experiments were built to search for new heavy particles, which included the Higgs boson at that time. The detectors of those three experiments are instrumented

around the corresponding interaction point, whereas the LHCb detector primarily covers one side of the interaction point, with its subdetectors residing along the beam axis. The LHCb experiment is explained in detail in Section 4.2.

The LHC first provided collisions to the experiments in late 2009 [90]. Many measurements surpassing the original purposes of the experiments have been performed since then, *e.g.* of soft QCD processes. The majority of the beam time in the years of LHC operation is dedicated to p - p collisions. Usually towards the end of those years, $^{208}\text{Pb}^{82+}$ ions are accelerated, allowing p -Pb and Pb-Pb collisions to be obtained. Another type of ion was once accelerated in 2017, namely xenon ions, which were collided for a few hours [97]. Currently, it is planned to accelerate oxygen ions for the first time in 2025.

4.2 The LHCb detector

The detector of the LHCb experiment [6, 7] covers the $2 < \eta_{\text{lab}} < 5$ range, where η_{lab} denotes the pseudorapidity in the laboratory system. Due to this setup, the detector is a forward spectrometer. Its design was chosen based on the original physics goals of the experiment, *i.e.* precision measurements based on decays of beauty and charm hadrons. At LHC energies, most pairs of heavy quarks and antiquarks produced in collisions are emitted under a small angle relative to the beam axis. This is caused by gluon-gluon fusion, which is the dominant production mechanism [98]. In the most likely case, one gluon carries a large fraction of the longitudinal momentum, whereas the other possesses a much smaller momentum fraction, inducing a boost along the beam axis. The hadrons forming from these heavy quarks inherit the boost. For the LHCb detector, only one of the two cones of increased production was instrumented because of cost reasons.

The LHCb detector is visualised in Figure 4.1. It has a length of approximately 20 m, a height of 10 m and a width of 12 m. The collisions happen inside the Vertex Locator (VELO), which permits the point of the first interaction, the primary vertex, to be measured. The trajectories of charged particles produced in a collision can be detected by the Tracker Turicensis (TT), placed upstream of a dipole magnet, and by three tracking stations downstream of the magnet. These subdetectors form the track-reconstruction system. The magnetic field is either parallel or antiparallel to the y axis of the LHCb coordinate system, indicated in the figure. Consequently, the magnet bends the paths of charged particles in the x - z plane. The species of charged particles are mainly distinguished based on information from two ring-imaging Cherenkov (RICH) detectors. Additional separation power is provided by a calorimeter system, which is composed of a scintillating-pad detector, a preshower detector, an electromagnetic calorimeter and a hadronic calorimeter. Dedicated stations for muon identification are located at the downstream end of the detector. All

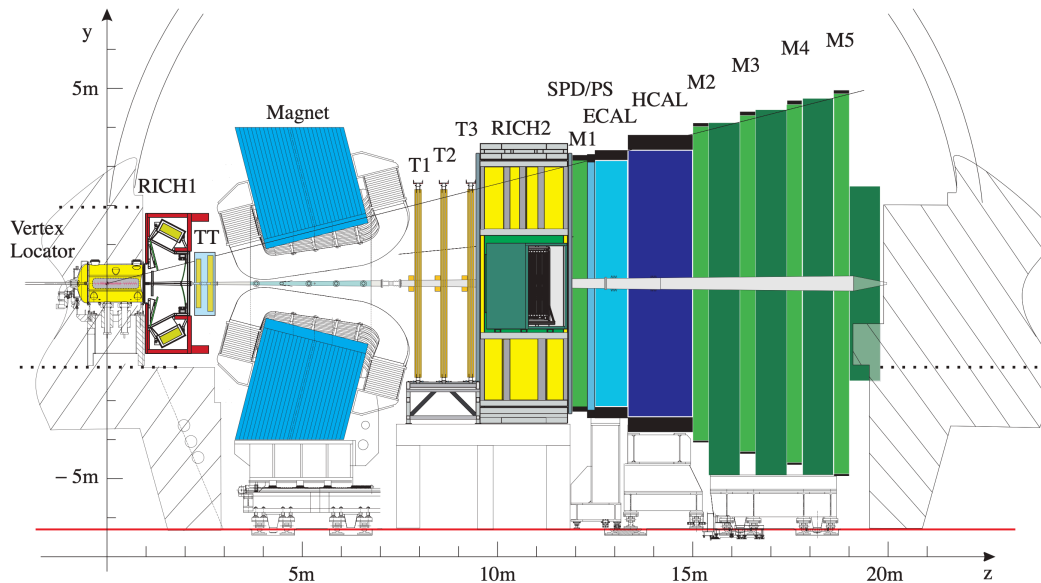


Figure 4.1: Schematic side view of the LHCb detector for an observer looking towards the outside of the LHC ring [99]. The following abbreviations are not used in the main text: T1–T3 for the tracking stations downstream of the magnet, M1–M5 for the muon stations, SPD for the scintillating-pad detector, PS for the preshower detector, ECAL for the electromagnetic calorimeter and HCAL for the hadronic calorimeter. The y and z axes of the LHCb coordinate system are also shown. This system is a right-handed Cartesian coordinate system, with the x axis pointing away from the observer.

these subdetectors serve as the particle-identification (PID) system. The subdetectors are described in Sections 4.2.1 and 4.2.2. The LHCb detector is the only detector at the LHC capable of track reconstruction and PID across its entire acceptance in η . It is thus a general-purpose detector in the forward η region [7].

During the majority of data recording, crossings of nonempty bunches occurred at 30 MHz [100]. Storing all of these events, accumulating over several years of operation, on the available computing resources is impossible. It would also be inefficient with regard to beauty and charm hadrons since these are produced at a much lower rate. Hence, the LHCb experiment is equipped with a trigger system [101], which considerably reduces the number of events retained. The trigger comprises a hardware stage and a two-level software stage. The hardware stage can employ information from the calorimeter system and the muon stations at a rate of 1 MHz to decide if an event is accepted or not. Muons with high p_T , for example, hint at decays of heavy hadrons. At the software stage, the accepted events are first partially and then fully reconstructed. Multiple sets of selection requirements exist

at the software stage to only accept events containing decay topologies of interest. The events retained by the trigger are persisted for data analysis.

Besides its instrumentation in the forward η region, another unique feature of the LHCb experiment among those at the LHC is its capability to be operated as fixed-target experiment [102]. Gas can be injected at low pressure into the VELO vacuum vessel. The beams then collide with the gas, which spreads along the LHCb section of the LHC beam pipe; at the experiments, the two beams travel in a joint beam pipe. Only noble gases are allowed to prevent damage to the VELO. The collisions with the gas targets result in nucleon–nucleon centre-of-mass energies of up to 110 GeV, which have not been reached before in a fixed-target experiment. Collisions of proton beams with helium, neon and argon were recorded as well as collisions of lead-ion beams with the latter two gases. The gas-injection system was originally designed to determine the integrated luminosity in colliding-beam mode via beam–gas imaging.

The LHCb detector in the setup described above and in Sections 4.2.1 and 4.2.2 recorded data from late 2009 to early 2013, with the majority collected in 2011 and 2012, and from mid-2015 to late 2018. Afterwards, the detector received a comprehensive upgrade to increase the instantaneous luminosity [103]. It was commissioned from 2022 onwards. The main changes are summarised in this paragraph. The hardware stage of the trigger system was removed. Consequently, the LHCb experiment is nowadays equipped with a software-only trigger, which has significantly higher efficiency than before, especially for fully hadronic and electronic final states at low p_T [104]. The new trigger strategy requires a readout of all subdetectors at 40 MHz. Moreover, they need to deal with higher occupancy, *i.e.* a larger number of hits from particles. The subdetectors of the track-reconstruction system were completely replaced. Various parts of the RICH detectors, the two calorimeters and the muon stations were upgraded as well. In addition, the fixed-target operation was improved in several ways by a new gas-injection system and a dedicated storage cell [105]. The cell possesses a length of 20 cm and was installed upstream of the VELO sensors in the VELO vacuum vessel. Gas is directly injected into the cell. Thus, the gas density is higher and more precisely known, and the interaction region is much better defined. The upgraded system is more flexible because nonnoble gases can also be injected, such as hydrogen, deuterium and possibly oxygen. Furthermore, the integrated luminosity can be determined with higher accuracy.

4.2.1 Track reconstruction

The VELO [106] is a silicon-microstrip detector consisting of 42 semicircular sensors with a diameter of 8.4 cm and a semicircular gap in the centre. The sensors are arranged along the beam axis in two retractable detector halves, located left and right of the beams inside a vacuum vessel. During data recording, the sensors

are positioned at a distance of only 7 mm from the beams. Before stable beam conditions in the LHC are reached, the detector halves are retracted from the beam axis to avoid damage. The storage cell for gases is therefore also composed of two cylindrical halves. There is no beam pipe inside the VELO; the sensors in each detector half are only shielded from the beams by a thin aluminium-alloy foil. When traversing the sensors, charged particles leave hits via the creation of electron-hole pairs, separated by the electric field applied and evoking pulses. The reconstruction software combines the hits into track segments, which are straight lines since the residual field from the LHCb magnet in the VELO is negligible. Searching for groups of track segments pointing to the same origin permits the primary vertex to be reconstructed as well as decay vertices, in particular of beauty and charm hadrons.

The TT [107] is a silicon-microstrip detector comprising four rectangular layers. Each of them covers an area of 2 m^2 . The microstrips in the first and last layers are assembled vertically, whereas those in the two middle layers are tilted by $\pm 5^\circ$ to constrain hit positions both horizontally and vertically.

The magnet is a warm dipole magnet with an integrated magnetic-flux density of approximately 4 T m. The magnetic field forces charged particles leaving the VELO to follow curved trajectories.

In the three tracking stations downstream of the magnet, two different detection techniques are used. The parts close to the beam pipe—three cross-shaped stations of four-layer silicon-microstrip detectors—form the Inner Tracker (IT) [107], while the remainder—three stations of four layers of straw drift tubes—constitutes the Outer Tracker (OT) [108]. The IT is constructed from silicon because of the high particle density in that region, requiring high resolution and large radiation hardness. The orientation of the microstrips and drift tubes in the four IT-OT layers per station is equal to that in the TT. Each of the IT-OT stations covers an area of 30 m^2 . Charged particles passing through the drift tubes induce hits via ionisation of the gas inside; the charge carriers created drift in the electric field maintained between the wire in the middle of the tube and its wall, which leads to pulses.

The momentum of a charged particle is determined from the measured radius of curvature, and the particle charge can be deduced from the bending direction. If its momentum is less than approximately $2\text{ GeV}/c$, a particle is deflected too strongly to reach the IT-OT stations. This is one reason for the existence of various track types [7]. Tracks traversing the entire magnetic field have the most precise momentum estimate and are called *long*. To form long tracks, the reconstruction software searches for hits in the IT-OT stations that can be combined with VELO track segments. If hits are found in the TT that are compatible with an assumed trajectory, they are also included in the subsequent track fit to improve the momentum resolution. The relative uncertainty on the long-track momentum ranges from 0.5% below $20\text{ GeV}/c$ to 1.0% at $200\text{ GeV}/c$. It is not guaranteed that the hits assigned to a track were caused by only one real particle passing through the detector. In the case

of high occupancy, the probability is increased of incorrectly combining VELO and IT–OT hits. This results in fake tracks. The presence of TT hits helps to discriminate real from fake tracks. The fraction of fake tracks relative to all tracks can be reduced with a neural-network-based variable [109] computed by the reconstruction software for each track. The neural-network input comprises the occupancy of each track-reconstruction subdetector, the track kinematics, the numbers of hits and the quality of the track fit.

4.2.2 Particle identification

The two RICH detectors [110] measure the Cherenkov photons induced by charged particles travelling faster than the speed of light in the detector media. The media used in the first and second RICH detectors are the C_4F_{10} and CF_4 gases, respectively, enabling the detectors to cover low and high momenta via the different refractive indices. The photons, emitted on cones, are reflected onto detection planes outside the acceptance, which creates ring images. By employing information from the track-reconstruction system on the trajectories in the RICH detectors, photon rings can be linked to tracks. This allows the emission angle of the Cherenkov radiation to be determined for each track. If a mass hypothesis is assigned to a track, the dependence of the Cherenkov angle on the track momentum is known. Consequently, the momentum information is needed to distinguish charged-particle species. The reconstruction software first assumes the charged-pion-mass hypothesis for all tracks and computes a likelihood from the photon-hit and track distributions. The charged-kaon-, (anti)proton-, electron- and muon-mass hypotheses are then inserted for each track, and the differences of the logarithmic likelihoods relative to that for the charged-pion-mass hypothesis are determined. These differences constitute PID variables. They are indispensable for the identification of charged hadrons. The nature of the Cherenkov radiation implies that there is a momentum threshold for each charged-particle species below which no photon emission occurs. In the first RICH detector, this threshold amounts to $9.3 \text{ GeV}/c$ for charged kaons. At lower momentum, both charged kaons and (anti)protons can be recognised by the absence of photon rings, but separating them is impossible.

The calorimeter system [111] measures the energy and position of particles. It also provides hardware-trigger and PID information. The detection proceeds via electromagnetic and hadronic particle showers developing in dense materials, which are similar to air showers. The scintillating-pad and preshower detectors are separated by a thin lead absorber and consist of square scintillator cells. A hit in the scintillating-pad detector indicates the passage of a charged particle, whereas an energy deposit in the preshower detector hints at the start of an electromagnetic shower. The electromagnetic and hadronic calorimeters are sampling calorimeters composed of square cells of alternating layers of an absorber material and a

scintillator. The absorber materials used are lead and iron in the electromagnetic and hadronic calorimeters, respectively, which have depths of in total 25 radiation lengths and 5.6 nuclear-interaction lengths. The depth of the hadronic calorimeter is limited owing to space constraints in the cavern. The scintillation light deposited by the shower particles is read out and converted into an energy estimate. The calorimeter system is the only subdetector capable of detecting neutral particles, which are primarily photons.

Four of the five muon stations [112, 113] are located downstream of the calorimeter system. They contain numerous multiwire proportional chambers to measure hits from charged particles not absorbed by the calorimeters. The majority of particles reaching the chambers is muons as they solely interact as minimum-ionising particles. However, the limited depth of the hadronic calorimeter can also cause hadrons to enter the stations. To mitigate the risk of misidentification and to isolate muons with high momentum, the four stations are interleaved with iron absorber blocks, each possessing a depth of 80 cm. The first muon station is placed upstream of the scintillating-pad detector to improve the p_T estimate in the hardware stage of the trigger. Its outer part is also covered by multiwire proportional chambers, but its inner part is equipped with gas electron multipliers due to the high-radiation environment. The only particle species that can even escape the muon stations is neutrinos, which cannot be measured directly with the LHCb detector.

Part II

Measurements

5 Prompt charged-particle production

The first measurement performed within this thesis is presented in this chapter. It has been published in Reference [8]. The differential cross-section of prompt production of long-lived charged particles in p - p collisions at $\sqrt{s} = 13$ TeV has been measured with the LHCb experiment. To date, this is the highest energy at which charged-particle production in p - p collisions has been measured in the forward η region. The analysis is conducted in intervals of η and p_T . Such a double-differential measurement had not yet been done. Moreover, positively and negatively charged particles are studied separately to reveal potential charge asymmetries. Previously, charged-particle production in p - p collisions at $\sqrt{s} = 7$ TeV was measured by the LHCb collaboration in intervals of either η or p_T [31]. The CMS and TOTEM [114] collaborations analysed the η dependence of charged-particle production in the forward region of p - p collisions at $\sqrt{s} = 8$ TeV [115]. This was also studied by the ALICE collaboration, who additionally investigated p - p collisions at $\sqrt{s} = 0.9$ TeV and $\sqrt{s} = 7$ TeV [116]. A double-differential measurement in η and p_T of charged-particle production in p - p collisions at $\sqrt{s} = 5$ TeV [117] was performed in parallel to that presented in this thesis by an independent group of LHCb collaborators. It was published slightly later.

5.1 Data and simulation samples

Two data samples are used in the analysis. They were recorded in July and August 2015 in p - p collisions at $\sqrt{s} = 13$ TeV, where both beam energies amounted to 6.5 TeV. Each data sample was collected during a single period of stable beam conditions in the LHC. At the hardware stage of the trigger system, events were accepted at a fixed rate and without any requirement on activity in the detector. The software stage only retained events from crossings of leading bunches of the LHC bunch trains, which avoids background in the detector from preceding events. Thanks to this unbiased trigger selection, the data samples are perfectly suited for absolute measurements of production cross-sections. The direction of the field of the LHCb magnet is regularly reversed throughout data recording to minimise the impact of detection asymmetries. Each of the two data samples was recorded in one of these two magnetic-field configurations. Analysing data from both configurations enables a cross-check between results independently obtained from the two samples. The

magnet-up and magnet-down data samples approximately comprise 226×10^6 and 134×10^6 events, respectively, and correspond to integrated luminosities of 3.0 nb^{-1} and 2.4 nb^{-1} .

Simulation samples are needed as well. In contrast to data, the species and origin of each track are known in simulation. Therefore, the signal efficiency of the selection requirements imposed can be determined, and the amount of residual background can be estimated. The p - p collisions are generated with PYTHIA 8.186 [118, 119]. The interactions of the generated particles with the detector and its response are handled by the GEANT4 toolkit [120, 121]. The same software as on data is then run to reconstruct tracks from the hits in simulation. Furthermore, the tracks at reconstruction level are associated with particles at generator level if possible or otherwise classified as fake. There are several sources of imperfection in simulation; the kinematics of the generated particles can be incorrectly modelled, the description of the detector material is partially simplified, and not all variations of experimental conditions are taken into account. Two simulation samples are used in the analysis, one for each magnetic-field configuration and each approximately containing 10^7 events. The events were generated without requiring the presence of any specific particle or decay.

The two statistically independent pairs of one data sample and one simulation sample are analysed separately throughout.

5.2 Strategy

The concept of prompt production of long-lived charged particles has already been mentioned several times. A measurement is ideally based on a definition of this concept that is independent of the detector and can easily be applied to the output of event generators. The widely used definition [122], proposed by the ALICE collaboration, is employed; a particle is classified as long-lived if its mean lifetime exceeds 30 ps, and it is prompt if it is either directly produced in the first interaction or does not have any long-lived ancestor particle. According to this definition, long-lived charged particles are charged pions, charged kaons, (anti)protons, electrons and muons as well as Σ^+ , $\bar{\Sigma}^-$, Σ^- , $\bar{\Sigma}^+$, Ξ^- , $\bar{\Xi}^+$, Ω^- and $\bar{\Omega}^+$ baryons. If produced promptly, they are part of the signal in this analysis.

The number, n , of signal particles enters the differential cross-section, which is determined in each interval of η and p_T as

$$\frac{d^2\sigma}{d\eta dp_T} = \frac{n}{\mathcal{L} \Delta\eta \Delta p_T}, \quad (5.1)$$

where $\Delta\eta$ and Δp_T denote the kinematic-interval widths. The integrated luminosity of the data sample under study is denoted by \mathcal{L} . The number of signal particles is

not readily accessible in data. Long tracks as introduced in Section 4.2.1 are hence chosen as candidates. Their selection is described in Section 5.3. The number, n_{cand} , of candidate tracks is related to n according to

$$n_{\text{cand}} = \varepsilon n + \sum_i n_i. \quad (5.2)$$

The efficiency of the track reconstruction and selection in the kinematic interval considered is denoted by ε . The sum includes the number, n_i , of background tracks from each source, i , that contributes. One source of background is fake tracks. Another source is nonprompt tracks from interactions of particles with the detector material or from decays of long-lived particles. The values of ε and n_i are taken from simulation, but control studies are performed to correct for imperfect modelling in simulation. For each source of background, a proxy observable, \mathcal{D}_i , is defined that is proportional to n_i and determinable both in data and simulation. With the assumption that the ratio of the numbers of background tracks in data and simulation is equal to the ratio of the proxies in data and simulation,

$$\frac{n_i}{n_{i,\text{sim}}} = \frac{\mathcal{D}_i}{\mathcal{D}_{i,\text{sim}}} =: R_i, \quad (5.3)$$

the unknown number of background tracks in data can be estimated as $n_i = R_i n_{i,\text{sim}}$. The proxies constructed are presented in Section 5.4. The number, n_{BG} , of background tracks from beam–gas interactions can directly be obtained from data. Thus, the number of signal particles in Equation (5.1) can be expressed as

$$n = \frac{1}{R_\varepsilon \varepsilon_{\text{sim}}} \left(n_{\text{cand}} - \sum_i R_i n_{i,\text{sim}} - n_{\text{BG}} \right), \quad (5.4)$$

where R_ε is a correction to the efficiency in simulation, explained in Section 5.5. The differential cross-section is measured separately for positively and negatively charged particles. Therefore, the charge ratio,

$$\frac{d^2\sigma_+}{d\eta dp_T} \bigg/ \frac{d^2\sigma_-}{d\eta dp_T} = \frac{n_+}{n_-}, \quad (5.5)$$

can be computed, which offers the advantage that the uncertainty on the integrated luminosity and further uncertainties cancel. The charge is given in units of the elementary charge throughout.

The analysis code is entirely written in the PYTHON programming language and uses standard scientific packages as well as packages designed for high-energy particle physics, which are listed in the published paper of the analysis [8].

5.3 Selection

The number of selection requirements is kept as small as possible, which avoids the need to estimate additional efficiencies and corresponding systematic uncertainties. Only two requirements are applied to the long tracks in the events accepted by the trigger system. The fit of each track is required to be of good quality. In addition, a requirement is imposed on the maximum value of the fake-track probability, P_{fake} , which is the neural-network-based variable mentioned in Section 4.2.1. The requirement rejects between 40 % and 80 % of the fake tracks, depending on the kinematic interval, and retains 99 % of the real tracks. The selection is applied to the data and simulation samples. The long tracks passing the selection are referred to as candidate tracks.

The differential cross-section slowly varies in η . Hence, a coarse η grid with six intervals in the $\eta \in [2.0, 4.8)$ range is chosen. The widths of the first five intervals amount to $\Delta\eta = 0.5$, whereas the last interval is reduced in size. This decision was made to reflect that the track-reconstruction efficiency decreases in the $\eta \in [4.8, 5.0)$ range. In p_T , the differential cross-section has a power-law shape. A logarithmic p_T grid is therefore used with 50 intervals in the $p_T \in [10, 10\,000)$ MeV/ c range and an interval width of $\Delta\log_{10}[p_T/(\text{MeV}/c)] = 0.06$. There are two momentum thresholds that restrict the lower end of this grid and cause the minimum accessible p_T value to depend on η . First, a particle needs a minimum momentum of approximately 2 GeV/ c to reach the IT–OT stations. Second, the correction to the efficiency in simulation introduces a requirement on the minimum track momentum of 5 GeV/ c .

The detector occupancy observed in data is not well reproduced in simulation. Since the analysis is not performed in intervals of an occupancy proxy, the simulation samples have to be adjusted, which is achieved by weights. To compute the weight for each event in simulation, the ratio of the distributions of the numbers of long tracks per event in data and simulation is formed. The value of this ratio at the occupancy of the event under consideration then represents the weight. Each of the two simulation samples is weighted to reproduce the occupancy of the data sample with the same magnetic-field configuration. The simulation weighting also normalises the simulation samples to the data samples, which is required due to the different numbers of events.

5.4 Background contributions

The composition of the background is shown in Figure 5.1 besides the signal, *i.e.* prompt long-lived charged particles. It is determined on simulation by enquiring the ancestor particle of each candidate track associated with a long-lived charged particle. The signal prevails; around $p_T \approx 2000$ MeV/ c , its fraction is greater




Figure 5.1: Origins of the candidate tracks associated with long-lived charged particles in simulation in intervals of η and p_T . The white areas correspond to the signal. Background contributions with a fraction of less than 0.5% across all intervals are included in the *other* category. The right ends of the light-grey rectangles at low p_T correspond to a momentum of 2 GeV/c, the approximate minimum momentum needed for the long-track reconstruction. In the following, these limits of the acceptance are shown in all plots displaying η and p_T dependence.

than 90%. Towards lower and higher p_T , the background increases. The largest background contribution is fake tracks, which occur across the entire p_T ranges and become more relevant with growing η . Tracks originating from photon conversions in the detector material are nonnegligible at low p_T and especially important at high η . Another type of background from material interactions is tracks stemming from interactions of charged pions with the detector, mostly present at low p_T . In the same kinematic region, tracks produced in decays of K_S^0 mesons occur. Three proxies are constructed to adjust the sizes of these background contributions in simulation to those in data. Because of the simulation weighting, the value of a proxy ratio in data and simulation can be interpreted; if the proxy ratio is unity, the background in data is appropriately described in simulation. The other sources of background are small and not adjusted. Instead, their contributions in simulation are combined, and the proxy ratio used for them is $R_{\text{other}} = 1.0 \pm 0.5$, where the assumed uncertainty is systematic and fully correlated between the two charges.

5.4.1 Fake tracks

The P_{fake} variable [109] is well suited for the study of fake tracks, which are the largest background contribution. The distribution spans the $P_{\text{fake}} \in [0, 1]$ range by construction. Tracks from real particles accumulate at small values, while fake tracks are approximately uniformly distributed across the entire range. The requirement from the selection discards tracks above $P_{\text{fake}} = 0.3$. To estimate the residual background from fake tracks in data, a control sample dominated by fake tracks has to be identified. Hence, the distribution is considered up to $P_{\text{fake}} = 1$. It is divided into ten intervals with a width of 0.1. In each of them, the purity of fake tracks is computed in simulation. The lowest interval above $P_{\text{fake}} = 0.3$ that has a purity of at least 80 % is chosen. The number of tracks in this interval constitutes the proxy, which is thus computable both in data and simulation. By choosing an interval close to the range from the selection, potential differences in the shapes of the P_{fake} distributions between data and simulation are small.

This procedure is applied to the P_{fake} distribution in each kinematic interval. The resulting proxy ratio is shown in Figure 5.2. It exceeds unity in almost all kinematic intervals, meaning that the number of fake tracks is underestimated in simulation. Values up to 2.8 are found, but the impact of these larger adjustments is still small according to Equation (5.4) as the number of fake tracks in simulation is small in those kinematic intervals. Discontinuities in the p_T dependence are caused by alterations in the P_{fake} interval chosen. The statistical uncertainty on the proxy ratio is obtained by propagating the uncertainties on the proxy counts in data and simulation. A systematic uncertainty is determined by computing the proxy in an alternative way. This alternative consists in counting the tracks in the P_{fake} interval chosen and all intervals above. Consequently, it is more affected by potential shape differences between data and simulation and can be regarded as an extreme variation. The uncertainty stemming from such a variation is reflected by a uniform distribution. The endpoint of this distribution is given by the alternative proxy ratio, whereas its mean value is set to the default proxy ratio. Both charges are combined in this study. Thus, the statistical and systematic uncertainties are fully correlated between the two charges.

Clone tracks are another reconstruction artefact. They can occur in close vicinity to the trajectory of only one real particle if two or more tracks are inferred from the hits created by this particle. Dedicated algorithms reduce the fraction of clone tracks to the level of 0.1 % [123, 124]. Nevertheless, their residual fraction is investigated separately in each kinematic interval. It is determined from two counts. First, all candidate tracks in simulation are counted. Second, a weighted count is initialised. For each candidate track in simulation, the associated particle is evaluated. If a particle was reconstructed as only one track, the weight for this track is set to zero since no clone track is present. If two tracks are associated with the same




Figure 5.2: Proxy ratio for fake tracks in intervals of η and p_T . The bars and boxes indicate the statistical and systematic uncertainties, respectively.

particle, each track is assigned a weight of $1/2$. Three-track associations and so forth are treated analogously, but become increasingly unlikely. In general, the weight for a track is $(j - 1)/j$, with j denoting the number of tracks associated with the corresponding particle. Dividing the second count by the first one gives the residual fraction of clone tracks. This fraction is found to be less than 0.06 % in each kinematic interval, rendering clone tracks negligible.

5.4.2 Material interactions

The second-largest source of background is tracks produced in material interactions. One type of this background is tracks originating from material interactions of charged pions. These interactions occur in the VELO because at least some of the interaction products are reconstructed as long tracks. The method developed in this background study was inspired by Reference [125], in which a high-precision material map of the VELO was created. The proxy for material interactions of charged pions is defined as the number of tracks belonging to candidate vertices that are formed by three tracks and located in the VELO material. The method is detailed below. Vertices consisting of three tracks are chosen since hadronic

interactions frequently produce three or more charged particles, while decays into three or more charged particles are rare in the data samples analysed.

In each event, all unique combinations of three candidate tracks are formed. Each point of closest approach is computed by a fit that minimises the sum of the distances of the tracks. This point constitutes the candidate vertex of a material interaction. An initial reduction of random combinations of three unrelated tracks is achieved by requiring the radial coordinate of the vertex in the x - y plane to be greater than 3 mm. The requirement does not reject any real material interaction as no material is present in that region.

The purity of the obtained set of candidate vertices is further increased by applying requirements optimised on simulation. Firstly, a simultaneous optimisation is performed of requirements on the minimum z coordinate of the vertex and on the maximum sum of the squared track distances to their vertex. A minimum- z requirement further discards regions near the interaction point, which are dominated by random track combinations. The simultaneous optimisation is done without dividing the simulation samples into kinematic intervals. The optimal requirements are determined in a grid search by maximising the significance of real interactions via a figure of merit. The two-dimensional distributions of the z and radial coordinates in data and simulation as well as the optimised minimum- z requirement are shown in Figure 5.3. The method allows the material of the VELO sensors and the aluminium-alloy foil shielding them to be reliably traced. Some simplifications of the material description in simulation are discernible; *e.g.*, the modelled shape of the foil is less curved. Moreover, the increased track and vertex densities near the interaction point around $z \approx 0$ mm are evident, like the requirement on the minimum radial coordinate. This requirement is not tightened by an optimisation since the VELO positions in data and simulation slightly differ for negative x values. Applying a requirement optimised on simulation would reject relevant regions in data and thus bias the proxy ratio.

Afterwards, an optimisation of a requirement on the minimum invariant mass of the track triplet of each candidate vertex is performed. As the species of a single track is not identified, a zero-mass hypothesis is assigned to each track. The optimisation is executed separately in each kinematic interval. To reduce fluctuations of the proxy ratio to be determined, the p_T grid is made coarser by subsequently merging three adjacent intervals starting from the lower end. Example mass distributions with the corresponding optimised minimum-mass requirements are shown in Figure 5.4. The shapes of the distributions in data and simulation are found to be in reasonable agreement. The requirements mainly suppress random track combinations. The distributions indicate that the contribution of material interactions of charged pions to the set of candidate vertices is approximately equal in size to that of material interactions of other particles, for example kaons, (anti)protons and (anti)neutrons.

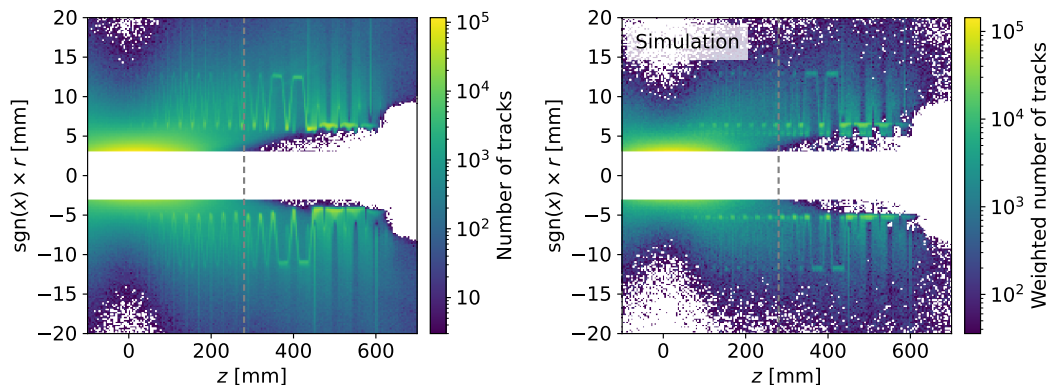


Figure 5.3: Two-dimensional distributions of the z and radial coordinates of track triplets from potential material interactions in data (left) and simulation (right). The radial coordinate is multiplied by the sign of the x coordinate to enable a separate depiction of the two VELO halves. The grey dashed lines indicate the optimised minimum- z requirement.

Since material interactions of charged pions cannot be selected without other relevant contributions, the purity of the set of candidate vertices needs to be studied in simulation. The efficiency of the candidate-vertex selection is computed as well to ensure that it is not too low, in which case the proxy ratio might be biased. The efficiency and the purity are shown in Figure 5.5. The efficiency is around 10%. The purity mostly ranges from 20% to 45% and decreases towards high η . Low purity can render the proxy ratio unreliable due to the other contributions selected, which are not proportional to the background to be adjusted. It is found that the proxy ratio is independent of the purity above a purity of 30%. Consequently, the impact of the nonproportional contributions is negligible if the purity exceeds this threshold. The value of the proxy ratio, determined separately for each of the two charges, is therefore only used in the kinematic intervals with a purity above the threshold, which are termed accepted intervals. The statistical uncertainty on the proxy ratio in the accepted intervals is obtained by propagating the uncertainties on the numbers of selected tracks in data and simulation. To estimate the proxy ratio in the nonaccepted intervals, the median value of all values in the accepted intervals is computed, and the uncertainty given by a uniform distribution covering these values is assigned to it.

A systematic uncertainty on the proxy ratio is determined by once loosening and tightening the default requirements on the z coordinate and on the sum of the squared track distances. Because of the structure of the VELO material, the z distribution is not smooth. Hence, the minimum- z requirement is only moderately varied. The requirements on the invariant mass are not altered as the shapes of




Figure 5.4: Invariant-mass distributions of track triplets in two kinematic intervals with different levels of background from material interactions of charged pions: large background (left) and small background (right). The species of the tracks are not identified, but most of them are charged hadrons; the invariant mass is computed with zero-mass hypotheses. The three stacked distributions are from simulation; besides the tracks originating from material interactions of charged pions and those from random track combinations, the tracks stemming from material interactions of other particles are shown, for which the direct ancestor particles are listed. The grey dashed lines indicate the optimised minimum-mass requirements.

the distributions in data and simulation are similar. For each of the four resulting variations, an alternative proxy ratio, R'_{mat} , is computed analogously to the default proxy ratio, R_{mat} ; the same purity threshold is used, and the median value is determined in the nonaccepted intervals. Some variations that cannot be explained by fluctuations are found, which implies that the descriptions of the corresponding efficiencies in data and simulation are not in perfect agreement. The default and alternative ratios are assumed to be equally plausible. Thus, the mean and variance of the default ratio and those alternative ratios that differ significantly from the default one are computed. The deviation of an alternative ratio is considered to be significant if the condition

$$(R'_{\text{mat}} - R_{\text{mat}})^2 > |u^2(R'_{\text{mat}}) - u^2(R_{\text{mat}})| \quad (5.6)$$

is fulfilled [126]. The difference of the variances, u^2 , appears since the same data set is employed for the determination of the default and alternative ratios. The mean of the proxy ratios then replaces the default proxy ratio, and their variance estimates the uncertainty on the mean. This procedure is applied both in the accepted and nonaccepted intervals of the default proxy ratio. In the accepted intervals, the uncertainty on the mean represents the systematic uncertainty on the proxy ratio, while its statistical uncertainty arises from the numbers of selected tracks in the




Figure 5.5: Efficiency and purity of the material-interaction selection in simulation in intervals of η and p_T . The bars indicate the statistical uncertainties.

default case. In the nonaccepted intervals, the uncertainty on the proxy ratio is purely systematic and given by the sum in quadrature of the uncertainty on the mean and the uniform uncertainty in the default case mentioned above.

The proxy ratio obtained is shown in Figure 5.6. In almost all kinematic intervals, the background is overestimated in simulation. There are some charge asymmetries in the proxy ratio at low η . The relative uncertainty on the median value amounts to 36%. The statistical and systematic uncertainties in the accepted intervals are uncorrelated between the two charges. The systematic uncertainty in the nonaccepted intervals is fully correlated between the two charges by construction. Nevertheless, it is assumed to be uncorrelated to avoid cancellation of the uncertainty in the charge ratio of the differential cross-section, which is explained in Section 5.7.2.

The other type of background from material interactions is conversions of photons into e^+e^- pairs, which substantially contribute at high η and low p_T . The construction of a separate proxy for this background was studied, but found to be unreliable without the inclusion of PID information, which is described in the following. Pairs of oppositely charged candidate tracks were formed based on requirements on the position of the candidate conversion vertex in the VELO and on the invariant mass. Since the photon is massless, the invariant mass of the track pairs is expected to peak close to zero. Such a peak was found in the mass distributions in data

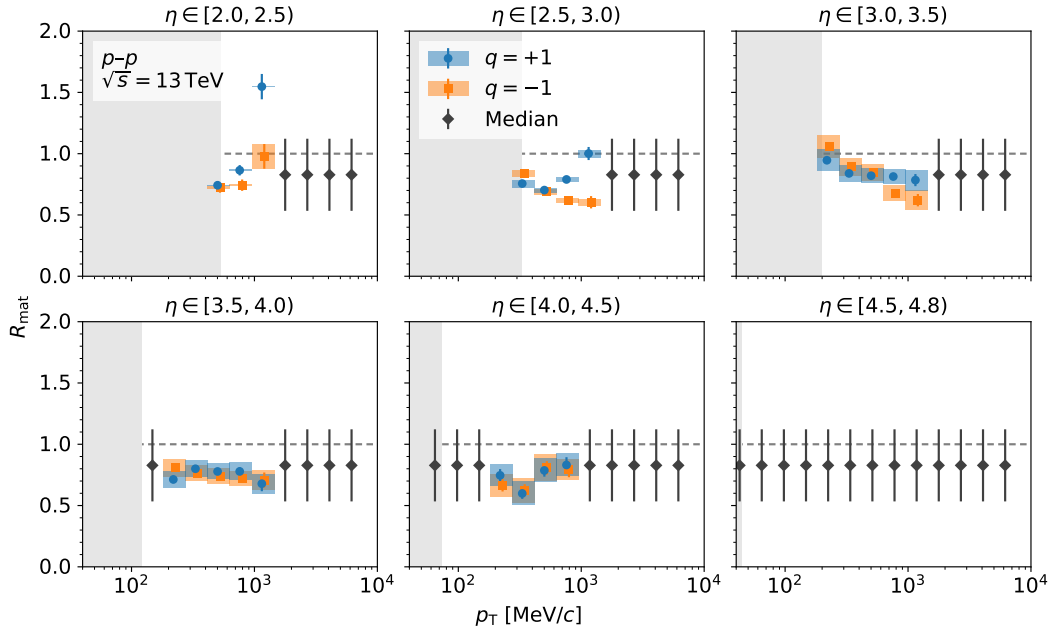


Figure 5.6: Proxy ratio for material interactions in intervals of η and p_T . The bars and boxes on the markers for the two charges indicate the statistical and systematic uncertainties, respectively. The bars on the markers for the median value represent its systematic uncertainty.

and simulation. However, when the track pairs associated with conversions were isolated in simulation, it was observed that the peak cannot be solely attributed to conversions. Due to this ambiguity, the integral of the peak was not used to define a proxy for photon conversions. Instead, the proxy for material interactions of charged pions is employed to adjust the background from photon conversions in simulation as well. This is justified as most of the photons originate from decays of neutral pions, of which the production is coupled to that of charged pions because of isospin symmetry. Furthermore, the conversion probability of photons and the interaction probability of charged pions are both proportional to the amount of material traversed.

5.4.3 Strange-hadron decays

Another source of nonprompt tracks, besides material interactions, is decays of long-lived particles. The main contribution is $K_S^0 \rightarrow \pi^+ \pi^-$ decays happening in the VELO. Candidate decays are obtained by selecting in each event all unique pairs of oppositely charged candidate tracks that possess a distance of closest approach of less than 1 mm. In addition, the point of closest approach must be downstream of the

mean primary-vertex position by at least 150 mm along the z axis. This requirement takes into account the long lifetime of the K_S^0 meson and further discards random combinations of two unrelated tracks. The candidate-decay selection is deliberately loose to keep its efficiency high. The topologies of the $\Lambda \rightarrow p\pi^-$ and $\bar{\Lambda} \rightarrow \bar{p}\pi^+$ decays closely resemble that of the $K_S^0 \rightarrow \pi^+\pi^-$ decay. Consequently, those decays can be selected in the same way and are therefore included as well although their contribution to the background is smaller. The proxy ratio is deduced from the yields of these three decays in data and simulation.

The invariant masses of the track pairs are computed with charged-pion-mass hypotheses assigned to both tracks for K_S^0 candidates and with (anti)proton-mass and charged-pion-mass hypotheses assigned to the tracks for Λ and $\bar{\Lambda}$ candidates. To be retained as a candidate, the invariant mass is required to lie in a window around the known mass [17] of the corresponding decaying particle. The candidates are then sorted into kinematic intervals according to the η and p_T values of the decaying particle. The p_T grid is made coarser like in Section 5.4.2 to reduce the statistical uncertainty on the number of candidates per kinematic interval.

In each kinematic interval with sufficiently many candidates, extended-maximum-likelihood fits are performed to the invariant-mass distributions in data and simulation. The K_S^0 , Λ or $\bar{\Lambda}$ component is modelled with a nonstandardised Student's t distribution. This distribution, which interpolates between a nonrelativistic Breit-Wigner and a Gaussian distribution, has three parameters: a location shift, a width modification and the number of degrees of freedom. In the fits, the distribution is multiplied by a parameter reflecting the yield of the K_S^0 , Λ or $\bar{\Lambda}$ component. No other peaking structure is present in the invariant-mass distributions. The random track combinations form a flat component, which is modelled with a second-degree Bernstein polynomial. Bernstein polynomials possess the excellent numerical properties of Chebyshev polynomials, but guarantee that the density is nonnegative if the coefficients are nonnegative. Thus, issues during numerical minimisation of the cost function are avoided. Besides limits that are placed on the parameters, weak Gaussian constraints are employed in the cost function to permit the fit to converge even if the peak of the K_S^0 , Λ or $\bar{\Lambda}$ component is small. A kinematic interval is only further considered if that peak has a significance of at least three standard deviations. Example fits are shown in Figure 5.7. Each fit returns the K_S^0 , Λ or $\bar{\Lambda}$ yield and its statistical uncertainty. The quality of a fit is reflected by its χ^2/n_{DOF} value, where χ^2 denotes the sum of the squared normalised residuals and measures the deviation of the fit from the distribution fitted. The number of degrees of freedom in the fit is denoted by n_{DOF} . A χ^2/n_{DOF} value around unity generally indicates a good fit, which is found to be the case.

The yield ratios are shown in Figure 5.8. A pattern is observed in the p_T dependence across all η intervals, which is similar for the ratios of the K_S^0 , Λ and $\bar{\Lambda}$ yields.




Figure 5.7: Fits to the invariant-mass distributions of $K_S^0 \rightarrow \pi^+ \pi^-$ candidates (top) and $\Lambda \rightarrow p \pi^-$ candidates (bottom) in data (left) and simulation (right) in one kinematic interval.

This suggests a general discrepancy between data and simulation in the number of strange quark–antiquark pairs produced. The yield ratios do not cover all kinematic intervals. Since they do not change strongly with p_T , their dependence can be modelled. The values of the ratios of the K_S^0 , Λ and $\bar{\Lambda}$ yields from all η intervals are combined and fitted with a monotone cubic Hermite spline [127]. This type of cubic spline has the advantage that it cannot overshoot the data points. The spline is based on six support points, which are manually placed. The values of the spline at the support points are fitted with a cost function of least squares. A weak Gaussian constraint is added that guides the spline to unity at the upper end of the p_T grid, where no yield ratio is present. The fitted model describes by which amount the background from strange-hadron decays needs to be adjusted in simulation depending on the strange-hadron kinematics. The statistical uncertainty on the fitted model is neglected as it is much smaller than the deviations of the yield ratios from the model. Instead, a systematic uncertainty of 0.15 is estimated to cover the

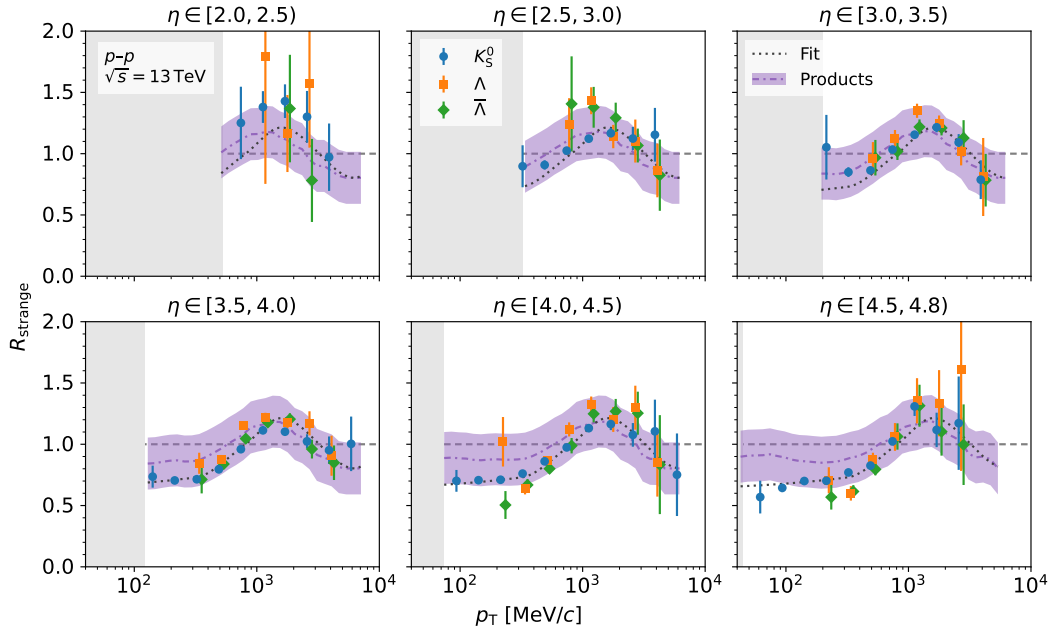


Figure 5.8: Proxy ratio for strange-hadron decays in intervals of η and p_T . The purple dash-dotted lines and the bands surrounding them indicate the proxy ratio and its systematic uncertainty, respectively, which have been mapped to the kinematics of the products of the K_S^0 , Λ and $\bar{\Lambda}$ decays. The dark-grey dotted lines represent the spline fit to the ratios of the K_S^0 , Λ and $\bar{\Lambda}$ yields in data and simulation. These ratios are indicated by the markers, with the bars representing the propagated statistical uncertainties from the invariant-mass fits.

deviations appropriately. Lastly, the fitted model has to be mapped to the kinematics of the products of the strange-hadron decays, which constitute the background to the candidate tracks. This is achieved based on simulation, where the connections are known between the decaying particles and their real decay products. The fitted and mapped models are also shown in the figure. The mapped model represents the proxy ratio for strange-hadron decays. It is closer to unity than the fitted model because the kinematic distributions of the decay products dilute the discrepancies. The uncertainty on the proxy ratio, which is purely systematic, is fully correlated between the two charges.

5.4.4 Beam–gas interactions

The beam pipe and the VELO vessel are evacuated. Nevertheless, residual gas can be present. If the beams interact with it, additional particles are produced. In the case of reconstruction as long tracks, such particles can pass the selection in the

analysis. However, they must not be included since they do not stem from a p - p collision at $\sqrt{s} = 13$ TeV. The presence of such tracks is studied on two control data samples comprising beam-gas events, in which only bunches from one of the two beams traverse the interaction point. Both possible configurations are included: that with the bunches travelling from the VELO towards the muon stations and the opposite one. These samples were recorded concurrently with the two default data samples.

The number of tracks from beam-gas events passing the selection is normalised to the number of default events. The fraction relative to the number of candidate tracks is then computed. For the magnet-down data sample, the fraction of tracks amounts to 1 % in all kinematic intervals. As this cannot be neglected, the background from beam-gas interactions is subtracted as prescribed by Equation (5.4). For the magnet-up data sample, the fraction is below 0.1 % and hence negligible.

5.5 Efficiency

The efficiency of reconstructing and selecting signal particles in a given kinematic interval depends on track and particle loss. Track loss can primarily be attributed to the detector occupancy. Particle loss is determined by the particle composition and the amount of material traversed. The particle composition is relevant because different particle species possess different lifetimes and different interaction lengths. To investigate the efficiency, the particle composition is represented by four classes: prompt charged pions, prompt charged kaons, prompt (anti)protons and other prompt long-lived charged particles. In the following, the latter are referred to as other particles.

In each kinematic interval and for each of the two charges, the efficiency is computed in simulation by dividing the number of candidate tracks at reconstruction level associated with signal particles by the number of signal particles at generator level. The efficiencies for the four particle classes are shown in Figure 5.9. The charged-pion, charged-kaon and (anti)proton efficiencies are similar and reach values of approximately 80 % in the $\eta \in [2.0, 4.0)$ range at high p_T . They decrease towards low p_T because such particles are more likely to interact with the detector material. Especially at low p_T , there are differences between the three hadron species due to their different interaction lengths. Asymmetries between the two charges are evident for (anti)protons, which are induced by different interaction probabilities of quarks and antiquarks with the material; the efficiency for antiprotons is generally lower than that for protons. In the $\eta \in [3.0, 3.5)$ interval, dips in the efficiencies start to develop around $p_T \approx 500$ MeV/ c , becoming more pronounced in the subsequent η intervals. This reduction in efficiency is because of an increased amount of material that particles with trajectories corresponding to these kinematic intervals pass.




Figure 5.9: Efficiencies in simulation for the four particle classes in intervals of η and p_T . The efficiencies are multiplied by the charge to visually separate the corresponding values.

Details on the material budget of the detector can be found in Reference [128]; the interface section between the VELO vacuum vessel and the beam pipe increases the material budget in the $\eta \in [3.5, 4.3)$ range, and the conic beam pipe in the first RICH detector limits the acceptance at small angles, around $\eta \approx 4.4$. The shape of the other-particle efficiency differs much from the shapes of the charged-pion, charged-kaon and (anti)proton efficiencies. This is caused by the fractions of the particles in this class and their efficiencies. Electrons are the dominant species at low p_T and have an efficiency in this kinematic region of approximately 30% to 50%. This explains the increased other-particle efficiency at low p_T compared to the region around $p_T \approx 2000$ MeV/c, where the short lifetimes of the baryons in this class decrease the efficiency. Muons exhibit the highest efficiency among all particle species owing to their low interaction probability, but their contribution to the other-particle efficiency is small since the fraction of prompt muons is small. The fractions of the baryons are small at low p_T and slowly grow towards high p_T . Their efficiencies are low up to around $p_T \approx 3000$ MeV/c and then commence to increase due to the Lorentz boost of their lifetimes, enabling these baryons to be reconstructed as long tracks. This causes the other-particle efficiency to rise towards high p_T . The combined efficiency for prompt long-lived charged particles is close to

and slightly lower than the charged-pion efficiency as the particle composition is dominated by charged pions.

Deviations of the efficiency in simulation from that in data are taken into account by a correction according to Equation (5.4). This correction is the product of a track-reconstruction correction and a particle-composition correction. Hadronic interactions of charged pions, charged kaons and (anti)protons with the detector material induce an additional uncertainty on the correction because of the uncertainty on the material budget in simulation. The uncertainty on the noncorrected efficiency in simulation is purely statistical and originates from the sizes of the simulation samples. It is uncorrelated between the two charges.

5.5.1 Track reconstruction

Hit inefficiencies, which can be caused by nonfunctional detector channels, are included in simulation [129]. The positions of the active detector elements agree between data and simulation at the level of 0.5 mm. These small discrepancies do not affect the track-reconstruction efficiency. To investigate potential residual mismodelling in simulation, the track-reconstruction efficiency needs to be measured on data.

Such a control study was performed centrally by LHCb collaborators and is documented in Reference [129]. The result is a correction factor to be applied to the track-reconstruction efficiency in simulation, which is adopted in the analysis. The method uses $J/\psi \rightarrow \mu^+ \mu^-$ decays to probe the reconstruction in the VELO and the IT-OT stations, on which the long-track reconstruction is based. The decay chosen has a large yield and small background contamination. One of the two muon candidates is required to be a long track, whereas the other is obtained by combining track segments from the muon stations with hits in the TT. Two invariant-mass distributions are then fitted: J/ψ candidates for which the muon-station-TT track can be matched to a long track via a minimum overlap of hits and J/ψ candidates for which both decay products are reconstructed as long tracks. Dividing the J/ψ yields by each other results in the efficiency of the long-track reconstruction. This efficiency is computed in intervals of the kinematics of the muon-station-TT track both in data and simulation, and the efficiency ratio is the correction factor reported. Its uncertainty is of statistical origin, but constitutes a systematic uncertainty in the analysis because it stems from control samples. The efficiencies in data were determined separately for each of the two charges and found to be consistent. Consequently, a combined value for the two charges is reported. Another relative systematic uncertainty of 0.8% contributes to the uncertainty on the correction factor. It was estimated in the control study by using alternative occupancy proxies for the simulation weighting [129].

The correction factor is reported in two intervals of η and five intervals of momentum, while the analysis is conducted in intervals of η and p_T . Therefore, a linear mapping is established that translates the initial correction factor into the kinematic intervals in the analysis. The mapping is defined by a four-index generalised matrix constructed from the candidate tracks in data and quantifying the overlap fractions of the initial intervals with the intervals of η and p_T . This approach takes into account the nonuniform track density. The mapped correction factor is shown in Figure 5.10. It represents the track-reconstruction correction to the efficiency in simulation. The covariance matrix of the correction is obtained by propagating the uncertainty on the initial correction factor. In many kinematic intervals, the correction is compatible with unity within one standard deviation, reflecting that the track reconstruction is in general reliably modelled in simulation. The lowest-momentum interval of the initial correction factor was extrapolated into and assigned an efficiency ratio of 1.00 ± 0.05 . The step-like structures discernible in the figure are caused by the wide momentum intervals of the initial correction factor. The mapping smoothens the structures, but they remain visible. The minimum momentum of $5 \text{ GeV}/c$ required for the muon-station- TT track restricts the lower ends of the accessible p_T ranges in the analysis compared to the momentum threshold of $2 \text{ GeV}/c$ induced by the field of the LHCb magnet. Since each initial interval overlaps with several intervals of η and p_T , the covariance matrix contains positive statistical correlations although the statistical uncertainty on the initial correction factor is uncorrelated between the η and momentum intervals. The uncertainty on the correction is fully correlated between the two charges.

5.5.2 Particle composition

The efficiency depends on the particle composition because of the different lifetimes and different interaction lengths of the various long-lived charged particles. Hence, the particle composition in simulation is corrected to resemble that in data. At the time of the analysis, the composition had not yet been measured in p - p collisions at $\sqrt{s} = 13 \text{ TeV}$ in the forward η region. This has been done for the first time only later, namely in the second measurement performed within this thesis. For this reason, the correction is based on an LHCb measurement of ratios of prompt hadron production in p - p collisions at $\sqrt{s} = 0.9 \text{ TeV}$ and $\sqrt{s} = 7 \text{ TeV}$ [130].

Firstly, the measured charged-kaon-to-charged-pion, (anti)proton-to-charged-pion and antiproton-to-proton ratios are extrapolated to $\sqrt{s} = 13 \text{ TeV}$. In each interval of η and p_T of the control study, each of the ratios is extrapolated separately with a function that is linear in $\ln \sqrt{s}$. A comparison of the extrapolated ratios with the corresponding ratios in simulation reveals that the (anti)proton-to-charged-pion ratio in simulation is larger by up to 40% than the expectation from data. Better agreement is found for the other two ratios. The control study does not cover the




Figure 5.10: Track-reconstruction correction to the efficiency in simulation in intervals of η and p_T . The bars indicate the propagated uncertainty. The gaps between the right ends of the light-grey rectangles at low p_T , representing the minimum momentum of 2 GeV/c needed for the long-track reconstruction, and the left markers are caused by the minimum-momentum requirement of 5 GeV/c applied in the control study [129]. In the following, these gaps occur in all plots in which the track-reconstruction correction enters.

lowest- and highest- η intervals in the analysis and was only performed in three wide p_T intervals. Therefore, double ratios are formed of the extrapolated ratios and those in simulation, which mitigates the dependence on η and p_T . The uncovered η intervals are equipped with the double ratios from the adjacent η intervals.

Afterwards, a corrected number of particles in simulation is defined for each of the four particle classes. This number is given by the product of the initial number in simulation and a correction function that is linear in $\log_{10} p_T$. Another set of double ratios is computed from the hadron ratios of the corrected and initial numbers. In each η interval and for each of the two charges, these double ratios are fitted to the double ratios from the extrapolations. The fits use a cost function of least squares and weak Gaussian constraints to suppress deviations of the corrected numbers from the initial ones above 5%. Moreover, it is respected that the total number of prompt long-lived charged particles in simulation must not be changed.

The corrected efficiency is obtained by adding the products of the efficiencies for the four classes in simulation and their corrected fractions relative to the total

number of particles. The ratio of the corrected efficiency and the initial efficiency in simulation represents the particle-composition correction. It is shown in Figure 5.11. The correction ranges between 0.96 and 1.02 in almost all kinematic intervals. The deviation from unity is largest for negatively charged particles at high p_T . This is due to a reduction in the antiproton fraction and an increase in the fraction of the negatively charged other particles. Statistical uncertainties are neglected in the determination of this correction as a considerable number of assumptions are involved. A systematic uncertainty on the correction is hence estimated as half the absolute value of the deviation of the correction from unity. This uncertainty is uncorrelated between the two charges.

5.5.3 Hadronic interactions

The control study for the track-reconstruction efficiency is based on muons. Thus, effects from hadronic interactions cannot be taken into account. However, such effects are relevant for charged pions, charged kaons and (anti)protons because the material budget of the detector, which induces these interactions, has a relative systematic uncertainty of 10 % in simulation. This value is explained in detail in Reference [129], where it was found that approximately 14 % of the charged pions and 11 % of the charged kaons cannot be reconstructed due to hadronic interactions occurring upstream of the third IT–OT station. Both of these loss fractions are mean values for the two charges. Since the loss fractions are proportional to the material budget, the track-reconstruction efficiency has an additional systematic uncertainty of 1.4 % and 1.1 % of the muon efficiency in simulation for charged pions and charged kaons, respectively.

For (anti)protons, the loss can be obtained from the (anti)proton and muon efficiencies in simulation. As (anti)protons do not decay, any additional loss compared to muons must be caused by hadronic interactions. The loss fraction computed ranges from 10 % to 30 %, depending on the kinematic interval. More antiprotons than protons are lost, which is expected in a detector made of matter. The loss is larger at low p_T since low-momentum particles are more likely to interact.

The combined uncertainty for long-lived charged particles is determined by multiplying the individual uncertainties for the three hadron species by their corrected fraction from Section 5.5.2 and forming the sum in quadrature. No individual uncertainty is assumed for the class of other particles; they are either leptons, or their loss is dominated by decay. The combined relative uncertainty is almost independent of the kinematic interval and amounts to 1 %. This systematic uncertainty on the track-reconstruction correction from hadronic interactions is fully correlated between the two charges.




Figure 5.11: Particle-composition correction to the efficiency in simulation in intervals of η and p_T . The values for the two charges are shown. The bars indicate the systematic uncertainty.

5.5.4 Explicit unfolding of interval-to-interval migration

The finite resolution of the detector distorts the real η and p_T distributions. Hence, the measured values differ from the real ones by random offsets. This leads to smearing and can induce migration between the kinematic intervals. Thanks to the high momentum resolution of the detector, the amount of smearing is small. It is described by [131]

$$f(\eta', p_T') = \int K(\eta', p_T'; \eta, p_T) g(\eta, p_T) d\eta dp_T, \quad (5.7)$$

where f and g respectively denote the (background-free) measured spectrum and the real spectrum of the signal particles. The resolution kernel, K , is a probability density function over the measured kinematics, which generally depends on the real kinematics.

The resolution kernel is estimated on simulation via the reconstructed and generated kinematics of the candidate tracks associated with signal particles. The amount of distortion induced by K is given by $f_{\text{sim}}/g_{\text{sim}}$, where the numerator is determined by replacing the integral in Equation (5.7) by a sum. The ratio deviates from unity by up to 2%, meaning that interval-to-interval migration cannot be neglected.

However, this effect is already taken into account, namely by the efficiency in simulation. By definition, the inverse of the efficiency, which appears in Equation (5.4), is equal to $g_{\text{sim}}/f_{\text{sim}}$. This ratio can be used to correct for distortion. Consequently, interval-to-interval migration is implicitly corrected for by the determination of the efficiency in simulation. Applying another unfolding in addition would cover the effect twice.

The implicit approach is not guaranteed to be unbiased. The drawback is that the distortion does not only depend on K , but also on the shape of g , corresponding to n in Equation (5.4). A steeper shape of g increases interval-to-interval migration, and assuming that the shape of g is perfectly reproduced in simulation is not justified. Therefore, the resolution kernel is also applied to g , resulting in an approximation of f . The $(f/g)/(f_{\text{sim}}/g_{\text{sim}})$ double ratio quantifies the bias caused by implicitly using $f_{\text{sim}}/g_{\text{sim}}$ instead of f/g , with the latter estimating the correct value. The deviations of the double ratio from unity are below 1%. This signifies that the relative distortion of up to 2% is at most biased by 1%. Thus, the relative systematic uncertainty of employing g_{sim} instead of g is less than 0.02% and hence negligible.

5.6 Propagation of uncertainty

The analysis aims at accurately including all statistical and systematic correlations. The statistical and systematic uncertainties on the results are obtained with analytical propagation of uncertainty. The quantities entering Equations (5.1) and (5.4), on which the results are based, are stored as three-index generalised vectors of the values in all kinematic intervals and for the two charges. The covariances thus need to be represented by six-index generalised matrices. This allows correlations between kinematic intervals and the two charges to be included as off-diagonal elements in the covariance matrices. Since the equations do not contain terms that mix contributions from various kinematic intervals or the two charges, the propagation of all individual uncertainties can be independently performed for each combination of η , p_T and charge. The results are given as the charge-combined differential cross-section and the charge ratio in intervals of η and p_T . In the matrix representation chosen, the covariance matrices of both the sum and the ratio can be succinctly obtained via contractions over the charge indices.

The types of correlation of the statistical and systematic uncertainties between the two charges have already been mentioned in Sections 5.4 and 5.5. The uncertainties are either fully correlated or uncorrelated. The integrated luminosities of the two data samples possess a relative uncertainty of 2.0%, which is systematic as it was determined independently of the analysis in a calibration study. This uncertainty is fully correlated between the two charges and between the two data samples. Correlations also exist between the kinematic intervals, *e.g.* due to the proxy ratio

for strange-hadron decays. They are not explicitly stated, but propagated into the results as well.

5.6.1 Candidate-count correlations

The prompt production of a long-lived charged particle is not a rare process, but frequently occurs in each event. It leads to a candidate count that varies from event to event. Consequently, the counts that each event contributes to various kinematic intervals are not independently Poisson distributed. Positive statistical correlations between the kinematic intervals (and between the two charges) are induced.

An estimate of the covariance matrix of the counts that reflects this effect is obtained from a subsampling approach. The subsampling is based on the independence and identical distribution of the events. The sample under consideration is divided into 100 subsamples by assigning the 1st, 101st, 201st *etc.* event to the first subsample; the 2nd, 102nd, 202nd *etc.* event to the second subsample and so forth. This interleaved definition of the subsamples avoids time ordering and makes them equivalent and independent realisations of the underlying random process. Thus, an unbiased sample covariance matrix can be determined according to

$$C = \frac{1}{k-1} \sum_{j=1}^k (\vec{n}_j - \langle \vec{n} \rangle) \cdot (\vec{n}_j - \langle \vec{n} \rangle)^\top. \quad (5.8)$$

The number of subsamples is denoted by k , the three-index generalised vector formed by the counts assigned to each subsample, j , in all kinematic intervals and for the two charges is denoted by \vec{n}_j , and the sample mean is denoted by

$$\langle \vec{n} \rangle = \frac{1}{k} \sum_j \vec{n}_j. \quad (5.9)$$

The covariance matrix, $C(\langle \vec{n} \rangle)$, of the sample mean is defined by C/k . The quantity of interest is the covariance matrix of the sum vector, $\vec{n} = \sum_j \vec{n}_j = k \langle \vec{n} \rangle$. It is given by

$$C(\vec{n}) = \left(\frac{d\vec{n}}{d\langle \vec{n} \rangle} \right)^2 C(\langle \vec{n} \rangle) = k^2 \frac{C}{k} = kC, \quad (5.10)$$

which follows from propagation of uncertainty. The subsampling variance of the candidate counts is increased by up to 50% compared to the Poisson expectation, caused by the positive correlations between the kinematic intervals. The largest correlations amount to 0.4 and occur around $p_T \approx 400 \text{ MeV}/c$, where the track density is largest. Some negative correlations are observed, most of them at high η and high p_T . They must be due to fluctuations since the underlying random process is only capable of creating positive correlations. The subsampling covariance matrix is also

computed for the other counts that enter Equation (5.4), which are the numbers of background tracks in simulation and the number of background tracks from beam–gas interactions. It follows that the statistical uncertainties on the counts are not uncorrelated, but partially correlated between the kinematic intervals and between the two charges.

5.6.2 Total uncertainty

The total uncertainty on the number of signal particles comprises all statistical and systematic uncertainties except for that on the integrated luminosity, which only enters when the differential cross-section is determined. A breakdown of the total uncertainty on the charge-combined number into the largest individual contributions is shown in Figure 5.12. The leading uncertainty in most kinematic intervals is that on the correction to the efficiency in simulation. It has a relative contribution of 5.1% at low p_T , which mainly stems from the uncertainty on the track-reconstruction correction. The largest individual uncertainty of 8.2% occurs at high p_T and is caused by the proxy ratio for fake tracks. A similarly large contribution exists at high η and low p_T due to the proxy ratio for material interactions. It amounts to 8.1% as the background from photon conversions is large in this kinematic region as well as the uncertainty on the proxy ratio, initially only designed for material interactions of charged pions. The remaining sources of uncertainty are each smaller than 2% in all kinematic intervals. The total relative uncertainty ranges from 1.1% to 10.9%.

In almost all kinematic intervals, the statistical uncertainties are small. They solely contribute at high p_T , where the uncertainties on the efficiency in simulation, on the number of candidate tracks and on the number of fake tracks in simulation are nonnegligible. Consequently, the analysis is dominated by systematic uncertainties. The uncertainties on the results are therefore not divided into statistical and systematic contributions. Instead, the total uncertainty is given, *i.e.* the sum in quadrature of statistical and systematic uncertainties.

5.6.3 Data-sample combination

All steps of the analysis are executed separately for each of the two data samples. This permits a cross-check between the results obtained independently for the two magnetic-field configurations. The results should agree within the systematic uncertainties assigned because a physics quantity must not depend on a detector configuration. The cross-check requires the charge-combined differential cross-sections to be computed for the magnet-up and magnet-down data samples and for the combination of both data samples.

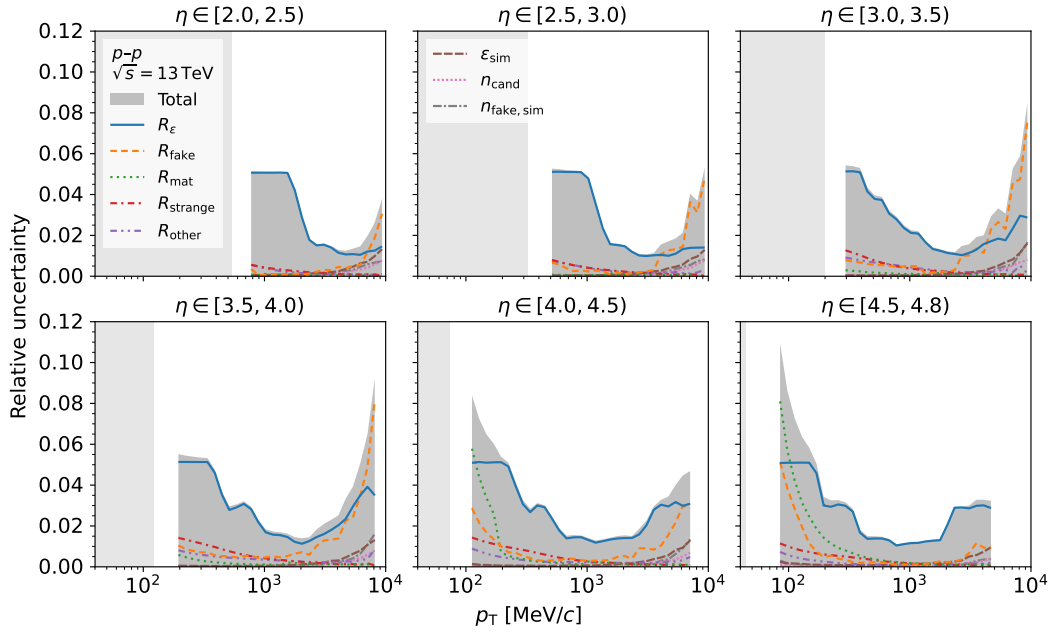


Figure 5.12: Breakdown of the total relative uncertainty on the number of prompt long-lived charged particles into the largest individual contributions in intervals of η and p_T . The uncertainties refer to the charge-combined number. Statistical and systematic uncertainties are added in quadrature if both exist. The contributions shown are the correction to the efficiency in simulation, given by the product of the track-reconstruction and particle-composition corrections; the proxy ratios constructed for fake tracks, material interactions and strange-hadron decays; the relative systematic uncertainty of 50% assumed for the nonadjusted other sources of background as well as the statistical uncertainties on the efficiency in simulation, on the number of candidate tracks and on the number of fake tracks in simulation. The relation between these quantities is defined by Equation (5.4).

The combined result is determined by adding the numbers of signal particles obtained from the two data samples and dividing by the sum of the integrated-luminosity values as well as by the product of the kinematic-interval widths. The statistical uncertainties are uncorrelated, and even independent, between the two data samples. Thus, the combined statistical covariance matrix is computed by adding the two individual statistical covariance matrices. The systematic uncertainties on the numbers of signal particles are assumed to be fully correlated between the two data samples. Furthermore, the systematic uncertainty on the integrated luminosity is by definition fully correlated between the two data samples. The corresponding correlation terms are added to the sums of the individual systematic covariance matrices to compute the combined systematic covariance matrix. The

relative deviations of the individual results from the combined one are found to be less than 0.2% in most kinematic intervals, one order of magnitude smaller than the combined relative systematic uncertainty. Differences between the individual results at high p_T are explained by the statistical uncertainties and hence represent fluctuations. At high η and low p_T , systematic deviations of up to 3% from the combined result occur, but they are well covered by the combined systematic uncertainty of up to 11%. This means that the individual results are consistent. Consequently, the cross-check is successful, and using the combined result is justified.

5.7 Results

The charge-combined differential cross-section and the charge ratio are determined from the combination of both data samples. They are thus based on a sample of p - p collisions at $\sqrt{s} = 13$ TeV that corresponds to an integrated luminosity of 5.4 nb^{-1} . Since the measured quantities are fully corrected for inefficiency and background, they can directly be compared with predictions from hadronic-interaction models.

5.7.1 Differential cross-section

The measured charge-combined differential cross-section of prompt production of long-lived charged particles, computed according to Equations (5.1) and (5.4), is shown in Figure 5.13. The data cover p_T values from 80 MeV/ c to 10 000 MeV/ c , where the largest p_T range is accessible in the $\eta \in [4.5, 4.8)$ interval. The total relative uncertainty on the measured differential cross-section, which includes the uncertainty on the integrated luminosity, is between 2.3% and 11.1%.

The result is compared with predictions from PYTHIA 8.303 [38] with the default Monash 2013 tune [40], EPOS-LHC [26], QGSJET II.04 [33] and SIBYLL 2.3d [35]. Each predicted differential cross-section is calculated as

$$\frac{d^2\sigma_{\text{model}}}{d\eta dp_T} = \frac{\sigma_{\text{inel}} n_{\text{model}}}{N_{\text{inel}} \Delta\eta \Delta p_T} \quad (5.11)$$

from the inelastic cross-section, σ_{inel} ; the number, N_{inel} , of generated inelastic events and the number, n_{model} , of signal particles obtained. The values of σ_{inel} implemented in the models are listed in Table 5.1. These values are found to be compatible with measurements, given in the same table, within less than one standard deviation. They are therefore inserted in Equation (5.11). Residual deviations in magnitude of the model predictions from the data can then only be caused by the predicted distributions of n_{model} . The number of generated inelastic events [132] is 10^8 for each model. These events are also used for the model comparisons in Section 5.7.2. The models mostly overestimate the measured differential cross-section. The overall

5 Prompt charged-particle production

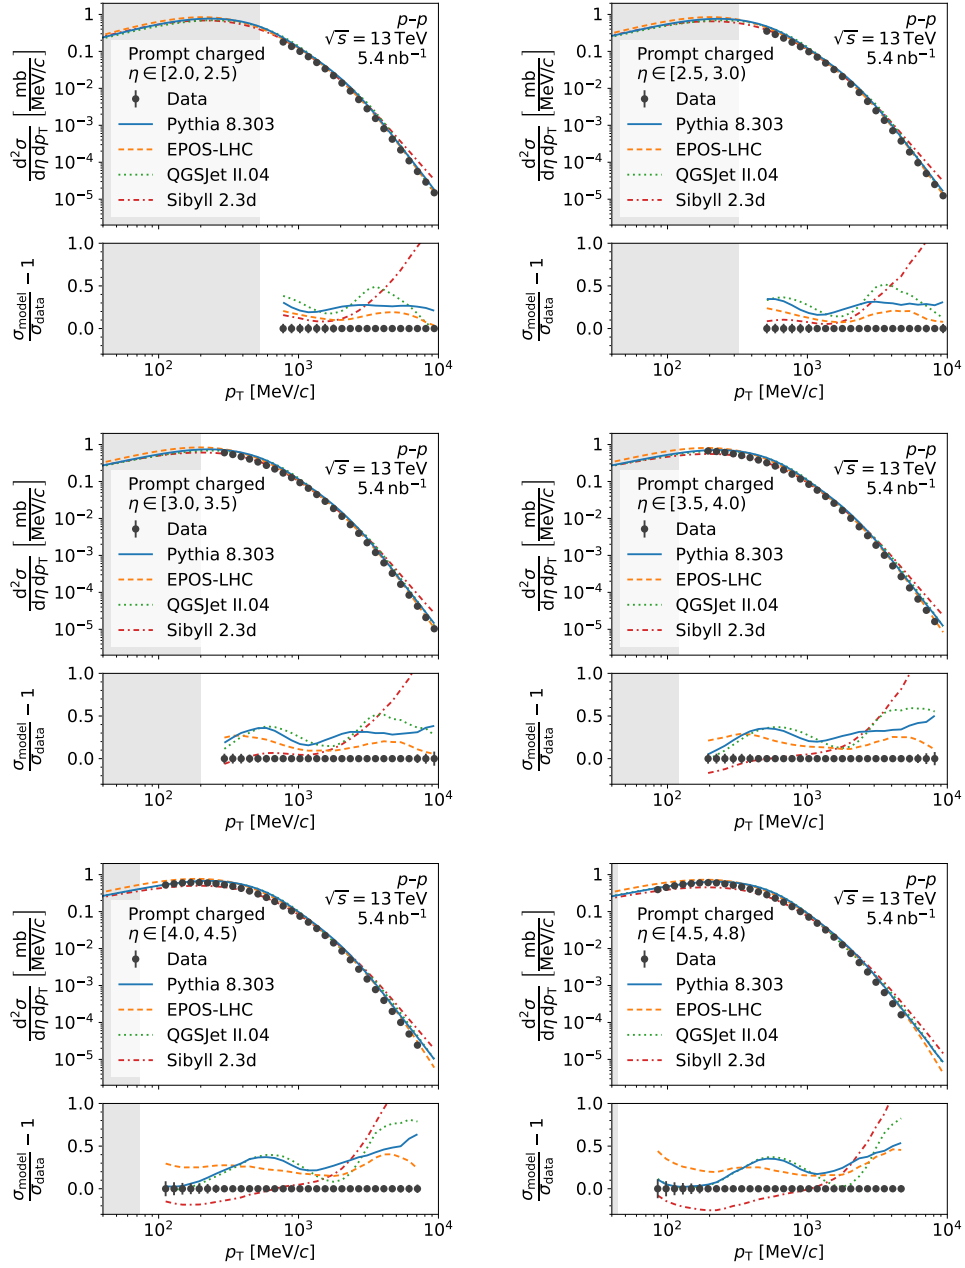


Figure 5.13: Differential cross-section of prompt production of long-lived charged particles in p - p collisions at $\sqrt{s} = 13$ TeV in intervals of η (from top left to bottom right) and p_T measured in the analysis and predicted by various hadronic-interaction models. The charge-combined values are shown. The relative deviations of the models from the data are also displayed. The bars indicate the total uncertainty.

Table 5.1: Inelastic cross-sections implemented in various hadronic-interaction models and measured by various experiments for p - p collisions at $\sqrt{s} = 13$ TeV.

Model or experiment	Inelastic cross-section [mb]
PYTHIA 8.303 [38, 40]	78.05
EPOS-LHC [26]	78.98
QGSJET II.04 [33]	80.17
SIBYLL 2.3d [35]	79.86
ATLAS [133]	78.1 ± 2.9
LHCb [134]	75.4 ± 5.4
TOTEM [135]	79.5 ± 1.8

best description of the data is provided by EPOS, with the relative deviations lying between +4 % and +46 %; the modelled shape reasonably agrees with the data, and the offset of the modelled magnitude weakly depends on η . PYTHIA and QGSJET tend to exhibit larger deviations from the measured differential cross-section than EPOS at high p_T , respectively up to +64 % and +83 %, but they are closer to the data at low p_T . The largest deviations occur for SIBYLL, amounting to -25 % around $p_T \approx 200$ MeV/ c and $+172$ % around $p_T \approx 9000$ MeV/ c . Except for a few kinematic intervals, all predictions differ significantly from the data. The total uncertainty is also small compared to the spread of the predictions, signifying that this measurement constitutes an important input to improve the models in the future.

The total correlation matrix of the measured differential cross-section of prompt production of long-lived charged particles, comprising all statistical and systematic contributions, is shown in Figure 5.14. The correlations are positive as the analysis is dominated by the systematic uncertainties, which are often fully correlated between the two charges or between kinematic intervals. The rectangular structures of increased correlation are due to the wide η and momentum intervals of the correction factor on which the track-reconstruction correction is based. Those wide intervals induce large correlations between adjacent kinematic intervals in the analysis; the large correlations between the two charges are also primarily caused by this correction. The correlation matrix distinctively reflects the effort required to accurately propagate all correlations.

5.7.2 Charge ratio

The ratio of the measured differential cross-sections of prompt production of long-lived positively and negatively charged particles, computed according to Equations (5.4) and (5.5), is shown in Figure 5.15. All uncertainties that are fully correl-

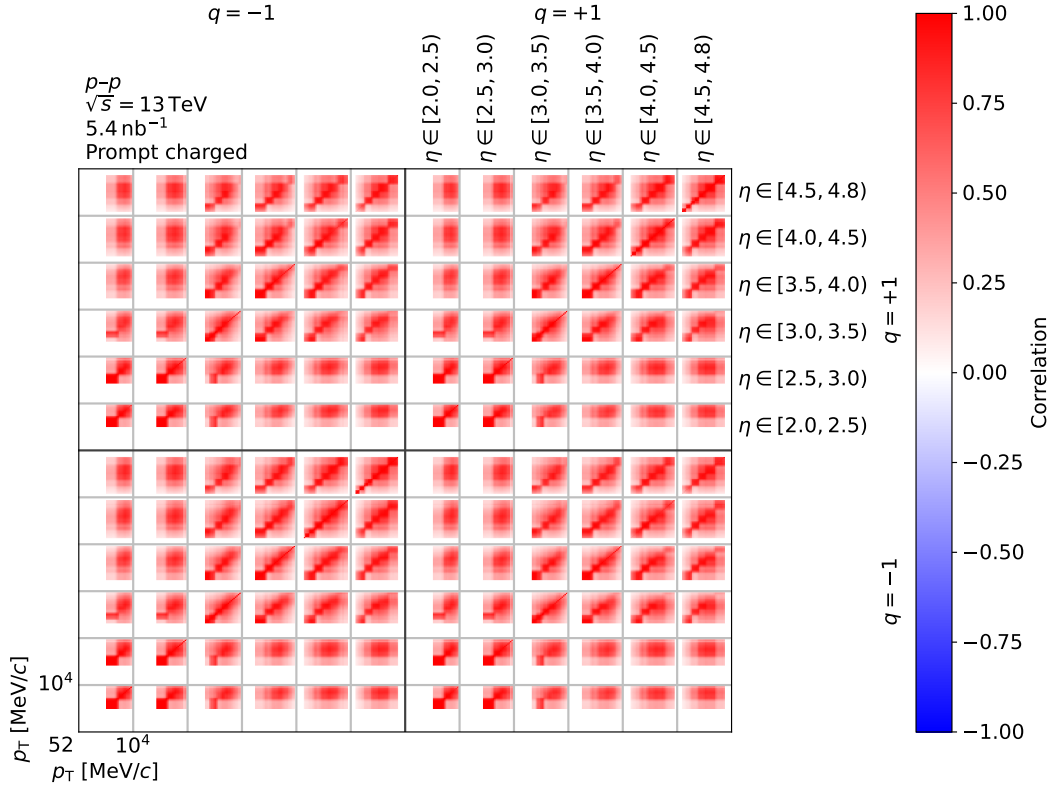


Figure 5.14: Total correlation matrix of the measured differential cross-section of prompt production of long-lived charged particles. The correlations are arranged in three dimensions. The four large quadrants correspond to the correlations between the two charges; the 36 cells in each quadrant indicate the correlations between the η intervals; lastly, in each cell, the correlations between the p_T intervals are shown.

ated between the two charges cancel and hence decrease the propagated uncertainty on the charge ratio. This applies, for example, to the uncertainties on the integrated luminosity and on the track-reconstruction correction. The impact of the latter is evident at low η and low p_T , where the total uncertainty on the charge ratio is much smaller than that on the charge-combined differential cross-section. The measured charge ratio is close to unity in many kinematic intervals, which is predicted by the models. Although the total uncertainty grows towards high p_T , a systematic rise of the data is observed, which is most pronounced in the $\eta \in [4.0, 4.8)$ range. This can be explained by the excess of positive charge in the initial state of a p - p collision, transferring to the final state in the forward η region. The measured rise is reproduced by PYTHIA to some extent. By contrast, the increase predicted by QGSJET is too steep and starts to emerge too low in p_T . There is tension between the flat EPOS

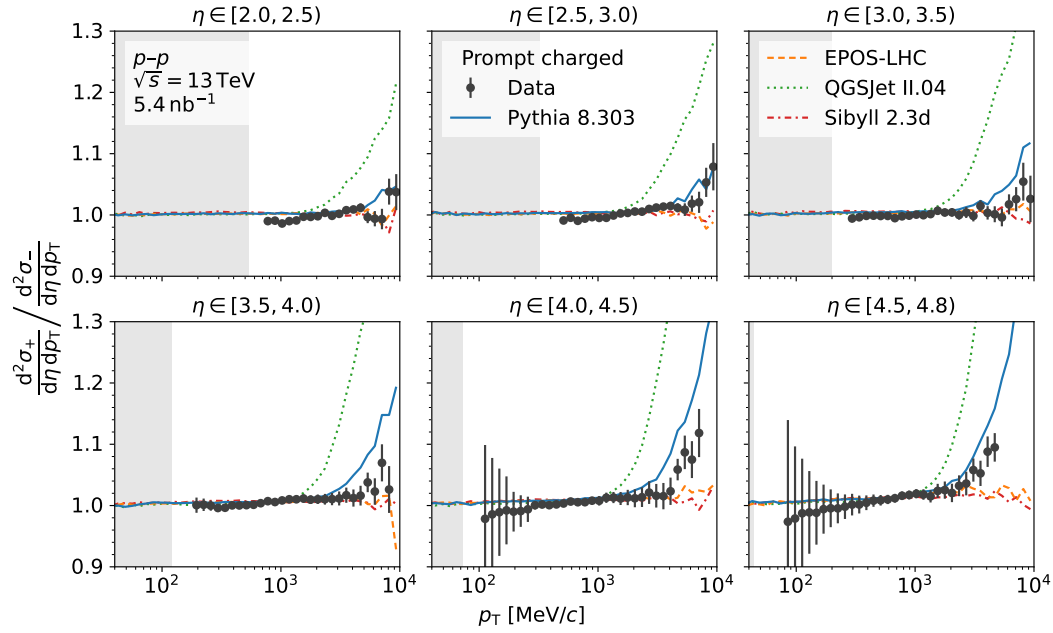


Figure 5.15: Ratio of the differential cross-sections of prompt production of long-lived positively and negatively charged particles in intervals of η and p_T measured in the analysis and predicted by various hadronic-interaction models. The bars indicate the total uncertainty.

and SIBYLL predictions and the data at high η and high p_T . Moreover, the measured charge ratio decreases towards low p_T in the $\eta \in [4.0, 4.8)$ range. This is not expected from the models and probably caused by the proxy ratio for material interactions. The background from photon conversions is large in this kinematic region. It is adjusted with the proxy ratio designed for material interactions of charged pions, which cannot be directly measured at high η , but is instead assigned a median value as described in Section 5.4.2. To reflect that potential charge asymmetries are not taken into account, the uncertainty on the median value is assumed to be uncorrelated between the two charges. Thus, it does not cancel in the measured charge ratio, but leads to a total uncertainty that renders the data compatible with unity.

6 Prompt identified-hadron production and nuclear modification

In this chapter, the second measurement performed within this thesis is presented. It has not yet been published, but is currently in an advanced state of the LHCb-internal review. The differential cross-sections of prompt production of charged pions, charged kaons, (anti)protons and long-lived charged particles in p -Pb collisions at $\sqrt{s_{\text{NN}}} = 8.16$ TeV and p - p collisions at $\sqrt{s} = 13$ TeV are measured with the LHCb experiment. For the first time, identified-hadron production in p -Pb collisions is measured in the forward η region, where η denotes the pseudorapidity (in the centre-of-mass system). Moreover, the analysis is the highest-energy measurement of identified-hadron production in the forward η region of p - p collisions to date. Like the first measurement of this thesis, it is conducted in intervals of η and p_{T} . The two particle charges are again studied separately. Furthermore, charged-kaon-to-charged-pion, (anti)proton-to-charged-pion and (anti)proton-to-charged-kaon ratios are measured as well as nuclear-modification factors and forward-to-backward ratios. Previously, identified-hadron production in the forward η region was measured by the LHCb collaboration at lower energies; ratios were determined of identified-hadron production in p - p collisions at $\sqrt{s} = 0.9$ TeV and $\sqrt{s} = 7$ TeV [130]. By contrast, absolute production cross-sections are determined in this measurement as well. Charged-particle production in p -Pb collisions at $\sqrt{s_{\text{NN}}} = 5$ TeV was also measured by the LHCb collaboration [117].

The second measurement of this thesis extends the first one, which was solely performed with information from the track-reconstruction system, in two ways. Analysing the hadron composition is now enabled by also using PID information. A comprehensive calibration method is developed to disentangle charged pions, charged kaons and (anti)protons based on data. With PID information, a separate proxy is constructed for photon conversions. In addition, data from p -Pb collisions are included to study light-hadron production in interactions of protons with heavier ions. Several improvements to techniques established in the first measurement are made as well, which concern the proxy ratios for fake tracks and strange-hadron decays and the particle-composition correction. The production of charged particles in p - p collisions at $\sqrt{s} = 13$ TeV has already been analysed in the first measurement, but is reanalysed to benefit from the improvements.

6.1 Data and simulation samples

The analysis is based on two data samples of p -Pb collisions, which were recorded in November 2016 at $\sqrt{s_{\text{NN}}} = 8.16$ TeV during two single periods of stable beam conditions in the LHC. The hardware stage of the trigger system accepted events at a fixed rate. At the software stage, events were only retained if minimum activity was registered in the detector; at least one track segment had to be reconstructed in the VELO. In addition, a requirement was imposed on the maximum number of VELO hits to veto events with considerable processing time. This trigger selection is found to be fully efficient on data. Thus, the p -Pb data samples can be regarded as trigger unbiased. The p -Pb collision system is asymmetric, and the production of particles in the direction of the proton beam is expected to differ from that in the direction of the lead-ion beam. However, the detector only covers one side of the interaction point. Therefore, events are collected in two beam configurations, obtained by switching the protons and ions between the two LHC beam pipes. The configuration with the proton beam travelling from the VELO towards the muon stations is referred to as *forward*, whereas that with the lead-ion beam travelling in this direction is termed *backward*. Each of the two p -Pb data samples was recorded in one of these two beam configurations. The forward and backward p -Pb data samples approximately comprise 187×10^6 and 139×10^6 events, respectively, and correspond to integrated luminosities of $90 \mu\text{b}^{-1}$ and $71 \mu\text{b}^{-1}$. Both p -Pb data samples were collected in the magnet-down configuration of the LHCb magnet. In simulation, the p -Pb collisions are generated with EPOS-LHC [26]. The remainder of the simulation chain does not differ from the description in Section 5.1. Two p -Pb simulation samples are part of the analysis, one for each beam configuration. Each of these samples approximately contains 4×10^6 events.

For p - p collisions, the same magnet-up data sample as in the first measurement is used, for which the background from beam-gas interactions was found to be negligible. The previous magnet-down data sample is replaced, however, by a sample for which this background is negligible as well. The present magnet-down data sample, stemming from the same period of data recording as the previous one, consists of 102×10^6 events and corresponds to an integrated luminosity of 1.6 nb^{-1} . The same two p - p simulation samples as in the first measurement are employed.

6.2 Strategy

The strategy of the first measurement, which has been presented in Section 5.2, is followed. The differential cross-section and the charge ratio are determined according to Equations (5.1) and (5.5), respectively, but the relation for the number of signal particles has to be modified. Since no contamination arises from beam-gas

interactions, this background does not need to be subtracted from the number of candidate tracks. In addition, an acceptance correction, C_{acc} , occurs for p -Pb collisions as the centre-of-mass system, in which the results shall be given, does not coincide with the laboratory system. This is detailed in Sections 6.3 and 6.7. Hence, the expression in Equation (5.4) is replaced by

$$n = \frac{C_{\text{acc}}}{R_{\varepsilon} \varepsilon_{\text{sim}}} \left(n_{\text{cand}} - \sum_i R_i n_{i, \text{sim}} \right), \quad (6.1)$$

with the meaning of the other variables remaining unchanged. In this measurement, either prompt charged pions, prompt charged kaons, prompt (anti)protons or prompt long-lived charged particles constitute the signal particles. Like in the first measurement, the number of candidate tracks for prompt long-lived charged particles is determined by counting the tracks passing the selection. The numbers of candidate tracks for prompt charged pions, charged kaons and (anti)protons are obtained from template fits to two-dimensional distributions of PID variables. This procedure is described in Section 6.5. The corresponding numbers of signal particles allow the charged-kaon-to-charged-pion, (anti)proton-to-charged-pion and (anti)proton-to-charged-kaon ratios to be computed according to

$$\begin{aligned} & \left(\frac{d^2 \sigma_{K^+; p; p}}{d\eta dp_T} + \frac{d^2 \sigma_{K^-; \bar{p}; \bar{p}}}{d\eta dp_T} \right) \bigg/ \left(\frac{d^2 \sigma_{\pi^+; \pi^+; K^+}}{d\eta dp_T} + \frac{d^2 \sigma_{\pi^-; \pi^-; K^-}}{d\eta dp_T} \right) \\ &= \frac{n_{K^+; p; p} + n_{K^-; \bar{p}; \bar{p}}}{n_{\pi^+; \pi^+; K^+} + n_{\pi^-; \pi^-; K^-}}. \end{aligned} \quad (6.2)$$

In these hadron ratios, uncertainties cancel that are fully correlated between the three hadron species. Nuclear effects are quantified by the nuclear-modification factor, resulting from the differential cross-sections in p -Pb and p - p collisions as

$$R_{p\text{-Pb}}(\eta, p_T) = \frac{d^2 \sigma_{p\text{-Pb}}}{d\eta dp_T} \bigg/ \left(208 \times \frac{d^2 \sigma_{p\text{-}p}}{d\eta dp_T} \right), \quad (6.3)$$

where 208 is the lead mass number. If both differential cross-sections are measured at the same (nucleon–nucleon) centre-of-mass energy, a value of $R_{p\text{-Pb}} = 1$ signifies the absence of nuclear effects; the production in p -Pb collisions can then be regarded as superposition of the production in 208 single p - p collisions. Another widely used quantity is the forward-to-backward ratio,

$$R_{\text{FB}}(|\eta|, p_T) = \frac{d^2 \sigma_{p\text{-Pb}}(|\eta|, p_T; \eta > 0)}{d\eta dp_T} \bigg/ \frac{d^2 \sigma_{p\text{-Pb}}(|\eta|, p_T; \eta < 0)}{d\eta dp_T}, \quad (6.4)$$

which compares the production in the $|\eta|$ intervals that can be accessed in both p -Pb beam configurations. The differential cross-section, the charge ratio, the nuclear-modification factor and the forward-to-backward ratio are each determined for the three hadron species and long-lived charged particles. This implies, together with the three hadron ratios and the p -Pb and p - p data samples, a large number of results.

6.3 Selection

The events retained by the trigger system can contain several primary vertices, which are formed by the reconstruction software. The number of primary vertices is a proxy for the number of p -Pb or p - p interactions per event, but not every real interaction corresponds to one primary vertex. An interaction producing substantially fewer particles than on average can be missed, while a high-multiplicity interaction might be misreconstructed as two primary vertices. In general, imposing a requirement on the number of primary vertices to increase the purity of events with exactly one first interaction is not needed as the integrated luminosity takes into account the number of first interactions. However, it is found for the p -Pb data samples that the spatial distributions of the primary vertices contain reconstruction artefacts, which are not present for the p - p data samples. An increased primary-vertex density is observed outside the expected interaction region. This is caused by secondary interactions that produce sufficiently many VELO track segments to be reconstructed as primary vertices. Requiring exactly one primary vertex per event removes these artefacts. This requirement is therefore applied to the p -Pb data and simulation samples. Since it reduces the number of events, the fraction of first interactions discarded has to be quantified to correspondingly decrease the integrated luminosity of the data sample under consideration. This fraction is deduced from a Poisson distribution for the number of simultaneous first interactions and a binomial distribution for the primary-vertex-reconstruction efficiency in simulation. It approximately amounts to 8% for both beam configurations and has a negligible uncertainty. The values of the integrated luminosity for the p -Pb data samples given in Section 6.1 have already been corrected by this fraction and can directly be inserted in Equation (5.1).

The track selection for the p -Pb samples closely resembles that for the p - p samples, which has been described in Section 5.3. Long tracks are only used as candidate tracks if their track fit is of good quality. The number of fake tracks is reduced by a requirement on the P_{fake} variable.

The beam energies of the p -Pb collisions are asymmetric; the protons possess an energy of 6.5 TeV, whereas that for the lead ions is 2.56 TeV per nucleon. Consequently, the centre-of-mass system of the collision moves relative to the laboratory system in the direction of the proton beam, *i.e.* approximately along the z axis. To

translate the measured track kinematics into the centre-of-mass system, a Lorentz boost needs to be applied, which depends on the kinematics and the particle mass. As the transverse momentum is perpendicular to the z axis, only the pseudorapidity in the laboratory system is affected by the boost. In data, this boost cannot be applied because the species of a single track is not known. Hence, an approximation is made. The η_{lab} value of each candidate track is shifted by the constant rapidity shift, $\Delta_y = 0.5 \ln(208/82) \approx 0.465$, for massless particles, where 82 is the lead atomic number. The η values are defined relative to the direction of the proton beam. Thus, the forward beam configuration allows positive η values to be accessed, and the relation $\eta = \eta_{\text{lab}} - \Delta_y$ holds. The backward beam configuration gives access to negative η values according to $\eta = -(\eta_{\text{lab}} + \Delta_y)$. The acceptance correction takes into account that the rapidity shift is an approximation.

The η grid for the p - p samples is slightly modified compared to the first measurement; the highest interval is enlarged by 0.2, permitting the $\eta \in [2.0, 5.0)$ range to be covered, with all six interval widths amounting to $\Delta\eta = 0.5$. This ensures a large overlap with the η grids for the p -Pb samples, which is relevant for the determination of the nuclear-modification factors and the forward-to-backward ratios. Since the rapidity shift is close to 0.5, those η grids are obtained by shifting that for the p - p samples by 0.5. Consequently, the $\eta \in [1.5, 4.5)$ and $\eta \in [-5.5, -2.5)$ ranges are covered for the forward and backward beam configurations, respectively. The same p_{T} grid as in the first measurement is employed. Due to the track-reconstruction correction, the minimum accessible p_{T} values are restricted by a minimum track momentum of 6 GeV/ c for the p -Pb samples.

Each of the four simulation samples is weighted to reproduce the detector occupancy of its corresponding data sample. The number of long tracks per event is used as occupancy proxy, and the ratio of the distributions in data and simulation determines the weights. A second simulation weighting is found to mitigate the unexpected decrease towards low p_{T} in the charge ratio for prompt long-lived charged particles in p - p collisions at high η_{lab} , which has been discussed in Section 5.7.2. This weighting ameliorates the agreement of the spatial distributions of the primary vertices in simulation with those in data. Unlike the first weighting, the second one is not applied throughout, but only in the highest- η_{lab} interval when the efficiency is computed.

6.4 Background contributions

Simulation is used to study the composition of the background in p -Pb collisions. The largest sources of background for prompt long-lived charged particles are the same as in p - p collisions: fake tracks, photon conversions, material interactions of charged pions and K_{S}^0 decays. The kinematic regions in which they accumulate are

also similar. More fake tracks contribute to the candidate tracks in p -Pb collisions because of the detector occupancy, which is on average larger than in p - p collisions. For each of these four sources of background, a proxy is constructed to adjust the contribution in simulation to that in data. The background for prompt (anti)protons comprises tracks produced in Λ , $\bar{\Lambda}$, Σ^+ and $\bar{\Sigma}^-$ decays, which are adjusted analogously to those from K_S^0 decays. The remaining sources of background, which are small, are not adjusted and again assigned a relative systematic uncertainty of 50 %, which is fully correlated between the two charges.

6.4.1 Fake tracks

In the first measurement, the proxy for fake tracks was determined by counting the tracks in one of several intervals of the P_{fake} distribution above 0.3, where the selection requirement for the p - p samples is placed, with sufficiently high purity. A systematic uncertainty was computed by additionally including all P_{fake} intervals above the chosen one. Moreover, the proxy ratio was assumed to be charge symmetric as described in Section 5.4.1. The method presented below improves on the previous one as the entire shape of the P_{fake} distribution is now taken into account to obtain the value of the proxy ratio depending on the charge, rendering the proxy ratio more reliable.

The P_{fake} requirement from the selection is removed for the p - p samples to increase the accessible range. This is not possible for the p -Pb samples because the requirement at 0.4 is centrally applied by the reconstruction software. In each kinematic interval and for each of the two charges, the P_{fake} distribution in data is divided into intervals with a width of 0.1. A template fit with two components, one for fake tracks and one for real tracks, is performed to each distribution. The template fits employ a recently devised technique, which is based on Reference [136], supports weighted templates and is fast and stable [137]. The templates are taken from simulation, meaning that they are weighted due to the occupancy weighting. Example fits in a kinematic interval with large background from fake tracks are shown in Figure 6.1. Each fit returns the number of fake tracks in data and its statistical uncertainty. For the p - p data samples, this number must be limited to the $P_{\text{fake}} < 0.3$ range, accepted in the selection. The χ^2/n_{DOF} values of the fits are often much larger than unity. Since the number of tracks in each kinematic interval is large, comparably small deviations already substantially contribute to the χ^2 value. These deviations occur as the shapes of the two components are not perfectly modelled in simulation, which is reflected by a systematic uncertainty.

To determine the proxy ratio, the fit result is divided by the known number of fake tracks in simulation. The proxy ratio is shown in Figure 6.2. It is greater than unity in all kinematic intervals, signifying that the number of fake tracks is too small in simulation. In the lowest- η interval, values above four are obtained, which

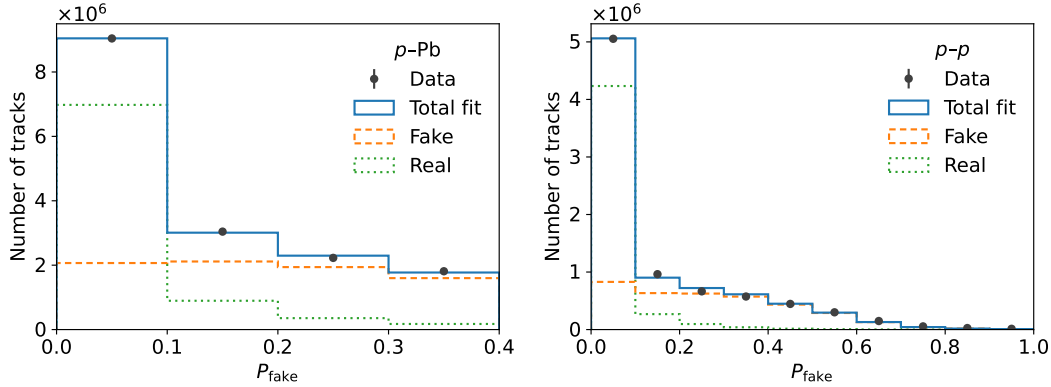


Figure 6.1: Template fits to the P_{fake} distributions in p -Pb data (left) and p - p data (right) in one kinematic interval for one charge.

constitute a large adjustment of the number of fake tracks in simulation. However, this number is small in that kinematic region, reducing the impact on the results. A systematic uncertainty on the proxy ratio is inferred from the deviations of the fits from the data. If the condition $\chi^2/n_{\text{DOF}} > 1$ holds, the statistical uncertainty, u_{stat} , is multiplied by the square root of this number. The increase in uncertainty is taken as the systematic uncertainty, which is thus given by

$$u_{\text{sys}} = u_{\text{stat}} \sqrt{\chi^2/n_{\text{DOF}} - 1}. \quad (6.5)$$

The values of the proxy ratio for the two charges are consistent in many kinematic intervals, but there are also regions in which they differ significantly. Computing the proxy ratio separately for each of the two charges is observed to improve the agreement of the charge ratio with unity at high η and low p_{T} , which is discussed in Section 6.9.2. By construction, the statistical and systematic uncertainties are uncorrelated between the two charges.

The residual fraction of clone tracks is investigated like in Section 5.4.1 and found to be negligible for the p -Pb samples as well.

6.4.2 Photon conversions

The selection of candidate photon conversions into e^+e^- pairs attempted based on topological and kinematic requirements in the first measurement was not suited to define a corresponding proxy as described in Section 5.4.2. Therefore, the background from photon conversions was adjusted with the proxy ratio designed for material interactions of charged pions. In this analysis, a separate proxy for photon conversions is constructed, which does not rely on track pairs, but on single tracks. This is possible because PID variables are now considered. With the separate proxy,




Figure 6.2: Proxy ratio for fake tracks in p -Pb collisions in intervals of η and p_T . The values for the two charges are shown. The bars and boxes indicate the statistical and systematic uncertainties, respectively. The right ends of the light-grey rectangles at low p_T correspond to a momentum of $2 \text{ GeV}/c$, the approximate minimum momentum needed for the long-track reconstruction. In the following, these limits of the acceptance are shown in all plots displaying η and p_T dependence.

the uncertainty on the results is reduced at high η and low p_T . Electrons are mostly identified by the calorimeter system via a likelihood computed from the calorimeter response for the various particle species. This PID information is available in a combined likelihood, which is the product of the likelihoods from the RICH detectors, the calorimeter system and the muon stations, with the former explained in Section 4.2.2. The difference of the logarithms of the combined likelihoods for the electron- and charged-pion-mass hypotheses represents the $\Delta_{e-\pi}$ variable. The electron contribution is expected to peak at positive $\Delta_{e-\pi}$ values, while the tracks from all other long-lived charged particles concentrate at negative values.

The $\Delta_{e-\pi}$ distribution enables the proxy for photon conversions to be determined similarly to the proxy for fake tracks. In each kinematic interval and for each of the two charges, a template fit [137] with two components is performed to the $\Delta_{e-\pi}$ distribution of the candidate tracks in data. The templates are taken from simulation. The conversion component consists of the candidate tracks that are associated with nonprompt electrons and have a photon as ancestor particle. The second component contains all other candidate tracks. Each fit returns the number of tracks from

photon conversions in data and its statistical uncertainty. The χ^2/n_{DOF} values of the fits tend to be large, which is taken into account by a systematic uncertainty on the proxy ratio as detailed in Section 6.4.1.

The background from photon conversions decreases towards high p_T . If the conversion component in the $\Delta_{e-\pi}$ distribution almost vanishes, the template fit cannot accurately determine its size. Hence, the fit result is only used in kinematic intervals with a fraction of tracks from photon conversions in data of at least 1 %, referred to as accepted intervals. The proxy ratio is obtained by dividing the fit result by the known number of tracks from photon conversions in simulation. The proxy ratio in the nonaccepted intervals is estimated as the median value of all values in the accepted intervals. The corresponding systematic uncertainty is deduced from a uniform distribution between the minimum and maximum values in the accepted intervals. The resulting proxy ratio is shown in Figure 6.3. Significant differences are observed between the values for the two charges at high η and low p_T , where the background from photon conversions is largest. The uncertainty on the median value, which is by chance located at unity, is considerable, but its impact on the uncertainty on the results is small since the background is small in the nonaccepted intervals. The statistical and systematic uncertainties in the accepted intervals are uncorrelated between the two charges. The systematic uncertainty in the nonaccepted intervals is fully correlated between the two charges.

6.4.3 Material interactions of charged pions

The entire method for the selection of tracks originating from material interactions of charged pions, which has been developed in the first measurement as presented in Section 5.4.2, is now applied to the p -Pb samples. The performance found closely resembles that for the p - p samples. The values obtained for the proxy ratio demonstrate that the background from material interactions of charged pions is even more overestimated in simulation for the p -Pb samples than for the p - p samples.

6.4.4 Strange-hadron decays

In the first measurement, tracks produced in $K_S^0 \rightarrow \pi^+\pi^-$, $\Lambda \rightarrow p\pi^-$ and $\bar{\Lambda} \rightarrow \bar{p}\pi^+$ decays were adjusted with a combined proxy ratio. For this purpose, the candidate decays selected were sorted into the kinematic intervals according to the kinematics of the decaying particles, and the invariant-mass fits were performed in those kinematic intervals. A combined spline fit to the ratios of the yields in data and simulation was then executed. Lastly, the fitted model was mapped to the kinematics of the decay products based on simulation. Now, the sorting into the kinematic intervals is modified to make the mapping superfluous and permit a smaller systematic uncertainty to be assigned.




Figure 6.3: Proxy ratio for photon conversions in p -Pb collisions in intervals of η and p_T . The bars and boxes on the markers for the two charges indicate the statistical and systematic uncertainties, respectively. The bars on the markers for the median value represent its systematic uncertainty.

Candidate decays are selected similarly to the first measurement with topological requirements chosen to retain high efficiency. The K_S^0 , Λ and $\bar{\Lambda}$ candidates are sorted into the kinematic intervals according to the kinematics of their decay products. Consequently, six sets of candidates are obtained: $K_S^0 \rightarrow \pi^+ X$, $K_S^0 \rightarrow X \pi^-$, $\Lambda \rightarrow p X$, $\Lambda \rightarrow X \pi^-$, $\bar{\Lambda} \rightarrow \bar{p} X$ and $\bar{\Lambda} \rightarrow X \pi^+$, where X denotes the decay product of which the kinematics is not considered. For each set of candidates, the invariant mass of the decaying particle is fitted in the kinematic intervals with sufficiently many entries in data and simulation. The fit strategy described in Section 5.4.3 is used. From the fit results, the ratios of the yields in data and simulation are computed. The p_T dependence observed in the yield ratios is the same irrespective of the charge of the corresponding decay product. This is expected if the deviations of the yield ratios from unity are caused by a general discrepancy between data and simulation in the number of strange quark–antiquark pairs produced, which has been concluded in the first measurement. Therefore, the six sets of candidates can be reduced to three categories: π -from- K_S^0 , p -from- Λ and π -from- Λ .

The three categories of yield ratios are shown in Figure 6.4. They do not cover all kinematic intervals, but their p_T dependence can be modelled. For each category, a

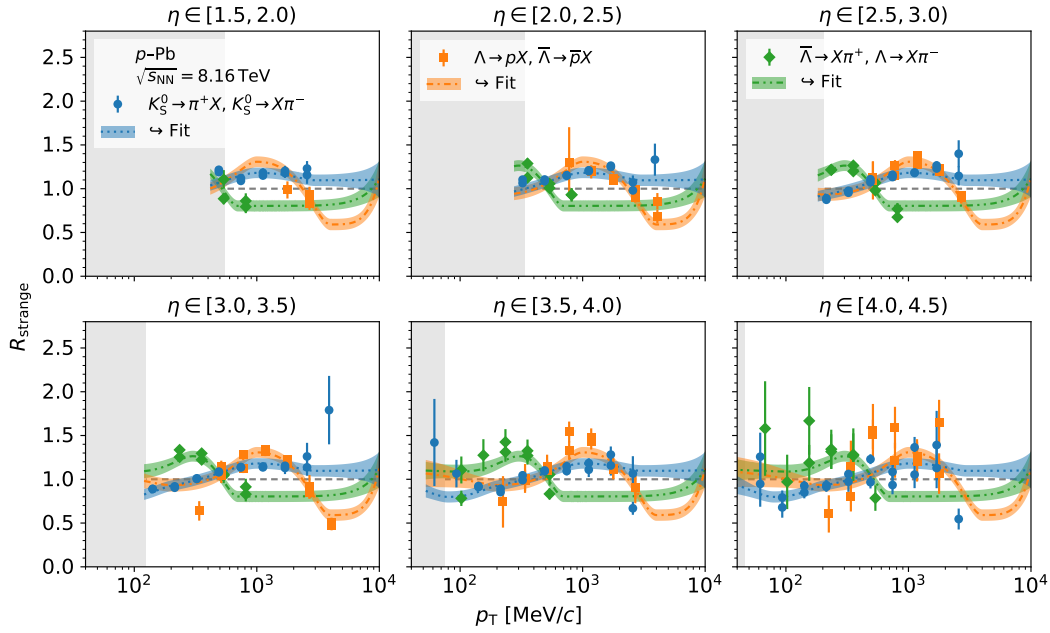


Figure 6.4: Proxy ratio for strange-hadron decays in p -Pb collisions in intervals of η and p_T . The blue dotted, orange dash-dotted and green dash-double-dotted lines indicate the three categories of the proxy ratio, with the surrounding bands representing the sum in quadrature of the statistical and systematic uncertainties. The lines result from the three spline fits to the yield ratios obtained for the sets of candidates listed above the corresponding line in the legend. The bars indicate the propagated statistical uncertainties from the invariant-mass fits in data and simulation.

spline fit is performed to the values of the yield ratio in all η intervals. Two weak Gaussian constraints are added that guide the spline to the median value of the corresponding yield ratio at the lower and upper ends of the p_T grid. The three fitted models determine the proxy ratio and its statistical uncertainty and are also shown in the figure. The p_T dependence of the p -from- Λ and π -from- Λ categories of the proxy ratio differs significantly, which can be attributed to the unequal (anti)proton and charged-pion masses. In the first measurement, these structures were diluted by considering the kinematics of the decaying particles. A comparison of the yield ratios with the models reveals that there are deviations not covered by the statistical uncertainty. Hence, this uncertainty is increased until at least 68% of the values are contained. The increase in uncertainty is taken as a systematic uncertainty on the proxy ratio. The value of this uncertainty ranges between 0.06 and 0.10, depending on the data sample. The statistical uncertainty is comparably small. The statistical and systematic uncertainties are fully correlated between the two charges.

The background from (anti)protons produced in $\Sigma^+ \rightarrow p\pi^0$ and $\bar{\Sigma}^- \rightarrow \bar{p}\pi^0$ decays is also adjusted with the p -from- Λ category of the proxy ratio because those decays are similar to the $\Lambda \rightarrow p\pi^-$ and $\bar{\Lambda} \rightarrow \bar{p}\pi^+$ decays, respectively, both in kinematics and topology. The masses and lifetimes of the decaying particles resemble each other, and the neutral-pion mass only slightly differs from the charged-pion mass.

6.5 Signal composition

The contributions of charged pions, charged kaons and (anti)protons to the hadron composition are disentangled with template fits to two-dimensional distributions of PID variables. This approach is inspired by two LHCb measurements of antiproton production in fixed-target p -He collisions at $\sqrt{s_{\text{NN}}} = 110$ GeV [138, 139]. It exploits the entire ranges of PID distributions of the candidate tracks; no requirements are imposed to isolate single hadron species. Thus, no efficiency needs to be estimated, which is in accord with the general approach of the analysis of applying as few selection requirements as possible. The main work consists in deriving templates from data. Taking them from simulation would bias the analysis as the PID variables are not reliably modelled. This is because, *e.g.*, the occupancy of the RICH detectors as well as experimental conditions, such as temperature and pressure of the gas media and variations thereof, are not accurately described [140]. The template creation on data requires calibration samples to be selected, which are not completely representative of the candidate tracks, however. In the recent LHCb measurement of antiproton production [139], a neural-network method was developed [141] to map the templates to the candidate tracks. That measurement is the only paper so far in which this method was used. In this analysis, the method is applied to data from p -Pb and p - p collisions for the first time. It is found that modifications are needed to achieve competitive performance, caused by the higher centre-of-mass energies and the larger occupancy.

Charged hadrons are primarily identified by the RICH detectors. Therefore, PID variables are employed that are based on the RICH likelihoods for the three hadron species as explained in Section 4.2.2. The $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ variables are respectively defined as the differences of the logarithms of those likelihoods for the charged-kaon- and (anti)proton-mass hypotheses relative to that for the charged-pion-mass hypothesis. The two-dimensional distributions in the $\Delta_{K-\pi}$ - $\Delta_{p-\pi}$ plane exhibit three distinct areas for the three hadron species. This is exemplified in Figure 6.5. At negative $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ values, the charged-pion contribution prevails. Charged kaons and (anti)protons reside at positive $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ values, respectively. Differences are evident in the shapes of the areas between data and simulation. The three distinct areas only occur in those kinematic intervals in which the Cherenkov momentum threshold is reached for all three hadron species. This further restricts the




Figure 6.5: Two-dimensional distributions of the $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ variables of negatively charged candidate tracks in p -Pb data (top left) and simulation (top right and bottom) in one kinematic interval. The three plots from simulation are shown for comparison and only display the candidate tracks associated with prompt negatively charged pions (top right), prompt negatively charged kaons (bottom left) and prompt antiprotons (bottom right).

minimum accessible p_T values compared to the limit from the track-reconstruction correction. With growing momentum, the Cherenkov angles of the three hadron species approach each other as the particle mass becomes negligible. This degrades the separation power and restricts the maximum accessible p_T values. Consequently, measuring the signal composition is only possible at intermediate p_T values.

6.5.1 Calibration samples

In simulation, the individual $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ distributions of charged pions, charged kaons and (anti)protons can easily be obtained by isolating candidate tracks associated with the corresponding particle at generator level. In data, these distributions

cannot be selected based on single tracks. Instead, decays into the hadron species of interest have to be identified, which is achievable via the invariant masses of the decaying particles. It needs to be ensured that no requirements on PID variables are applied to the decay products to be used for calibration. Then, the sets of calibration tracks represent the individual $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ distributions of the desired hadron species.

The calibration tracks for charged pions, charged kaons, protons and antiprotons are obtained from $K_S^0 \rightarrow \pi^+\pi^-$, $\phi \rightarrow K^+K^-$, $\Lambda \rightarrow p\pi^-$ and $\bar{\Lambda} \rightarrow \bar{p}\pi^+$ decays, respectively, where ϕ is the compact notation for the $\phi(1020)$ meson. In each event with exactly one primary vertex (also for the $p-p$ samples), all unique pairs of oppositely charged candidate tracks are formed to select candidate decays. The invariant mass of a track pair must lie in a window around the known K_S^0 , ϕ , Λ or $\bar{\Lambda}$ mass [17]. This procedure is similar to that in Section 5.4.3. The requirement of exactly one primary vertex per event is necessitated by six geometric variables employed to increase the purity of the selection. These variables belong to an irreducible set of seven variables that contain the entire information on a two-body decay. The derivation of the variables is presented in Appendix A. Three of the six geometric variables are ideally zero: the distance, d_{CA} , of closest approach of the decay products; the distance, d_{DP} , of the primary vertex from the decay plane and the in-decay-plane impact parameter, I , of the decaying particle at the primary vertex. Finite lifetime of the decaying particle is probed by the other three variables: the squared in-decay-plane impact parameters, I_{11} and I_{22} , of the decay products at the primary vertex and the remaining contribution, I_{12} , to the flight distance of the decaying particle that is not covered by I_{11} and I_{22} . The variables are perfectly suited to be used in a machine-learning classifier, which profits from a small number of independent input variables.

In this subsection, the K_S^0 , ϕ , Λ and $\bar{\Lambda}$ decays and the calibration tracks obtained from them are referred to as signal, whereas all other tracks from random track combinations are termed background. Most separation power for K_S^0 , Λ and $\bar{\Lambda}$ candidates is present in the I_{11} , I_{22} and I_{12} variables due to the long lifetimes of those particles, which cause the decay vertices to be detached from the primary vertex. A solely geometric selection of the ϕ decay is more challenging since the ϕ meson decays promptly and thus does not possess a flight signature. This reduces the separation power in I_{11} , I_{22} and I_{12} . To further reject background in the ϕ selection, the $\Delta_{K-\pi}$ variable of the track not used for calibration is additionally included. For consistency, the $\Delta_{K-\pi}$ variable is employed in the K_S^0 , Λ and $\bar{\Lambda}$ selections as well. The invariant mass, which is the seventh variable from the irreducible set, serves as a discriminating variable between the peaking signal and the residual flat background. Consequently, it is not included in the selections.

The selections are defined by training a gradient-boosted decision tree as implemented in the XGBoost package [142] for each of the three species of calibration

tracks and each of the two charges. The training is executed on simulation. The known signal and background categories of the candidate calibration tracks determine the two classes to be distinguished. The classifiers employ the six geometric variables and the $\Delta_{K-\pi}$ variable as input. Two preprocessing operations are performed; the decimal logarithm is applied to the geometric variables to render their distributions Gaussian-like, and the absolute value of I_{12} is used as this variable is not guaranteed to be nonnegative by construction. The candidate calibration tracks are randomly assigned to training and testing subsamples that shall comprise 75% and 25% of the tracks, respectively. The default XGBoost settings are employed for almost all hyperparameters, but the maximum tree depth is reduced from its default value of six to four, which decreases the risk of overtraining the classifier. Overtraining arises if a classifier closely adapts to characteristics of the training subsample. Furthermore, the weights of the signal class are multiplied by the ratio of the numbers of background and signal entries in the training subsample, which takes into account that the input is dominated by background. The distributions of the classifier response are shown in Figure 6.6. The difference in magnitude between the signal and background distributions at values near unity reflects that the charged-pion signal is better separable from the background than the charged-kaon signal. The testing distributions agree well with the training distributions even on a logarithmic scale, indicating that effects of overtraining are sufficiently small. Each testing subsample is used to determine a requirement on the response of the corresponding classifier, which defines the K_S^0 , ϕ , Λ and $\bar{\Lambda}$ selections. The response requirement is set by demanding a signal efficiency of 50%. Afterwards, the classifiers are applied to the candidate calibration tracks in data, and the response requirements are imposed.

The selected candidate calibration tracks are sorted into kinematic intervals. In each kinematic interval and for each of the six sets of candidate calibration tracks, the invariant-mass distribution of the decaying particle is fitted analogously to Section 5.4.3. Example fits to ϕ invariant-mass distributions are shown in Figure 6.7. The signal in simulation is narrower than that in data because the natural width of the ϕ meson is not simulated; thus, the peak width in simulation is only due to the momentum resolution of the detector. The background in the ϕ invariant-mass distribution is still substantial. By contrast, the background in the K_S^0 , Λ and $\bar{\Lambda}$ invariant-mass distributions is much smaller than in Section 5.4.3 since these selections are tighter.

The fitted signal and background components of the invariant mass, the discriminating variable, can be used to disentangle the signal and background components of another variable, the control variable, on a statistical basis. If the discriminating and control variables are statistically independent both for signal and background, a weight for each entry can be computed from the discriminating variable and




Figure 6.6: Distributions of the classifier response for the $K_S^0 \rightarrow \pi^+ \pi^-$ selection (left) and the $\phi \rightarrow K^+ K^-$ selection (right) in p -Pb simulation. The signal and combinatorial-background distributions obtained from the training and testing subsamples are shown for one charge.

applied to the control variable. The weighted distributions of the control variable then represent its signal and background components. This formalism was first described in Reference [143], became popular based on Reference [144] and has recently been revisited and extended in Reference [145]. The weights are usually called *sWeights*.

The statistical independence required is a stronger condition than lack of correlation, which only takes into account linear relations. Independence can be probed with a test using U -statistic permutation [146]. This test has high statistical power and is optimal in its construction. In the analysis, two control variables exist, which are the $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ variables. Consequently, the test needs to be applied four times in each kinematic interval for a given set of candidate calibration tracks. As the signal and background categories have to be known before the potential *sWeight* computation, the test must be performed on simulation. It returns a probability value. If this value is very small, evidence of dependence is found. In almost all kinematic intervals, the probability values are observed to be reasonably large. Single small values of approximately 10^{-3} occur, but only in kinematic intervals below the minimum accessible p_T values. Independence is thus demonstrated, and the *sWeights* for the signal component are determined according to Reference [145]. They peak at the known mass of the decaying particle and have negative values in the sidebands.

Applying the *sWeights* to the candidate calibration tracks enables the background-subtracted $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ distributions to be obtained of charged pions, charged kaons and (anti)protons in data. These weighted tracks form the calibration samples; in this section, the default data and simulation samples are termed initial samples




Figure 6.7: Fits to the invariant-mass distributions of $\phi \rightarrow K^+K^-$ candidates in p -Pb data (left) and simulation (right) in one kinematic interval for one charge.

to discriminate them from the calibration samples. The *sWeighted* $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ distributions are needed to construct the two-dimensional templates for the fits measuring the hadron composition in the initial data samples. Directly employing the *sWeighted* distributions as templates biases the fit results because of the properties of the calibration tracks and the corresponding decays. In particular, the charged-pion and (anti)proton calibration tracks are nonprompt, while most candidate tracks in the initial samples are prompt. The detector occupancy of the events that contribute does not agree either between the calibration and initial samples. In addition, the track slopes are different, which are important for the determination of the centres of the photon rings in the RICH detectors. Since such distributions affect the PID response, *i.e.* the two-dimensional distributions of the $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ variables, the properties of the candidate tracks must be taken into account as well. The connection between the calibration and initial samples is established by machine learning, namely by neural networks that are trained to discern the dependence of the PID response on the relevant variables.

The calibration-sample selection presented above is also executed purely on simulation. The calibration simulation samples obtained allow a cross-check of the entire PID-calibration method, which consists in comparing the hadron composition resulting from this method with that present in the initial simulation samples. This serves as a stringent closure test, and the differences are an estimate of the systematic uncertainty on the hadron composition in data, induced by all steps of the method.

6.5.2 Transfer of particle-identification response

The idea and implementation of the neural-network method designed to model the dependence of the PID response on track and event properties are taken from

References [139, 141]. The PID response is described by a Gaussian mixture model, which is a linear combination of two-dimensional Gaussian distributions in this case. Each component has six parameters: two means, two widths, a correlation angle and an amplitude. The correlation angle permits a rotation in the $\Delta_{K-\pi}-\Delta_{p-\pi}$ plane. The values of these parameters are predicted by a set of multilayer-perceptron networks as implemented in the TENSORFLOW package [147]. The network input comprises twelve variables that either affect the PID response significantly or reflect differences between prompt and nonprompt tracks. The variables are geometric, kinematic, indicate the track-reconstruction quality and measure the detector occupancy: the x , y and z coordinates of the track origin as well as the track $x-z$ and $y-z$ slopes; the track momentum and η values; the χ^2/n_{DOF} and n_{DOF} values of the track fit as well as the number of hits in each of the two RICH detectors and the total number of tracks per event.

Before the neural-network training, the momentum ranges covered by the calibration tracks are restricted to discard kinematic regions without separation power in the Cherenkov angle. Otherwise, the neural-network models would be distorted, degrading their performance in other kinematic regions as well. A minimum-momentum requirement of $9.5 \text{ GeV}/c$ is imposed because it is slightly above the momentum threshold for Cherenkov radiation of charged kaons, which is $9.3 \text{ GeV}/c$. Moreover, the momentum must not exceed $120 \text{ GeV}/c$. One model is trained for each of the six calibration data samples (*i.e.* separately for each of the three hadron species and each of the two charges). The *sWeights* are included in the training, ensuring that the networks only process the individual distributions of charged pions, charged kaons or (anti)protons although they are solely trained on data. Each trained model is validated by evaluating the model prediction in the kinematic intervals retained and comparing the distributions obtained with the corresponding $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ distributions from the calibration sample. As the networks shall learn the shapes of these distributions, the (dis)agreement observed is the main metric for the model performance.

Several modifications to the original method have been necessary as its performance was found to be suboptimal, which is detailed in the following. Six hyperparameters are present in the original method. They define the model complexity, the network architecture and the training process. Their settings are listed in Table 6.1. Firstly, the same settings were used to train models on the $p\text{-Pb}$ and $p\text{-p}$ calibration samples. However, the model performance was not comparable to that observed in the antiproton-production analysis [139, 141]. In particular, the loss functions of the models did not evolve as expected. The value of a loss function ideally decreases fast during the first training epochs, then less rapidly and moderately approaches a minimum. If a training sample is too large to be jointly processed in one epoch, it can be divided into batches. In this case, the loss function oscillates around the minimum value due to sample fluctuations. By contrast, the loss functions decreased

Table 6.1: Settings of the hyperparameters, H , in the original neural-network method, developed in the analysis of antiproton production in p -He collisions [139, 141], and in the modified method, established in this analysis, for the three hadron species. The variables listed are the number, n_c , of Gaussian components; the number, n_n , of layer nodes; the number, n_l , of layers; the number, n_e , of epochs; the batch size, b ; the initial learning rate, ℓ_i ; the learning-rate reduction factor, a ; the number, n_p , of patience epochs and the minimum learning rate, ℓ_{\min} . The numbers of nodes and layers refer to all network layers except for the output layer. *The number of epochs is not a hyperparameter in this case since it is determined by the early-stopping criterion of reaching the minimum learning rate. **The learning-rate reduction factor is set to 0.9 for the training on simulation. ***The minimum learning rate is not a hyperparameter in this case as it results from the reduction of the initial learning rate by the fixed value of five orders of magnitude.

H	Setting					
	Antiproton-production analysis			This analysis		
	π^-	K^-	\bar{p}	π^+, π^-	K^+, K^-	p, \bar{p}
n_c	64	64	20	32	28	16
n_n	128	128	128	128	96	64
n_l	3	3	3	6	4	4
n_e	380 000	50 000	130 000	$\approx 16\,000^*$	$\approx 16\,000^*$	$\approx 16\,000^*$
b	10 000	20 000	10 000	30 000	30 000	30 000
ℓ_i	10^{-3}	5×10^{-6}	10^{-4}	10^{-3}	10^{-3}	10^{-3}
a	—	—	—	0.99**	0.99**	0.99**
n_p	—	—	—	10	10	10
ℓ_{\min}	10^{-8***}	$5 \times 10^{-11***}$	10^{-9***}	10^{-10}	10^{-10}	10^{-10}

slowly across many epochs and suddenly much more steeply during the last epochs. Increasing the number of epochs did not change this behaviour, which was attributed to the implemented reduction of the learning rate; in the original method, the initial learning rate is reduced by five orders of magnitude across all epochs. Apart from the loss functions, the model predictions exhibited clear potential for improvement because their projections were often shifted compared to the $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ distributions to be learned. Furthermore, sharply rising structures were found in the projections in various kinematic intervals, induced by single Gaussian components of the models with implausibly large amplitudes. Another aspect is the training time. Like in the original method, the training is performed on a GPU, enabling a significant increase in speed compared to CPU-based training. An NVIDIA V100 GPU is used. The original hyperparameter settings lead to a training time of approximately 25 h for charged-pion calibration tracks. This duration is mainly

driven by the large number of epochs, which amounts to 380 000. While such a duration is still acceptable if a model needs to be trained only a few times, it is undesirable if extensive studies on hyperparameters are required. The number of models to be constructed in the analysis is 48, resulting from the six calibration samples for each initial sample and the eight initial samples. Thus, the training time needed to be further reduced.

Modifying the original method to achieve reasonable performance on the p -Pb and p - p samples required a significant amount of work. More than 180 different models were trained and validated to find appropriate hyperparameters and settings. The modified method, which was lastly established, comprises eight hyperparameters. Their settings are also listed in Table 6.1. The model performance was discovered to strongly depend on the initial learning rate and the learning-rate reduction across the epochs. An algorithm that dynamically adjusts the learning rate during the training proved to be suited. It is available in TENSORFLOW. If the loss function does not decrease across a given number of epochs, referred to as patience epochs, the learning rate is multiplied by a factor of less than unity. This reduction should be ceased as soon as the learning rate reaches values in a range where numerical effects become relevant. A minimum learning rate is therefore defined, which constitutes an early-stopping criterion. Consequently, the number of epochs is not a hyperparameter, but follows from this criterion. It approximately amounts to 16 000. This means that the modified method involves considerably fewer epochs than the original one, resulting in a much faster training of only 1.5 h. In addition, the loss functions have the expected plateauing behaviour. An initial learning rate of 10^{-3} is found to work successfully for all calibration tracks, demonstrating that this hyperparameter is more robust than in the original method, in which different initial learning rates were needed. For charged-pion and charged-kaon calibration tracks, the number of Gaussian components is greatly reduced compared to the original method, mitigating the risk of obtaining sharply rising artefacts in the predictions. The number of Gaussian components is chosen to be largest for charged pions, in which case the PID response has the most complex shape; the $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ distributions possess a two-peak structure in some kinematic intervals. For charged kaons and (anti)protons, the distributions always exhibit a single peak, but they are often asymmetric in the former case, which necessitates a larger number of Gaussian components. Projections of the model prediction for antiprotons are exemplified in Figure 6.8. Good agreement is achieved; the model is generally compatible with the data. The impact of the numbers of nodes and layers as well as of the batch size on the model predictions is smaller than that of the hyperparameters discussed above. Nevertheless, these hyperparameters were also manually optimised to further improve the model performance. The numbers of nodes were adapted to the numbers of Gaussian components. Four layers perform well for charged kaons and (anti)protons, whereas six layers are better for charged pions. The batch size is




Figure 6.8: Projections of the model prediction for \bar{p} calibration tracks in p -Pb data onto the $\Delta_{K-\pi}$ axis (left) and the $\Delta_{p-\pi}$ axis (right) in one kinematic interval.

increased to 30 000 in all cases. A larger batch size reduces the oscillations of the loss function around the minimum value. The performance is additionally ameliorated by extending the original set of input variables by p_T , which directly defines the kinematic intervals in the analysis together with η .

Overtraining the networks is unlikely since no classification occurs. Instead, the networks shall learn the dependence of the two-dimensional PID response on a set of input variables based on training samples, which are the calibration samples. The models are then applied to other samples, the initial samples. There is only mild correlation because the initial samples represent a large superset of the calibration samples. As a cross-check, the evolutions of the parameters of the Gaussian components with the input variables are assessed [139, 141]. Smooth evolutions without fluctuations are found, indicating that the models are not overtrained.

The model performance on the initial samples can only be evaluated on simulation by comparing the predicted distributions with the known PID response for the particle under consideration. If the model is transferable, it is able to predict the known PID response without significant differences. An example of this comparison is shown in Figure 6.9. When the complexity of the processed modelled is considered, the agreement is acceptable. There are residual discrepancies, which might affect the determination of the hadron composition, but this is taken into account by the closure test of the PID-calibration method.

6.5.3 Hadron fractions

The predictions of the neural-network models on the initial data samples are used as templates in the fits to the two-dimensional distributions of the $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ variables of the candidate tracks in data. These templates represent the three signal

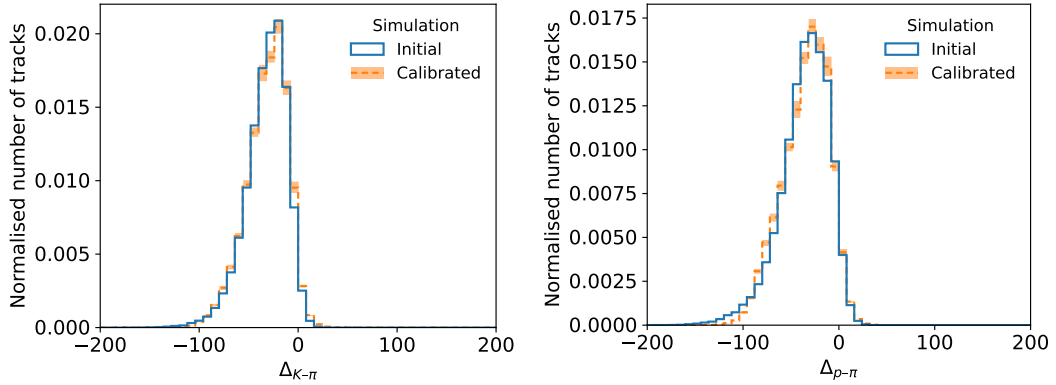


Figure 6.9: Distributions of the $\Delta_{K-\pi}$ variable (left) and the $\Delta_{p-\pi}$ variable (right) for negatively charged pions initially present in p -Pb simulation and obtained from the PID-calibration method executed on p -Pb simulation in one kinematic interval.

components in the fits: charged-pion, charged-kaon and (anti)proton candidate tracks. There are two further components, which are background: fake tracks and tracks from long-lived charged particles other than the signal components. The shapes of the background components are taken from simulation.

The template fit is performed separately in each kinematic interval and for each of the two charges. The numbers of charged-pion, charged-kaon and (anti)proton candidate tracks are the only parameters that are allowed to vary. The fractions of fake tracks and of tracks from other long-lived charged particles are fixed to their values in simulation. The former is multiplied by the proxy ratio for fake tracks, which quantifies the difference in the contribution of this background between data and simulation. The template fits use a χ^2 cost function for Poisson-distributed counts [148], taking into account the statistical uncertainty on the distribution fitted, but not that on the templates. Including this uncertainty destabilises the fits and tends to give fitted models describing the peaks of the distributions worse, but these regions are most important for the hadron fractions to be measured. The residual mismodelling of the templates, which is a larger effect than their statistical uncertainties, is covered by the systematic uncertainty to be assigned. The fits are exemplified in Figure 6.10. Each fit returns the numbers of charged-pion, charged-kaon and (anti)proton candidate tracks, their statistical uncertainties and their statistical correlations. These correlations are propagated into the results. The χ^2/n_{DOF} values of the fits are large because the signal templates are not perfect. However, they are able to model the distributions much better than signal templates obtained directly from the initial simulation samples, *i.e.* without the PID-calibration

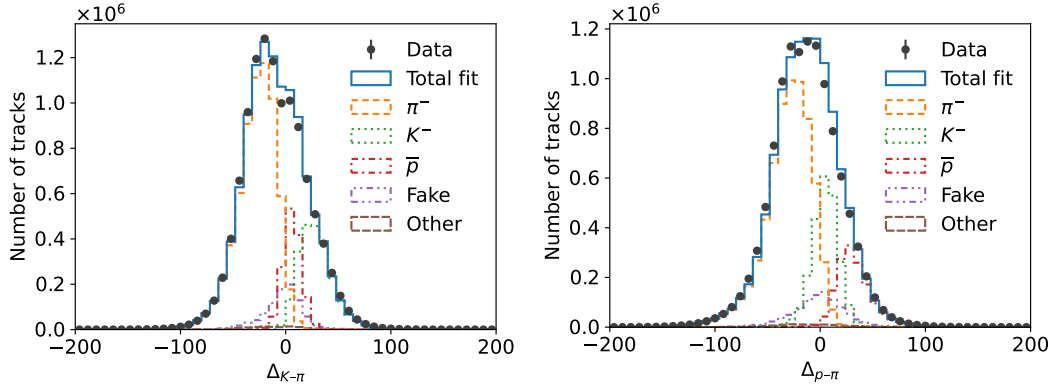


Figure 6.10: Projections of the template fit to the two-dimensional distribution of the $\Delta_{K-\pi}$ and $\Delta_{p-\pi}$ variables of negatively charged candidate tracks in p -Pb data onto the $\Delta_{K-\pi}$ axis (left) and the $\Delta_{p-\pi}$ axis (right) in one kinematic interval. The *other* component denotes tracks from long-lived charged particles other than charged pions, charged kaons and (anti)protons.

method; those templates induce much larger χ^2/n_{DOF} values. This demonstrates that the method is effective.

The fractions of the three hadron species are obtained by dividing the fit results by the sum of all components. They are shown in Figure 6.11. The statistical uncertainties are below 0.2% in most kinematic intervals. All steps of the PID-calibration method, including the template fits, are performed on simulation as well to enable the closure test introduced in Section 6.5.1. The absolute values are then determined of the differences of the fractions resulting from the fits and those present in the corresponding initial simulation sample, computed by counting the tracks associated with signal particles. These values amount to 4% on average, relative to the fractions in data, and are below 15% in most kinematic intervals. This size is reasonable in consideration of the difficulty of deriving appropriate templates. The closure test can thus be regarded as successful, and the absolute differences are assigned as the systematic uncertainties on the fractions in data. The statistical and systematic uncertainties are uncorrelated between the two charges. Moreover, the statistical uncertainty is partially anticorrelated between the three hadron species, while the systematic uncertainty is uncorrelated between them.

6.6 Efficiency

Simulation indicates that the efficiencies for prompt charged pions, prompt charged kaons, prompt (anti)protons and other prompt long-lived charged particles in p -Pb collisions closely resemble those in p - p collisions. The same dependence on η_{lab}




Figure 6.11: Fractions of charged-pion, charged-kaon and (anti)proton candidate tracks in p -Pb data in intervals of η and p_T . The fractions are multiplied by the charge to visually separate the corresponding values. The bars and boxes indicate the statistical and systematic uncertainties, respectively. The large gaps between the right ends of the light-grey rectangles at low p_T , representing the minimum momentum of 2 GeV/c needed for the long-track reconstruction, and the left markers are caused by the Cherenkov momentum thresholds in the RICH detectors. In the following, these gaps occur in all plots in which the hadron fractions enter.

and p_T is observed, including the dips caused by the material budget. Like in the first measurement performed within this thesis, the efficiency in simulation is corrected based on data to take into account mismodelling of the track reconstruction and the particle composition. The systematic uncertainty from interval-to-interval migration, investigated as described in Section 5.5.4, is found to be negligible for the p -Pb samples as well.

6.6.1 Track reconstruction

The track-reconstruction corrections for the p -Pb and p -p samples are obtained by mapping correction factors from a control study [129] into the kinematic intervals in the analysis as presented in Section 5.5.1. The only difference lies in the determination of the correction factors for the p -Pb samples. At low momentum, the invariant-mass distributions of $J/\psi \rightarrow \mu^+ \mu^-$ candidates suffer from a high level

of combinatorial background from randomly joined muon candidates. This would lead to an undesirably large statistical uncertainty in the kinematic intervals at low p_T , especially for the backward beam configuration. Hence, in the control study, the J/ψ candidates were only used to compute the correction factor at high momentum, whereas the low-momentum region was covered by $K_S^0 \rightarrow \pi^+\pi^-$ candidates. Since charged pions cannot be reconstructed as muon-station–TT tracks, the efficiency of the long-track reconstruction was obtained by separately probing the reconstructions in the VELO and the IT–OT stations. The resulting relative statistical uncertainties on the correction factors do not exceed 1.8% for tracks with momenta below 40 GeV/ c .

The corrections to the individual charged-pion, charged-kaon and (anti)proton efficiencies in simulation are given by the track-reconstruction correction. In the case of the combined efficiency for long-lived charged particles, this correction needs to be multiplied by the particle-composition correction, which is described in the following.

6.6.2 Particle composition

In the first measurement, the particle-composition correction was based on an LHCb measurement of ratios of prompt hadron production in p – p collisions at lower energies as explained in Section 5.5.2. These hadron ratios were extrapolated to $\sqrt{s} = 13$ TeV, and the corresponding hadron ratios in simulation were corrected to reproduce the extrapolated values. The corrections were parameterised as functions linear in $\log_{10} p_T$ to describe the dependence across the entire p_T ranges. Now, the main contributions to the particle composition—prompt charged pions, prompt charged kaons and prompt (anti)protons—are measured in this analysis. Below, the other prompt long-lived charged particles are termed other particles. Consequently, no extrapolation in energy needs to be performed. However, an extrapolation in p_T is still required as the hadron composition is not measurable across the same p_T ranges as the long-lived charged particles. This is due to the Cherenkov momentum thresholds in the RICH detectors. The p_T dependence is obtained from hadronic-interaction models. One advantage is that the other-particle fraction is now constrained by the sum of the charged-pion, charged-kaon and (anti)proton fractions, rendering the correction more reliable; no constraint was available from the extrapolation in energy in the first measurement. Furthermore, half the absolute value of the deviation of the correction from unity was assigned as systematic uncertainty. This is a valid estimate if several assumptions are made in the determination of a correction. The drawback is that the uncertainty vanishes if the correction is unity, which is now avoided.

The hadronic-interaction models studied are PYTHIA 8.308 [38] with the default Monash 2013 tune [40], EPOS-LHC [26], QGSJET II.04 [33] and SIBYLL 2.3d [35].

For each of them, 10^8 inelastic events are generated [32] to compute the predicted charged-pion, charged-kaon, (anti)proton and other-particle fractions. The measured fractions of charged-pion, charged-kaon and (anti)proton candidate tracks, which have been presented in Section 6.5.3, still include background from non-prompt tracks and do not incorporate the efficiency. The corresponding particle fractions, needed for a comparison with the model predictions, are obtained according to Equation (6.1). The choice of a reasonable extrapolation into the kinematic intervals at low and high p_T is driven by several considerations. The charged-pion fraction must not become too large at low p_T to prevent the other-particle fraction from falling below zero because of the constraint on the sum of the fractions to give unity. Simultaneously, the charged-pion fraction should not be too small at high p_T ; the other-particle fraction is implausibly large there otherwise. The (anti)proton fraction needs to be sufficiently large to avoid negative values at high p_T , which can occur due to the predicted decrease with p_T . In addition, the resulting particle-composition correction is expected to be below 10% in most kinematic intervals. This estimate stems from the LHCb measurement of prompt charged-particle production in p -Pb and p - p collisions at $\sqrt{s_{NN}} = 5$ TeV, where this correction reached at most 7% [117]. All these considerations are best fulfilled if SIBYLL is used to describe the p_T dependence.

In each η interval, for each of the two charges and for each of the three hadron species, the extrapolation is defined as product of a factor and the fraction predicted by SIBYLL. The predictions for p - p collisions are also employed for p -Pb collisions owing to the similar p_T dependence found. The factor is determined in a fit to the measured fractions using a cost function of least squares. Weak Gaussian constraints, which exploit the predictions from the other models as well, are needed at low and high p_T to reflect the considerations mentioned above. Each fit returns the factor and its statistical uncertainty. The fitted extrapolations represent the corrected fractions, from which the particle-composition correction is obtained. Various alternative definitions of the extrapolations, based on one of the other models or combinations of all four models, have been investigated, but none of them has improved on the default definition based on SIBYLL. Nevertheless, the extrapolations do not describe the measured fractions perfectly. This is taken into account by assigning half the absolute value of the deviations of the corrected charged-pion, charged-kaon and (anti)proton fractions from the corresponding initial fraction in simulation as the systematic uncertainty on the corrected fractions. This systematic uncertainty is a compromise between reflecting the modelling involved in the correction and not overestimating the uncertainty. The statistical uncertainty is negligible. The corrected other-particle fraction and its systematic uncertainty are computed from the corrected fractions of the three hadron species via the constraint on the sum of the fractions. The corrected fractions are shown in Figure 6.12. The deviations from the initial fractions are moderate in most kinematic intervals. The largest discrepancies




Figure 6.12: Fractions of prompt charged pions, prompt charged kaons, prompt (anti)protons and other prompt long-lived charged particles initially present in simulation and corrected to minimise differences compared to data in p -Pb collisions in intervals of η and p_T . The fractions are multiplied by the charge to visually separate the corresponding values, and they are vertically shifted to avoid overlaps. The bars indicate the systematic uncertainty.

occur around $p_T \approx 2000$ MeV/ c , where the charged-pion fraction is reduced, while the other-particle fraction increases. The charged-kaon and (anti)proton fractions are only slightly modified.

The corrected efficiency for prompt long-lived charged particles is obtained by adding the products of the four corrected fractions and the individual efficiencies in simulation. The systematic uncertainty is propagated into the corrected efficiency. The ratio of this efficiency and the initial efficiency in simulation represents the particle-composition correction. It is shown in Figure 6.13. The initial particle composition in simulation would lead to an efficiency that is overestimated by up to 10 %. This can be attributed to the other-particle fraction, which is initially too small. Since the efficiency for these particles is much lower than the charged-pion efficiency, an increase in their fraction reduces the combined efficiency. The charge asymmetries observed are small. The uncertainty, which is purely systematic, is uncorrelated between the two charges.




Figure 6.13: Particle-composition correction to the efficiency in p -Pb simulation in intervals of η and p_T . The values for the two charges are shown. The bars indicate the systematic uncertainty.

6.6.3 Hadronic interactions

The uncertainty on the material budget in simulation [129] induces an additional systematic uncertainty on the track-reconstruction correction as detailed in Section 5.5.3. The determination of the uncertainty remains unchanged compared to the first measurement. As the material budget in simulation is the same for all samples studied, the uncertainty is fully correlated between the two charges, between the three hadron species, between the p -Pb and p -p samples as well as between the forward and backward p -Pb samples.

6.7 Acceptance correction

The centre-of-mass system does not coincide with the laboratory system for the p -Pb collisions analysed. This effect is taken into account by applying the constant rapidity shift for massless particles to the η_{lab} value of each candidate track as described in Section 6.3. However, the rapidity shift is an approximation because the kinematics and the particle mass are ignored, which can induce migration between the kinematic intervals. To correct for this, a Lorentz boost into the centre-

of-mass system is performed, and the particle pseudorapidity is compared with that obtained from the rapidity shift. The boost can only be executed on simulation, where the species of a single particle is known.

The generated η_{lab} and p_{T} values of all prompt long-lived charged particles in simulation are studied. Since the boost points into the z direction, the transverse momentum is not affected by it, but the pseudorapidity in the laboratory system changes due to its dependence on the momentum z component. The acceptance correction is determined from two counts. First, the particles are counted and sorted into kinematic intervals according to the η_{lab} value shifted by Δ_y (the default η value in the analysis) and p_{T} . Second, the boosted η_{lab} value, η' , and p_{T} are used. To obtain η' , the four-momentum, p^ν , of the corresponding particle and the boost vector are needed. The latter describes the constant relative velocity, \vec{v} , of the two systems. The four-momentum is set by the momentum components of the particle in the laboratory system and the particle mass, from which the relativistic energy can be computed via the energy–momentum relation. The boost vector, $\vec{\beta} = \vec{v}/c$, has only one nonzero component, which is the z component. With the definition of the rapidity, this component is found to be

$$|\beta_z| = \frac{|v_z|}{c} = \tanh \Delta_y. \quad (6.6)$$

The contraction of the Lorentz matrix, $\Lambda^\mu{}_\nu(\vec{\beta})$, with p^ν results in the boosted four-momentum, p'^μ . The spatial components of p'^μ permit η' to be deduced.

The individual acceptance corrections for prompt charged pions, prompt charged kaons, prompt (anti)protons and other prompt long-lived charged particles are equal to the corresponding ratio of the second and first counts. The combined correction for prompt long-lived charged particles is determined by adding the products of the four individual corrections and the corrected fractions from Section 6.6.2. The acceptance correction is larger for heavier particles and at low p_{T} . It does not exceed 1 % for charged pions, charged kaons and long-lived charged particles in almost all kinematic intervals. For (anti)protons, the correction is at most 4 % at the lower ends of the accessible p_{T} ranges. No correction occurs for the p – p samples as both beam energies of the p – p collisions analysed are identical, causing the centre-of-mass and laboratory systems to be the same.

6.8 Propagation of uncertainty

The machinery used to analytically propagate all statistical and systematic uncertainties and correlations is taken from the first measurement of this thesis. It was designed to include correlations between η , p_{T} and charge for the computation of the charge-combined differential cross-section and the charge ratio. Besides the

charge ratio, three further ratios are present in this analysis: hadron ratios, nuclear-modification factors and forward-to-backward ratios. This respectively requires the correlations to be taken into account between the three hadron species, between the p -Pb and p - p samples as well as between the forward and backward p -Pb samples.

The types of correlation of the statistical and systematic uncertainties between the two charges and between the numerators and denominators of the other ratios are listed in Table 6.2. The statistical uncertainties are uncorrelated if they stem from independent (sub)samples. The type of systematic correlation depends on the evaluation or origin of the corresponding uncertainty. Some of the uncertainties on the track-reconstruction correction and on the proxy ratios are fully correlated by construction, but are assumed to be uncorrelated to avoid cancellation of the effects in the ratios. The statistical uncertainty on the numbers of charged-pion, charged-kaon and (anti)proton candidate tracks is partially anticorrelated between these three hadron species due to the template fits employed to determine their fractions. The non-Poisson nature of prompt charged-particle production, reflected by the subsampling approach presented in Section 5.6.1, induces positive statistical correlations of the numbers of candidate tracks and the numbers of background tracks in simulation between the kinematic intervals and between the two charges; the subsampling is also applied to the p -Pb samples. The relative uncertainties on the integrated luminosities, which are purely systematic, amount to 2.6 % and 2.5 % for the forward and backward p -Pb data samples, respectively. This uncertainty is assumed to have a partial correlation of 0.5 between the p -Pb and p - p data samples and between the forward and backward p -Pb data samples. This assumption was made in the corresponding calibration study, which is independent of the analysis, because the same techniques were used for the different data samples. The table indicates the majority of the effort invested in the accurate propagation of all correlations; the types of correlation between the kinematic intervals are not listed, but taken into account as well.

6.8.1 Total uncertainty

A breakdown of the total uncertainty on the charge-combined number of prompt long-lived charged particles, which does not include the uncertainty on the integrated luminosity, into the largest individual contributions is shown in Figure 6.14. The overall picture is similar to that obtained for p - p collisions in the first measurement. The leading uncertainty in almost all kinematic intervals is that on the correction to the efficiency in simulation, possessing a relative contribution of up to 5.9%. The increase in this uncertainty around $p_T \approx 1500$ MeV/ c occurs due to the particle-composition correction, whereas the rise towards high p_T is caused by the uncertainty propagated from the control study for the track-reconstruction efficiency. Another relevant source of uncertainty is the proxy ratio for fake tracks,

Table 6.2: Types of correlation of the statistical and systematic uncertainties between the two charges, between the three hadron species, between the p -Pb and p - p samples as well as between the forward and backward p -Pb samples. The uncertainties are either fully correlated (cor.), partially (anti)correlated [part. (anti)cor.] or uncorrelated (uncor.). The sources of uncertainty listed are the correction to the efficiency in simulation, given by the product of the track-reconstruction and particle-composition corrections, with the former possessing systematic uncertainties from the simulation weighting and hadronic interactions; the proxy ratios constructed for fake tracks, photon conversions, material interactions of charged pions and strange-hadron decays; the relative systematic uncertainty of 50 % assumed for the nonadjusted other sources of background; the noncorrected efficiency in simulation; the numbers of candidate tracks for the three hadron species and for long-lived charged particles; the numbers of background tracks in simulation as well as the integrated luminosity. The relation between these quantities is defined by Equations (5.1) and (6.1). *The uncertainty is of statistical origin, but was determined in a control study [129]; thus, it is systematic in this analysis. **The uncertainties are assumed to be uncorrelated to avoid cancellation in the ratio. ***The type of correlation depends on the kinematic interval.

Source	Type of correlation			
	Charges	Hadron species	p -Pb & p - p	Forward & backward
$u_{\text{stat}}(R_{\epsilon, \text{track}})^*$	Cor.	Cor.	Uncor.	Uncor.
$u_{\text{syst}, w}(R_{\epsilon, \text{track}})$	Cor.	Cor.	Uncor.**	Uncor.**
$u_{\text{syst}, h}(R_{\epsilon, \text{track}})$	Cor.	Cor.	Cor.	Cor.
$u_{\text{syst}}(R_{\epsilon, \text{comp}})$	Uncor.	—	Uncor.	Uncor.
$u_{\text{stat}}(R_{\text{fake}})$	Uncor.	Cor.	Uncor.	Uncor.
$u_{\text{syst}}(R_{\text{fake}})$	Uncor.	Cor.	Uncor.	Uncor.
$u_{\text{stat}}(R_{\text{mat}, \gamma})$	Uncor.	Cor.	Uncor.	Uncor.
$u_{\text{syst}}(R_{\text{mat}, \gamma})$	Uncor. or cor.***	Cor.	Uncor.	Uncor.
$u_{\text{stat}}(R_{\text{mat}, \pi})$	Uncor.	Uncor.**	Uncor.	Uncor.
$u_{\text{syst}}(R_{\text{mat}, \pi})$	Uncor. or cor.***	Uncor.**	Uncor.	Uncor.
$u_{\text{stat}}(R_{\text{strange}})$	Cor.	Uncor.	Uncor.	Uncor.
$u_{\text{syst}}(R_{\text{strange}})$	Cor.	Uncor.**	Uncor.	Uncor.
$u_{\text{syst}}(R_{\text{other}})$	Cor.	Uncor.**	Uncor.**	Uncor.**
$u_{\text{stat}}(\epsilon_{\text{sim}})$	Uncor.	Uncor.	Uncor.	Uncor.
$u_{\text{stat}}(n_{\text{cand}, h})$	Part. cor.	Part. anticor.	Uncor.	Uncor.
$u_{\text{syst}}(n_{\text{cand}, h})$	Uncor.	Uncor.	Uncor.	Uncor.
$u_{\text{stat}}(n_{\text{cand}, l})$	Part. cor.	—	Uncor.	Uncor.
$u_{\text{stat}}(n_{i, \text{sim}})$	Part. cor.	Uncor.	Uncor.	Uncor.
$u_{\text{syst}}(\mathcal{L})$	Cor.	Cor.	Part. cor.	Part. cor.

which prevails at high η and low p_T and contributes at most 7.4%. The statistical uncertainty on the efficiency in simulation reaches 3.1% at high p_T , but the impact on the total uncertainty is moderate since other uncertainties are also larger there. The remaining sources do not amount to more than 2%. The total relative uncertainty ranges between 2.0% and 8.0%.

For the numbers of prompt charged pions, prompt charged kaons and prompt (anti)protons, the largest individual uncertainty is that on the number of candidate tracks, which is dominated by the systematic contribution from the closure test of the PID-calibration method. In the kinematic intervals in which this uncertainty is smaller, the leading uncertainty is that on the correction to the efficiency in simulation. The remaining sources only marginally affect the total uncertainty. For the forward p -Pb samples, the total relative uncertainty is between 1.8% and 11.5% for charged pions, between 1.7% and 17.2% for charged kaons and between 2.1% and 17.6% for (anti)protons.

The individual and total uncertainties on the charge-combined number of prompt long-lived charged particles for p - p collisions can be compared with those obtained in the first measurement, which have been discussed in Section 5.6.2. The relative uncertainty from material interactions, induced by photons and charged pions, is now much smaller at high η and low p_T , namely 1% instead of 8%. This is enabled by the separate proxy for photon conversions, possessing a considerably smaller uncertainty in this region than the proxy for material interactions of charged pions, which was used to adjust the background from photon conversions in the first measurement. Moreover, the uncertainty on the proxy ratio for strange-hadron decays is now below 0.5% instead of 1.5% thanks to the three categories of the proxy ratio, allowing the systematic uncertainty to be reduced. The minimum value of the total relative uncertainty of 2.2% is now twice as large compared to the first measurement, in which it has amounted to 1.1%. The increase is caused by the larger uncertainty on the particle-composition correction, but this uncertainty is now estimated more reliably.

6.8.2 Data-sample combination

Potential bias from the configuration of the field of the LHCb magnet is investigated by a cross-check performed similarly to Section 5.6.3. The results should not depend on a detector configuration. Therefore, the differential cross-sections for the two charges are computed separately for the two p - p data samples, corresponding to the two magnetic-field configurations, and for the combination of both data samples. Unlike in the first measurement, the systematic uncertainties on the numbers of signal particles are not assumed to be fully correlated between the two data samples, but the mean systematic correlation is used for the combination. This choice is made because the individual results can be regarded as equally plausible.




Figure 6.14: Breakdown of the total relative uncertainty on the number of prompt long-lived charged particles in p -Pb collisions into the largest individual contributions in intervals of η and p_T . The uncertainties refer to the charge-combined number. Statistical and systematic uncertainties are added in quadrature if both exist. The contributions shown are the correction to the efficiency in simulation, given by the product of the track-reconstruction and particle-composition corrections; the proxy ratios constructed for fake tracks, photon conversions, material interactions of charged pions and strange-hadron decays, with the ratio for charged pions from K_S^0 decays displayed in the latter case; the relative systematic uncertainty of 50% assumed for the nonadjusted other sources of background as well as the statistical uncertainties on the efficiency in simulation, on the number of candidate tracks and on the number of fake tracks in simulation. The relation between these quantities is defined by Equation (6.1). The gaps between the right ends of the light-grey rectangles at low p_T , representing the minimum momentum of 2 GeV/c needed for the long-track reconstruction, and the left ends of the lines are caused by the minimum-momentum requirement of 6 GeV/c applied due to the p -Pb track-reconstruction correction. In the following, these gaps occur in all plots in which this correction enters, but not the hadron fractions (leading to larger gaps).

The cross-check is found to be successful. Any deviation of the individual results from the combined one is covered by the combined systematic uncertainty. Three steps of the analysis are essential for this cross-check to succeed at high η and low p_T . First, determining the proxy ratio for fake tracks separately for each of the two charges allows charge asymmetries to be considered. Second, the capability to measure the background from photon conversions results in an appropriate subtraction from the number of candidate tracks. Third, the spatial distributions of the primary vertices, which affect the efficiency at high η , are taken into account by the second simulation weighting as introduced in Section 6.3. Consequently, the combination of both p - p data samples is employed to compute all results besides the forward and backward p -Pb data samples.

6.9 Results

The numbers of particle species, charges and data samples analysed lead to a multitude of results; only one example plot is included in each of the following five subsections. The results are based on two samples of p -Pb collisions at $\sqrt{s_{NN}} = 8.16$ TeV respectively corresponding to integrated luminosities of $90 \mu\text{b}^{-1}$ and $71 \mu\text{b}^{-1}$ for positive and negative η values as well as on a sample of p - p collisions at $\sqrt{s} = 13$ TeV obtained from 4.6nb^{-1} of integrated luminosity. As the analysis is dominated by systematic uncertainties, the statistical contributions are not given separately, but the total uncertainty is reported, defined as the sum in quadrature of statistical and systematic uncertainties.

6.9.1 Differential cross-sections

The measured charge-combined differential cross-section of prompt production of (anti)protons, computed according to Equations (5.1) and (6.1), is shown in Figure 6.15. The minimum value of the total relative uncertainty on the measured differential cross-section, including the uncertainty on the integrated luminosity, amounts to 3.4%. Similar values are obtained for charged pions, charged kaons and long-lived charged particles as well as for the other data samples. These small uncertainties justify the measurement of absolute production cross-sections instead of only hadron ratios, in which the integrated luminosity cancels.

The results obtained for p -Pb collisions are compared with predictions from EPOS-LHC [26] and QGSJET II.04 [33], while those for p - p collisions are compared with predictions from PYTHIA 8.308 [38] with the default Monash 2013 tune [40], EPOS-LHC, QGSJET II.04 and SIBYLL 2.3d [35]. Each predicted differential cross-section is calculated as prescribed by Equation (5.11). In the first measurement, it was found that the values of the inelastic cross-section for p - p collisions imple-




Figure 6.15: Differential cross-section of prompt production of (anti)protons in p -Pb collisions at $\sqrt{s_{\text{NN}}} = 8.16$ TeV in intervals of η and p_{T} measured in the analysis and predicted by various hadronic-interaction models. The charge-combined values are shown. The bars indicate the total uncertainty.

mented in the models are compatible with measurements. They are therefore used in Equation (5.11). The implemented values of the inelastic cross-section for p -Pb collisions are listed in Table 6.3. They cannot be compared with measurements at the same nucleon-nucleon centre-of-mass energy since no experiment has ever measured the inelastic cross-section for p -Pb collisions at $\sqrt{s_{\text{NN}}} = 8.16$ TeV. Fortunately, an estimate of this cross-section can be computed from two measurements in p - p collisions at $\sqrt{s} = 8$ TeV, which are close in energy and also given in the table. To extrapolate from p - p to p -Pb collisions, the measured values are multiplied by $208^{2/3}$, where 208 is the mass number of the lead nucleus. This factor takes into account that the radius, r , is proportional to $208^{1/3}$ and the surface area (thus the cross-section) to r^2 if a spherical nucleus is assumed. Forming the weighted mean of the two extrapolated values with the inverse of the variance as weight results in an estimate of the inelastic cross-section for p -Pb collisions at $\sqrt{s_{\text{NN}}} = 8$ TeV of 2533.5 mb. The implemented values are significantly smaller than this mean extrapolated value. Hence, this value is inserted in Equation (5.11) to calculate the predictions for p -Pb collisions. Consequently, residual deviations in magnitude from the data can only be due to the predicted distributions of the numbers of signal

Table 6.3: Inelastic cross-sections implemented in various hadronic-interaction models and measured by various experiments. The weighted mean of the two extrapolated p -Pb cross-sections is computed with the inverse of the variance as weight.

Model or experiment	Inelastic cross-section [mb]		
	p -Pb at $\sqrt{s_{\text{NN}}} = 8.16$ TeV	p - p at $\sqrt{s} = 8$ TeV	p - p at $\sqrt{s} = 8$ TeV ($\times 208^{2/3}$)
EPOS-LHC [26]	1247.52	—	—
QGSJET II.04 [33]	2234.57	—	—
ATLAS [149]	—	71.7 ± 0.7	2518.1 ± 24.8
TOTEM [150]	—	74.7 ± 1.7	2622.4 ± 59.7
Weighted mean	—	—	2533.5

particles. The number of inelastic events generated for each model [32] is 10^8 . These events are also used for the model comparisons in Sections 6.9.2, 6.9.3, 6.9.4 and 6.9.5. The measured (anti)proton differential cross-section in p -Pb collisions for positive η values is reasonably described by EPOS at high η ; however, the prediction is too high at low η . QGSJET tends to underestimate the data towards high p_T , whereas it exceeds them otherwise.

The total correlation matrices of the measured differential cross-sections, containing all statistical and systematic contributions, resemble the matrix obtained in the first measurement. The correlations are positive as the systematic uncertainties prevail, which are often fully correlated between the two charges or between kinematic intervals.

6.9.2 Charge ratios

The ratio of the measured differential cross-sections of prompt production of long-lived positively and negatively charged particles, computed according to Equations (5.5) and (6.1), is shown in Figure 6.16. The total uncertainty grows towards high p_T ; nevertheless, a systematic rise of the measured charge ratio is observed in the $\eta \in [-5.5, -3.5)$ range, becoming more pronounced as $|\eta|$ increases. This can be attributed to the excess of positive charge in the initial state of a p -Pb collision, which transfers to the final state. At low p_T , the data are compatible with unity. The EPOS prediction only moderately rises with p_T . By contrast, the p_T dependence of QGSJET is similar to the data, but the model underestimates them at high p_T .




Figure 6.16: Ratio of the differential cross-sections of prompt production of long-lived positively and negatively charged particles in p -Pb collisions in intervals of η and p_T measured in the analysis and predicted by various hadronic-interaction models. The bars indicate the total uncertainty.

In the first measurement, the charge ratio for long-lived charged particles in p - p collisions decreased towards low p_T in the two highest intervals of η . However, it was still compatible with unity because of the large total uncertainty, evoked by assuming the uncertainty on the proxy ratio for material interactions in that kinematic region to be uncorrelated between the two charges. Thanks to the charge-asymmetric proxy ratio for fake tracks, the separate proxy for photon conversions and the second simulation weighting, that assumption can now be dropped. This leads to a well-controlled total uncertainty, which is at most 9% instead of more than 16%.

6.9.3 Hadron ratios

The ratio of the measured differential cross-sections of prompt production of charged kaons and charged pions, computed according to Equations (6.1) and (6.2), is shown in Figure 6.17. The data are best reproduced by SIBYLL, but at high p_T , they do not exhibit the predicted plateau. The PYTHIA and EPOS predictions are not sufficiently steep compared to the data. The rapid decrease towards high p_T modelled by QGSJET is not present in data in the $\eta \in [2.5, 4.0)$ range.




Figure 6.17: Ratio of the differential cross-sections of prompt production of charged kaons and charged pions in p - p collisions at $\sqrt{s} = 13$ TeV in intervals of η and p_T measured in the analysis and predicted by various hadronic-interaction models. The bars indicate the total uncertainty.

The agreement of the models with the measured (anti)proton-to-charged-pion ratio is good up to around $p_T \approx 1500$ MeV/c. At higher p_T , the predictions closest to the data are those from PYTHIA and EPOS; QGSJET underestimates the ratio, while SIBYLL overestimates it. The measured (anti)proton-to-charged-kaon ratio is well reflected by QGSJET at high p_T , where the other predictions are too high.

6.9.4 Nuclear-modification factors

The measured nuclear-modification factor of prompt production of long-lived charged particles, computed according to Equations (5.1) and (6.3), is shown in Figure 6.18. It can be determined in ten η intervals and is directly computed from the measured differential cross-sections in p -Pb and p - p collisions, which hold at $\sqrt{s_{NN}} = 8.16$ TeV and $\sqrt{s} = 13$ TeV, respectively. An interpolation is often used to estimate the denominator of the nuclear-modification factor at the same centre-of-mass energy as the numerator. Such an interpolation leads to a further systematic uncertainty from the assumptions about the energy dependence of the differential cross-section. Comparisons with predictions from hadronic-interaction models or dedicated theory calculations can also be performed if the energies are different. The




Figure 6.18: Nuclear-modification factor of prompt production of long-lived charged particles in p -Pb and p - p collisions in intervals of η and p_T measured in the analysis and predicted by various hadronic-interaction models. The (nucleon-nucleon) centre-of-mass energies of the two collision systems are given. The bars indicate the total uncertainty.

importance of measuring instead of interpolating the denominator is emphasised, for example, in References [151, 152]. Therefore, only the measured differential cross-sections enter the computation. For positive η values, the measured nuclear-modification factor falls to 0.3 at high η , indicating that the production in p -Pb collisions is suppressed relative to that in p - p collisions. The suppression is smaller at low η and at high p_T . For negative η values, the measured nuclear-modification factor reaches values above 0.7 around $p_T \approx 2000$ MeV/ c , signifying that the differential cross-sections per nucleon for the two collision systems are more similar. The η dependence is milder than for positive η values, but the p_T dependence is stronger. The total relative uncertainty achieved is small. It is between 4.4 % and 6.0 % around $p_T \approx 2000$ MeV/ c , and its minimum value across all kinematic intervals amounts to 3.6 %.

The model predictions are also shown in Figure 6.18. For positive η values and up to around $p_T \approx 3000$ MeV/ c , the data are reasonably reproduced by EPOS, their shape better than their magnitude; the predicted decrease towards high p_T is not observed in data. The behaviour of QGSJET disagrees with the data across the




Figure 6.19: Forward-to-backward ratio of prompt production of charged pions in p -Pb collisions in intervals of $|\eta|$ and p_T measured in the analysis and predicted by various hadronic-interaction models. The bars indicate the total uncertainty.

entire p_T ranges. For negative η values, the shapes of the two models are much more similar and resemble the measured shape. The predicted positions of the maximum values of the nuclear-modification factor coincide with the measured ones in the $\eta \in [-3.5, -2.5)$ range. In magnitude, the data lie between the two models. In most kinematic intervals, the total uncertainty on the measured nuclear-modification factor is sufficiently small to conclude that the EPOS and QGSJET predictions are incompatible with the data, which can help to improve the models in the future and demonstrates the impact of this measurement.

The measured suppression of light-hadron production in p -Pb collisions relative to p - p collisions for positive η values is often interpreted as a consequence of nuclear shadowing [153]. The behaviour of the nuclear-modification factor measured for negative η values might result from the Cronin effect [154]. However, also other cold-nuclear-matter effects could occur [151], or collective effects in dense-matter environments might be responsible [155] as presented in Section 3.2.

6.9.5 Forward-to-backward ratios

The measured forward-to-backward ratio of prompt production of charged pions, computed according to Equations (5.1) and (6.4), is shown in Figure 6.19. It can

be determined in those four $|\eta|$ intervals that overlap between the η grids for the forward and backward p -Pb samples. The measured forward-to-backward ratio is below unity, *i.e.* the production in p -Pb collisions for positive η values is suppressed relative to that for negative η values. The data tend to be flat in p_T , which is predicted by EPOS, but there are deviations in magnitude. Both the shape and the magnitude modelled by QGSJET do not agree with the data.

Part III
Implications

7 Validation of hadronic-interaction models

The prerequisites needed for the tuning of hadronic-interaction models, *i.e.* their systematic validation and improvement, are presented at the beginning of this chapter. Afterwards, a path towards global tuning is outlined; moreover, the hadron composition obtained in the second measurement of this thesis is interpreted with regard to the Muon Puzzle.

7.1 Comparisons with measurements

The performance of a hadronic-interaction model is assessed by comparing its predictions for certain observables with measurements. This requires the numerical values of the results to be documented, ideally in a machine-readable format and in a central database. Furthermore, selections have to be applied to the output of the corresponding event generator to isolate the process of interest. Storing such selections in a central place is beneficial to avoid reimplementations. For measurements from collider experiments, this preservation is realised by the HEPData [156] and RIVET [157, 158] systems.

HEPData is an open-access repository comprising data tables from published papers in experimental high-energy particle physics. More than 10 000 publications are already included, with approximately 20 % stemming from the LHC experiments. The results of a paper are stored in a record, which is a collection of tables and descriptions of their content. The tables can be exported in various formats.

RIVET is a library containing code to select the physics process analysed at generator level for each of numerous published papers. The selection code is written in the C++ programming language. Its main part is usually composed of iterations over the events and particles generated by a model, in which requirements are imposed on the species, ancestor and/or kinematics of the particles. This type of code is not to be confused with the analysis code of a measurement, which also applies selections, but these act on tracks and candidate decays in data to improve the signal-to-background ratio. In RIVET, a paper is preserved as a plugin, consisting of the selection code, metadata and plotting settings. The metadata declare, *e.g.*, the collision system and the beam energies. The plots created visualise

the measured observables from the paper and the predictions selected from the model. The values of the former are taken from HEPData, which can also provide definitions of kinematic intervals to RIVET. Thus, there is a strong link between the two systems. Despite the wide variety of measurements already included in RIVET, a long wish list still exists, in particular for the LHC experiments.

To reduce the length of the wish list, the author of this thesis contributed to the development of new RIVET plug-ins for LHCb measurements. The work was performed as part of two summer-student projects, taken by Gino Daniels and Joshua Friedman from Purdue University. Both of them worked under close supervision by Lars Kolk, doctoral student at TU Dortmund University, and the author. The preservation in RIVET was additionally motivated via the SFB 1491 collaborative research centre of the German Research Foundation, which was established in 2022. Two of the subprojects of SFB 1491 focus on hadronic interactions in particle and astroparticle physics. They require RIVET plug-ins of related measurements to pursue global tuning. Moreover, new HEPData records were created within one of the summer-student projects.

Six LHCb measurements that had not yet been completely preserved were identified by Lars Kolk and the author. It was ensured that the measurements were (highly) desired by the RIVET team; furthermore, they are relevant for hadronic interactions in the atmosphere. Three measurements of beauty production were chosen: beauty-quark production in p - p collisions [159], beauty-hadron fragmentation fractions in p - p collisions [160] and beauty-hadron production in p -Pb collisions [161]. In addition, two measurements of charm production were identified: D^0 -meson production in p -Pb collisions [162] and J/ψ -meson production in p -Pb collisions [163]. These five measurements cover heavy-quark production and hadronisation in the forward η region at high energies. Heavy hadrons contribute little to air showers, but neutrinos originating from their decays form the prompt atmospheric neutrino flux, which is an important background for high-energy neutrino observatories [2]. Lastly, the measurement of charged-particle production in p -Pb collisions [117] was chosen. Three of the six measurements did not possess a HEPData record, whereas no RIVET plug-in existed for five of the six measurements. The records and plug-ins were drafted within the summer-student projects. Lars Kolk and the author then took care of the review processes of the three new HEPData records and the five new RIVET plug-ins until their publication in References [164–171].

The first measurement performed within this thesis, charged-particle production in p - p collisions [8], was preserved in HEPData [172] and RIVET [173] by Lars Kolk in consultation with Hans Dembinski. Since the second measurement performed by the author has not yet been published, this preservation is not yet possible.

An analogue to the HEPData system also exists for astroparticle experiments, namely the CRDB system [174]. By contrast, a library similar to the RIVET system,

which would use air-shower simulation as input, does not yet exist, but would be beneficial for global tuning.

7.2 Towards a global picture

The authors of the hadronic-interaction models provide them in various versions or tunes with optimal values for the parameters. These values are determined based on various measurements, which are chosen by expert knowledge and usually stem from collider experiments. Most of those measurements cover the central η region [2]. Consequently, the parameter values might not be optimal for the modelling in the forward η region. The discrepancies observed in the comparisons of the results obtained within this thesis with the predictions hint at room for improvements in the models. Measurements from astroparticle experiments are in general not considered for tuning, which is unfortunate with regard to the Muon Puzzle. In global tuning, the models would be validated against data both from collider and astroparticle experiments. Such an approach is investigated within the SFB 1491 collaborative research centre and is presented in Reference [175]. So far, global tuning was only explored in Reference [176].

Including a multitude of measurements requires an automatic tuning procedure [175]. For collider data, the availability of a measurement as HEPData record and RIVET plug-in decides if it can be considered. An event generator would first be run with its default parameter values for the collision systems and beam energies of all the measurements. The deviations of the predictions from the data can be summarised in a χ^2 value. Infrastructure already exists for the creation of comparison plots and the χ^2 computation; the MCPlots website [177, 178] employs pregenerated model output for this purpose. The χ^2 value obtained would then be passed to a tuning algorithm [179–182], which determines new estimates for the parameters [175]. Afterwards, the loop restarts. Complications arise as not all measurements can be regarded as equally relevant due to the physics processes implemented in the model under study. Therefore, weights should be assigned to the measurements, requiring expert knowledge. This is also required for the parameters, which cannot be arbitrarily varied; there might also be important relations between them. Moreover, the feasibility of the tuning loop depends on the interface of the model. In the case of PYTHIA, all parameters are extensively documented and can be modified by the user to create custom tunes. For the EPOS, QGSJET and SIBYLL models, accessing and modifying the parameters is more difficult because these models were not originally designed for user-based tuning.

For astroparticle data, a future RIVET-like library or an extension of RIVET would not only steer a hadronic-interaction model, but full air-shower simulation, which additionally includes the particle transport in the atmosphere among other as-

pects [175]. The library would execute the simulation for the observation conditions of the measurements, for example ranges in cosmic-ray energy and zenith angle. The χ^2 value, computed from the predicted and measured observables, such as X_{\max} and $N_{\mu'}$, would enter the tuning algorithm together with the χ^2 value from the collider data to iterate towards a global minimum. It is not guaranteed that the optimised parameter values enable a reasonable description of all the data; global tuning might reveal that extensions of the models are required. This global approach, connecting particle and astroparticle physics, is in reach, but multiple steps still need to be developed. A viable possibility of establishing a connection between the LHCb measurements presented in this thesis and astroparticle physics lies in the R energy ratio, defined in Equation (3.1).

This ratio results from the mean electromagnetic- and hadronic-energy pseudorapidity densities. In the second measurement of this thesis, the hadron composition in p - p collisions at $\sqrt{s} = 13$ TeV is measured in the $\eta \in [2.0, 5.0]$ range at intermediate p_T values. For the particle-composition correction, an extrapolation of the measured hadron fractions is performed based on simulation to entirely cover the $p_T < 10$ GeV/ c range as described in Section 6.6.2. The minimum accessible p_T value is evoked by the minimum momentum of 2 GeV/ c needed for the long-track reconstruction and thus depends on η . The extrapolations represent the corrected fractions of charged pions, charged kaons and (anti)protons. The corrected fraction of other long-lived charged particles is computed via the constraint on the sum of the fractions to give unity. These fractions can be used in a hadronic-interaction model to estimate the R_{LHCb} energy ratio, where the index signifies that the evaluation is restricted to the LHCb acceptance via the η , p_T and momentum requirements. The densities are determined by adding the corrected energies of the corresponding particles in each event and averaging over a sufficiently large number of generated events. The corrected energies are obtained by multiplying the initial energies in the model by the ratio of the corrected and initial particle fractions. The electromagnetic contribution demands corrected fractions of photons and electrons, which are approximated by the corrected charged-pion fraction. The motivation for this is the same as in Section 5.4.2; most of the photons are produced in decays of neutral pions, of which the production is coupled to that of charged pions because of isospin symmetry. The hadronic contribution requires corrected fractions of all long-lived hadrons, for which several assumptions are made. The corrected charged-kaon fraction is used for K_S^0 and K_L^0 mesons as well as for Λ and $\bar{\Lambda}$ baryons due to the single strange valence (anti)quark. The (anti)neutrons are assigned the corrected (anti)proton fraction, while Ξ^0 and $\bar{\Xi}^0$ baryons are treated like the Σ^+ , Σ^- , Ξ^- and Ω^- baryons. The resulting R_{LHCb} values for all models studied are shown in Figure 7.1. They are compared with the values obtained from the initial energies in the models. It is found that the incorporation of the corrected




Figure 7.1: Ratio of the mean electromagnetic- and hadronic-energy pseudorapidity densities restricted to the LHCb acceptance initially predicted by various hadronic-interaction models and corrected for the hadron composition measured in p - p collisions at $\sqrt{s} = 13$ TeV.

particle fractions leads to a reduction in R_{LHCb} that is close to 15%. The R_{LHCb} value depends on the model, but the amount of reduction is similar for all models. As detailed in Section 3.2, a reduction in R of 15% has the potential to explain the Muon Puzzle [69]. Such a reduction implies that more energy is retained in the hadronic cascade of a simulated air shower and thus available for muon production. Fewer neutral and charged pions are then produced, which is indicated by the measurement presented. This change in the hadron composition might originate from collective effects in dense-matter environments, in which case the models would underestimate them in p - p collisions. No uncertainty on R_{LHCb} is given since the sole goal is to emphasise the consistent reduction. However, the findings should not be overinterpreted owing to the restriction to the LHCb acceptance and the assumptions on the particle fractions, inducing differences between R and R_{LHCb} .

8 Conclusion and outlook

Two measurements of prompt production of light hadrons performed with the LHCb experiment at CERN are presented in this thesis. In the first measurement, which has been published in Reference [8], long-lived charged particles are studied in p - p collisions at $\sqrt{s} = 13$ TeV. The differential production cross-section is measured in intervals of η and p_T and is determined separately for positively and negatively charged particles. The analysis constitutes the highest-energy study of charged-particle production in the forward η region of p - p collisions to date. Particular effort is invested in the accurate propagation of all statistical and systematic correlations. The most precise values of the differential cross-section across all intervals possess a total uncertainty of less than 3%, which includes the external uncertainty on the integrated luminosity of 2%. The total uncertainty is dominated by the systematic contribution, whereas the statistical contribution is small. The second measurement extends the first one in two ways; the three most abundant charged hadrons—charged pions, charged kaons and (anti)protons—are measured separately via PID information, and data from p -Pb collisions at $\sqrt{s_{NN}} = 8.16$ TeV are analysed as well, enabling nuclear effects to be investigated. Hadron ratios, nuclear-modification factors and forward-to-backward ratios are measured besides differential cross-sections and charge ratios. So far, an absolute measurement of identified-hadron production in the forward η region has not yet been conducted in p - p collisions at an energy as high as $\sqrt{s} = 13$ TeV and not yet at all in p -Pb collisions. A comprehensive PID-calibration method is established to take into account imperfections of the detector simulation compared to data in the determination of the charged-pion, charged-kaon and (anti)proton fractions. The method uses machine-learning techniques: gradient-boosted decision trees and multilayer-perceptron neural networks. The implementation of the latter is taken from another LHCb measurement at lower energy. Multiple modifications have been done to allow, for the first time, the application to data from p - p and p -Pb collisions, with a clear gain both in physics performance and run time. Various statistical techniques, *e.g.* template fits and *sWeights*, are employed as well. The results of the two measurements are compared with predictions from widely used hadronic-interaction models: PYTHIA 8.308, EPOS-LHC, QGSJET II.04 and SIBYLL 2.3d. None of the models accomplishes a satisfactory description of all the observables within the total measurement uncertainties; these are also much smaller than the spread of the model predictions in many kinematic intervals. The two measurements of light-hadron

production in the forward η region are important ingredients for future tuning of the models, from which both particle and astroparticle physics can benefit. In collider experiments, light-hadron production often forms the underlying event, and the correct simulation of those processes is required for reliable background estimation. Moreover, several astroparticle experiments observe a deficit in the number of muons produced in simulated high-energy air showers compared to data, known as the Muon Puzzle. The two measurements, performed in p - p and p -Pb collisions at the highest energies that are nowadays available, can constrain the processes in the atmosphere since the two collision systems enclose the dominant type of first interaction, which is protons interacting with nitrogen or oxygen, on a mass scale. The hadron composition in p - p collisions at $\sqrt{s} = 13$ TeV determined in the second measurement is placed in the context of the Muon Puzzle via the ratio of the mean electromagnetic- and hadronic-energy pseudorapidity densities in the models. The measured composition indicates a reduction in this energy ratio of approximately 15% in the LHCb acceptance, caused by a decrease in charged-pion production relative to the initial model predictions. It is notable that a reduction in the energy ratio of 15%, without restriction to the LHCb acceptance, was found to be capable of solving the Muon Puzzle. The reason for the decrease in charged-pion production might lie in mechanisms that alter the particle production at high energy density, which could be underestimated in the models in the case of small collision systems, like p - p collisions. Although the sizes of the reduction are similar, the interpretation requires care due to the limited η range and the assumptions made in the LHCb acceptance. These concern the extrapolations of the measured hadron composition towards low p_T and the unmeasured long-lived neutral particles. Full air-shower simulation with new tunes of the models would be needed to validate this size of the reduction. To draw firm conclusions on the Muon Puzzle, the hadron composition should be measured in a collision system that resembles the interactions in the atmosphere even more. The LHCb experiment, despite its unique forward coverage, does not extend as forward as the particles produced in showers. An experiment that can directly measure the energy ratio more forward would be ideal. These and other potential future studies are described in the following paragraph. Furthermore, the author of this thesis contributed to the preservation of other LHCb measurements in the HEPData and RIVET systems. Those measurements deal with beauty, charm and charged-particle production in p - p and p -Pb collisions and therefore provide information on heavy- and light-hadron production in the forward η region, which is relevant for interactions in the atmosphere.

The two measurements conducted could be developed further in the future. In the light of the proton–oxygen collisions currently planned for 2025, a desirable next step would be the application of the entire analysis workflows to those data. Measuring light-hadron production in proton–oxygen collisions would constrain hadronic-interaction models more directly as no interpolation between p - p and p -Pb

collisions would be needed. In addition, the measurements could be advanced by determining all observables not only in intervals of η and p_T , but also for various classes of a multiplicity proxy. This would permit another search for strangeness enhancement in the forward η region via the charged-kaon-to-charged-pion ratio. Unlike most measurements at the LHCb experiment, analyses of light-hadron production do not substantially profit from the new wealth of data recorded with the upgraded detector. This is because they are not limited by the statistical uncertainty. Nevertheless, larger data samples might help in calibration studies, for example of the track-reconstruction efficiency, and could then enable a smaller systematic uncertainty to be achieved. The RICH detectors restrict the accessible p_T ranges for the determination of the hadron composition, which is unfortunate at low p_T , where QCD processes with small momentum transfer reside. An exploitation of this region would probably require another technique for charged-hadron identification, *e.g.* based on energy loss in a material. The two measurements presented should be included in the tuning of the models in the future. A starting point could be the creation of new tunes using only these measurements as additional input to investigate if an increase in the number of muons in simulated air showers is found. Later, the optimum would be the inclusion in global tuning together with a variety of other collider and astroparticle measurements. Further studies are in general needed concerning the Muon Puzzle. Apparently, the observation of a muon deficit in simulation depends on the muon contamination in the cosmic-ray energy estimator, but this explanation is not entirely conclusive. Complementary information could be provided by future measurements conducted at the Pierre Auger Observatory. As part of the Auger upgrade, the Cherenkov stations were equipped with radioantennas. These can measure the energy in the electromagnetic cascade of a shower, which is an alternative to the fluorescence telescopes [183]. Moreover, the energy ratio discussed above should be directly measured at a collider as forward as possible. Such a measurement could be undertaken with electromagnetic and hadronic calorimeters covering the $\eta > 7$ range [2].

The work performed within this thesis demonstrates the connection between high-energy particle and astroparticle physics. Precision data from experiments at particle colliders can aid the interpretation of air-shower measurements with regard to the mass composition of ultrahigh-energy cosmic rays. This interpretation relies on hadronic-interaction models, which have to be extensively validated. Measurements of light-hadron production in the forward η region at the highest energies, which can be performed with the LHCb experiment, are particularly valuable for this purpose. A profound understanding of hadronic interactions is ultimately relevant for unveiling the sources of ultrahigh-energy cosmic rays and shedding light on physics processes under extreme conditions, occurring in distant places in the universe.

A Basis vectors for two-body decays

The variables explained in this chapter are used to select the samples for the calibration of the PID response as presented in Section 6.5.1. They were derived by Michael Schmelling, with whom the author of this thesis worked together. Since no published reference is available yet, the derivation is reproduced below.

In three dimensions and in events with exactly one primary vertex, the entire information on a two-body decay of a particle originating from the primary vertex is contained in 13 parameters: the position of the primary vertex (3 coordinates), the positions of the two decay-product trajectories at a fixed z coordinate (2×2 coordinates) and the momentum vectors, \vec{p}_1 and \vec{p}_2 , of the two decay products at these positions (2×3 coordinates). Translation and rotation invariance each decrease the number of parameters by three (3 coordinates and 3 Euler angles). Thus, seven parameters remain, which are irreducible and therefore optimal. They can be constructed by shifting the primary vertex to the origin of the coordinate system. Moreover, the trajectories are assumed to be straight lines, and the momentum is assumed not to change along the trajectory. The straight-line assumption is a good approximation in the VELO, in which the decays analysed happen, due to its location outside the field of the LHCb magnet. Consequently, the trajectory positions can be replaced by vectors, \vec{q}_1 and \vec{q}_2 , that are orthogonal to the momentum vectors. As the seven parameters have to be rotation invariant, all possible scalar products and scalar triple products, which are rotation-invariant quantities in three dimensions, of \vec{p}_1 , \vec{p}_2 , \vec{q}_1 and \vec{q}_2 are formed. This results in

$$\begin{aligned}
 p_{11} &= \vec{p}_1 \cdot \vec{p}_1, & p_{22} &= \vec{p}_2 \cdot \vec{p}_2, & p_{12} &= \vec{p}_1 \cdot \vec{p}_2, \\
 q_{11} &= \vec{q}_1 \cdot \vec{q}_1, & q_{22} &= \vec{q}_2 \cdot \vec{q}_2, & q_{12} &= \vec{q}_1 \cdot \vec{q}_2, \\
 Q_{11} &= \vec{q}_1 \cdot \vec{p}_1, & Q_{22} &= \vec{q}_2 \cdot \vec{p}_2, & Q_{12} &= \vec{q}_1 \cdot \vec{p}_2, & Q_{21} &= \vec{q}_2 \cdot \vec{p}_1
 \end{aligned} \tag{A.1}$$

as well as

$$\begin{aligned}
 A_1 &= \vec{q}_1 \cdot (\vec{p}_1 \times \vec{p}_2), & A_2 &= \vec{q}_2 \cdot (\vec{p}_1 \times \vec{p}_2), \\
 B_1 &= \vec{p}_1 \cdot (\vec{q}_1 \times \vec{q}_2), & B_2 &= \vec{p}_2 \cdot (\vec{q}_1 \times \vec{q}_2).
 \end{aligned} \tag{A.2}$$

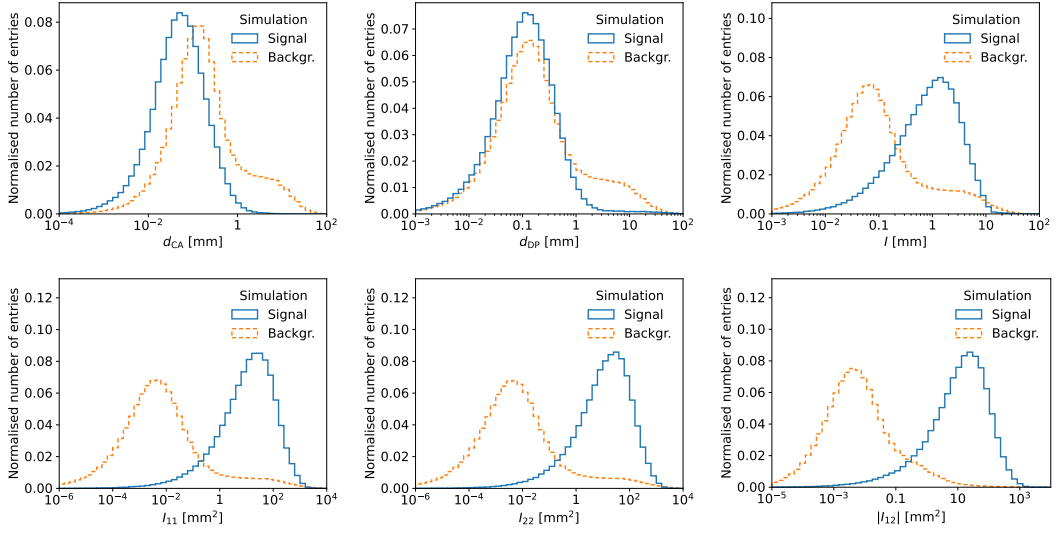


Figure A.1: Distributions of the six geometric variables from the two-body-decay basis for $K_S^0 \rightarrow \pi^+\pi^-$ candidates in p -Pb simulation. The signal and background distributions refer to $K_S^0 \rightarrow \pi^+\pi^-$ decays and random track combinations, respectively.

Out of these potential parameters, seven are chosen: $p_{11}, p_{22}, p_{12}, Q_{12}, Q_{21}, A_1$ and A_2 . The other parameters either vanish by construction or can be expressed through the seven parameters chosen according to

$$\begin{aligned}
 Q_{11} = Q_{22} = 0, \\
 q_{11} = \frac{A_1^2 + p_{11}Q_{12}^2}{N^2}, \quad q_{22} = \frac{A_2^2 + p_{22}Q_{21}^2}{N^2}, \quad q_{12} = \frac{A_1A_2 - p_{12}Q_{12}Q_{21}}{N^2}, \quad (\text{A.3}) \\
 B_1 = \frac{p_{11}Q_{12}A_2 + p_{12}Q_{21}A_1}{N^2}, \quad B_2 = \frac{p_{22}Q_{21}A_1 + p_{12}Q_{12}A_2}{N^2},
 \end{aligned}$$

where $N^2 = p_{11}p_{22} - p_{12}^2$ is used. The seven parameters chosen do not possess an intuitive meaning. Hence, six geometric variables are deduced. The distance, d_{CA} , of closest approach of the decay products; the distance, d_{DP} , of the primary vertex from the decay plane and the in-decay-plane impact parameter, I , of the decaying particle at the primary vertex are ideally zero. They are defined by

$$d_{CA} = \frac{|A_1 - A_2|}{N}, \quad d_{DP} = \frac{|A_1 + A_2|}{N}, \quad I = \frac{|p_{22}Q_{21} - p_{11}Q_{12}|}{N\sqrt{p_{11} + p_{22} + 2p_{12}}}. \quad (\text{A.4})$$

Finite lifetime of the decaying particle is probed by the squared in-decay-plane impact parameters, I_{11} and I_{22} , of the decay products at the primary vertex and by

the remaining contribution, I_{12} , to the flight distance of the decaying particle that is not covered by I_{11} and I_{22} . These variables are greater than zero for finite-lifetime particles and given by

$$I_{11} = \frac{p_{11}Q_{12}^2}{N^2}, \quad I_{22} = \frac{p_{22}Q_{21}^2}{N^2}, \quad I_{12} = \frac{p_{12}Q_{12}Q_{21}}{N^2}. \quad (\text{A.5})$$

The seventh variable deduced is not a geometric variable, but a kinematic one. It is chosen to be the invariant mass, which depends on p_{11} , p_{22} and p_{12} . In summary, the six geometric variables— d_{CA} , d_{DP} , I , I_{11} , I_{22} and I_{12} —and the invariant mass represent a basis for two-body decays because they are a linearly independent spanning set.

The distributions of the six geometric variables in simulation are shown in Figure A.1. The d_{CA} and d_{DP} distributions for the signal from $K_S^0 \rightarrow \pi^+ \pi^-$ decays behave as expected since they peak close to zero. There is limited separation power in each of these two variables as the background distributions largely overlap with the signal distributions. The signal I distribution peaks at larger values than the background distribution, which must be caused by the detector resolution since the I variable is ideally zero for a two-body decay. However, this observation is not relevant for the selection. The signal I_{11} , I_{22} and $|I_{12}|$ distributions occur at large values and are clearly separated from the background distributions. This is expected due to the long K_S^0 lifetime, causing the decay vertex to be detached from the primary vertex.

Bibliography

- [1] C.M.G. Lattes, G.P.S. Occhialini and C.F. Powell. 'Observations on the tracks of slow mesons in photographic emulsions'. In: *Nature* 160, 453 (1947). doi: 10.1038/160453a0.
- [2] J. Albrecht *et al.* 'The Muon Puzzle in cosmic-ray induced air showers and its connection to the Large Hadron Collider'. In: *Astrophys. Space Sci.* 367, 27 (2022). doi: 10.1007/s10509-022-04054-5.
- [3] F. Halzen and A.D. Martin. *Quarks and leptons: An introductory course in modern particle physics*. John Wiley & Sons, 1984.
- [4] T.K. Gaisser, R. Engel and E. Resconi. *Cosmic rays and particle physics: Second edition*. Cambridge University Press, 2016.
- [5] D. d'Enterria *et al.* 'Constraints from the first LHC data on hadronic event generators for ultra-high energy cosmic-ray physics'. In: *Astropart. Phys.* 35, 98 (2011). doi: 10.1016/j.astropartphys.2011.05.002.
- [6] LHCb collaboration (A.A. Alves Jr. *et al.*). 'The LHCb detector at the LHC'. In: *J. Instrum.* 3, S08005 (2008). doi: 10.1088/1748-0221/3/08/S08005.
- [7] LHCb collaboration (R. Aaij *et al.*). 'LHCb detector performance'. In: *Int. J. Mod. Phys. A* 30, 1530022 (2015). doi: 10.1142/S0217751X15300227.
- [8] LHCb collaboration (R. Aaij *et al.*). 'Measurement of prompt charged-particle production in pp collisions at $\sqrt{s} = 13$ TeV'. In: *J. High Energy Phys.* 01, 166 (2022). doi: 10.1007/JHEP01(2022)166.
- [9] S. Weinberg. 'Essay: Half a century of the Standard Model'. In: *Phys. Rev. Lett.* 121, 220001 (2018). doi: 10.1103/PhysRevLett.121.220001.
- [10] ATLAS collaboration (G. Aad *et al.*). 'Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC'. In: *Phys. Lett. B* 716, 1 (2012). doi: 10.1016/j.physletb.2012.08.020.
- [11] CMS collaboration (S. Chatrchyan *et al.*). 'Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC'. In: *Phys. Lett. B* 716, 30 (2012). doi: 10.1016/j.physletb.2012.08.021.

- [12] J.A. Formaggio, A.L.C. de Gouvêa and R.G.H. Robertson. ‘Direct measurements of neutrino mass’. In: *Phys. Rep.* 914, 1 (2021). doi: 10.1016/j.physrep.2021.02.002.
- [13] Belle collaboration (S.-K. Choi *et al.*). ‘Observation of a narrow charmoniumlike state in exclusive $B^\pm \rightarrow K^\pm \pi^+ \pi^- J/\psi$ decays’. In: *Phys. Rev. Lett.* 91, 262001 (2003). doi: 10.1103/PhysRevLett.91.262001.
- [14] LHCb collaboration (R. Aaij *et al.*). ‘Observation of $J/\psi p$ resonances consistent with pentaquark states in $\Lambda_b^0 \rightarrow J/\psi K^- p$ decays’. In: *Phys. Rev. Lett.* 115, 072001 (2015). doi: 10.1103/PhysRevLett.115.072001.
- [15] LHCb collaboration (R. Aaij *et al.*). ‘Observation of a narrow pentaquark state, $P_c(4312)^+$, and of the two-peak structure of the $P_c(4450)^+$ ’. In: *Phys. Rev. Lett.* 122, 222001 (2019). doi: 10.1103/PhysRevLett.122.222001.
- [16] T.S. Virdee. ‘Beyond the Standard Model of particle physics’. In: *Philos. Trans. Royal Soc. A* 374, 20150259 (2016). doi: 10.1098/rsta.2015.0259.
- [17] Particle Data Group (S. Navas *et al.*). ‘Review of particle physics’. In: *Phys. Rev. D* 110, 030001 (2024). doi: 10.1103/PhysRevD.110.030001.
- [18] T. Regge. ‘Introduction to complex orbital momenta’. In: *Il Nuovo Cimento* 14, 951 (1959). doi: 10.1007/BF02728177.
- [19] V.N. Gribov. ‘A reggeon diagram technique’. In: *Sov. Phys. JETP* 26, 414 (1968).
- [20] F.E. Low. ‘Model of the bare pomeron’. In: *Phys. Rev. D* 12, 163 (1975). doi: 10.1103/PhysRevD.12.163.
- [21] H.J. Drescher *et al.* ‘Parton-based Gribov–Regge theory’. In: *Phys. Rep.* 350, 93 (2001). doi: 10.1016/S0370-1573(00)00122-8.
- [22] K. Werner, F.-M. Liu and T. Pierog. ‘Parton ladder splitting and the rapidity dependence of transverse momentum spectra in deuteron-gold collisions at the BNL Relativistic Heavy Ion Collider’. In: *Phys. Rev. C* 74, 044902 (2006). doi: 10.1103/PhysRevC.74.044902.
- [23] S. Ostapchenko. ‘Nonlinear screening effects in high energy hadronic interactions’. In: *Phys. Rev. D* 74, 014026 (2006). doi: 10.1103/PhysRevD.74.014026.
- [24] R.S. Fletcher *et al.* ‘Sibyll: An event generator for simulation of high energy cosmic ray cascades’. In: *Phys. Rev. D* 50, 5710 (1994). doi: 10.1103/PhysRevD.50.5710.

-
- [25] Pierre Auger collaboration (A. Aab *et al.*). ‘Measurement of the fluctuations in the number of muons in extensive air showers with the Pierre Auger Observatory’. In: *Phys. Rev. Lett.* 126, 152002 (2021). doi: 10.1103/PhysRevLett.126.152002.
- [26] T. Pierog *et al.* ‘EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider’. In: *Phys. Rev. C* 92, 034906 (2015). doi: 10.1103/PhysRevC.92.034906.
- [27] S. Ostapchenko *et al.* ‘Constraining high energy interaction mechanisms by studying forward hadron production at the LHC’. In: *Phys. Rev. D* 94, 114026 (2016). doi: 10.1103/PhysRevD.94.114026.
- [28] T. Pierog. ‘Air shower simulation with a new generation of post-LHC hadronic interaction models in CORSIKA’. In: *Proc. Sci. ICRC2017*, 1100 (2018). doi: 10.22323/1.301.1100.
- [29] K. Werner. ‘Core-corona separation in ultrarelativistic heavy ion collisions’. In: *Phys. Rev. Lett.* 98, 152301 (2007). doi: 10.1103/PhysRevLett.98.152301.
- [30] K. Werner *et al.* ‘The physics of EPOS’. In: *Eur. Phys. J. Web Conf.* 52, 05001 (2013). doi: 10.1051/epjconf/20125205001.
- [31] LHCb collaboration (R. Aaij *et al.*). ‘Measurement of charged particle multiplicities and densities in pp collisions at $\sqrt{s} = 7$ TeV in the forward region’. In: *Eur. Phys. J. C* 74, 2888 (2014). doi: 10.1140/epjc/s10052-014-2888-1.
- [32] H. Dembinski, A. Fedynitch and A. Prosekin. ‘Chromo: An event generator frontend for particle and astroparticle physics’. In: *Proc. Sci. ICRC2023*, 189 (2024). doi: 10.22323/1.444.0189.
- [33] S. Ostapchenko. ‘Monte Carlo treatment of hadronic interactions in enhanced pomeron scheme: QGSJet-II model’. In: *Phys. Rev. D* 83, 014018 (2011). doi: 10.1103/PhysRevD.83.014018.
- [34] S. Ostapchenko. ‘Total and diffractive cross sections in enhanced pomeron scheme’. In: *Phys. Rev. D* 81, 114028 (2010). doi: 10.1103/PhysRevD.81.114028.
- [35] F. Riehn *et al.* ‘Hadronic interaction model Sibyll 2.3d and extensive air showers’. In: *Phys. Rev. D* 102, 063002 (2020). doi: 10.1103/PhysRevD.102.063002.
- [36] E.-J. Ahn *et al.* ‘Cosmic ray interaction event generator Sibyll 2.1’. In: *Phys. Rev. D* 80, 094003 (2009). doi: 10.1103/PhysRevD.80.094003.
- [37] J. Ranft. ‘Dual parton model at cosmic ray energies’. In: *Phys. Rev. D* 51, 64 (1995). doi: 10.1103/PhysRevD.51.64.

- [38] C. Bierlich *et al.* ‘A comprehensive guide to the physics and usage of Pythia 8.3’. In: *SciPost Phys. Codeb.* 8 (2022). doi: 10.21468/SciPostPhysCodeb.8.
- [39] M. Fieg *et al.* ‘Tuning Pythia for forward physics experiments’. In: *Phys. Rev. D* 109, 016010 (2024). doi: 10.1103/PhysRevD.109.016010.
- [40] P. Skands, S. Carrazza and J. Rojo. ‘Tuning Pythia 8.1: The Monash 2013 tune’. In: *Eur. Phys. J. C* 74, 3024 (2014). doi: 10.1140/epjc/s10052-014-3024-y.
- [41] C. Bierlich *et al.* ‘The Angantyr model for heavy-ion collisions in Pythia 8’. In: *J. High Energy Phys.* 10, 134 (2018). doi: 10.1007/JHEP10(2018)134.
- [42] T. Pierog, M. Perlin and K. Werner. ‘Hadronic interactions studies at the LHC’. In: *Eur. Phys. J. Web Conf.* 283, 05001 (2023). doi: 10.1051/epjconf/202328305001.
- [43] D. d’Enterria, T. Pierog and G. Sun. ‘Impact of QCD jets and heavy-quark production in cosmic-ray proton atmospheric showers up to 10^{20} eV’. In: *Astrophys. J.* 874, 152 (2019). doi: 10.3847/1538-4357/ab01e2.
- [44] G.V. Kulikov and G.B. Khristiansen. ‘On the size spectrum of extensive air showers’. In: *Sov. Phys. JETP* 35, 441 (1959).
- [45] D.J. Bird *et al.* ‘Evidence for correlated changes in the spectrum and composition of cosmic rays at extremely high energies’. In: *Phys. Rev. Lett.* 71, 3401 (1993). doi: 10.1103/PhysRevLett.71.3401.
- [46] High Resolution Fly’s Eye collaboration (R.U. Abbasi *et al.*). ‘First observation of the Greisen-Zatsepin-Kuzmin suppression’. In: *Phys. Rev. Lett.* 100, 101101 (2008). doi: 10.1103/PhysRevLett.100.101101.
- [47] Pierre Auger collaboration (J. Abraham *et al.*). ‘Observation of the suppression of the flux of cosmic rays above 4×10^{19} eV’. In: *Phys. Rev. Lett.* 101, 061101 (2008). doi: 10.1103/PhysRevLett.101.061101.
- [48] M. Unger, G.R. Farrar and L.A. Anchordoqui. ‘Origin of the ankle in the ultrahigh energy cosmic ray spectrum, and of the extragalactic protons below it’. In: *Phys. Rev. D* 92, 123001 (2015). doi: 10.1103/PhysRevD.92.123001.
- [49] K.-H. Kampert and M. Unger. ‘Measurements of the cosmic ray composition with air shower experiments’. In: *Astropart. Phys.* 35, 660 (2012). doi: 10.1016/j.astropartphys.2012.02.004.
- [50] W. Heitler. *The quantum theory of radiation*. Oxford University Press, 1944.
- [51] J. Matthews. ‘A Heitler model of extensive air showers’. In: *Astropart. Phys.* 22, 387 (2005). doi: 10.1016/j.astropartphys.2004.09.003.

-
- [52] Pierre Auger collaboration (A. Aab *et al.*). 'The Pierre Auger Cosmic Ray Observatory'. In: *Nucl. Instrum. Methods Phys. Res. A* 798, 172 (2015). doi: 10.1016/j.nima.2015.06.058.
- [53] T.H.-J. Mathes (for the Pierre Auger collaboration). 'The HEAT telescopes of the Pierre Auger Observatory: Status and first data'. In: *32nd International Cosmic Ray Conference* 153 (2011).
- [54] A.M. Botti (for the Pierre Auger collaboration). 'The AMIGA underground muon detector of the Pierre Auger Observatory - Performance and event reconstruction'. In: *Proc. Sci. ICRC2019*, 202 (2021). doi: 10.22323/1.358.0202.
- [55] T. Abu-Zayyad *et al.* 'The surface detector array of the Telescope Array experiment'. In: *Nucl. Instrum. Methods Phys. Res. A* 689, 87 (2012). doi: 10.1016/j.nima.2012.05.079.
- [56] H. Tokuno *et al.* 'New air fluorescence detectors employed in the Telescope Array experiment'. In: *Nucl. Instrum. Methods Phys. Res. A* 676, 54 (2012). doi: 10.1016/j.nima.2012.02.044.
- [57] Pierre Auger collaboration (A. Aab *et al.*). 'Muons in air showers at the Pierre Auger Observatory: Mean number in highly inclined events'. In: *Phys. Rev. D* 91, 032003 (2015). doi: 10.1103/PhysRevD.91.032003.
- [58] Pierre Auger collaboration (A. Aab *et al.*). 'Testing hadronic interactions at ultrahigh energies with air showers measured by the Pierre Auger Observatory'. In: *Phys. Rev. Lett.* 117, 192001 (2016). doi: 10.1103/PhysRevLett.117.192001.
- [59] Pierre Auger collaboration (A. Aab *et al.*). 'Depth of maximum of air-shower profiles at the Pierre Auger Observatory. I. Measurements at energies above $10^{17.8}$ eV'. In: *Phys. Rev. D* 90, 122005 (2014). doi: 10.1103/PhysRevD.90.122005.
- [60] T. Abu-Zayyad *et al.* 'Evidence for changing of cosmic ray composition between 10^{17} and 10^{18} eV from multicomponent measurements'. In: *Phys. Rev. Lett.* 84, 4276 (2000). doi: 10.1103/PhysRevLett.84.4276.
- [61] H.P. Dembinski *et al.* (for the EAS-MSU, IceCube, KASCADE-Grande, NEVOD-DECOR, Pierre Auger, SUGAR, Telescope Array and Yakutsk EAS Array collaborations). 'Report on tests and measurements of hadronic interaction properties with air showers'. In: *Eur. Phys. J. Web Conf.* 210, 02004 (2019). doi: 10.1051/epjconf/201921002004.

- [62] L. Cazon (for the EAS-MSU, IceCube, KASCADE-Grande, NEVOD-DECOR, Pierre Auger, SUGAR, Telescope Array and Yakutsk EAS Array collaborations). 'Working group report on the combined analysis of muon density measurements from eight air shower experiments'. In: *Proc. Sci. ICRC2019*, 214 (2021). doi: 10.22323/1.358.0214.
- [63] D. Soldin (for the EAS-MSU, IceCube, KASCADE-Grande, NEVOD-DECOR, Pierre Auger, SUGAR, Telescope Array and Yakutsk EAS Array collaborations). 'Update on the combined analysis of muon measurements from nine air shower experiments'. In: *Proc. Sci. ICRC2021*, 349 (2022). doi: 10.22323/1.395.0349.
- [64] F. Gesualdi *et al.* 'On the muon scale of air showers and its application to the AGASA data'. In: *Proc. Sci. ICRC2021*, 473 (2022). doi: 10.22323/1.395.0473.
- [65] J.C. Arteaga-Velázquez. 'A report by the WHISP working group on the combined analysis of muon data at cosmic-ray energies above 1 PeV'. In: *Proc. Sci. ICRC2023*, 466 (2024). doi: 10.22323/1.444.0466.
- [66] L. Cazon *et al.* 'The muon measurements at Haverah Park and their connection to the Muon Puzzle'. In: *Proc. Sci. ICRC2023*, 431 (2024). doi: 10.22323/1.444.0431.
- [67] R. Ulrich, R. Engel and M. Unger. 'Hadronic multiparticle production at ultrahigh energies and extensive air showers'. In: *Phys. Rev. D* 83, 054026 (2011). doi: 10.1103/PhysRevD.83.054026.
- [68] Pierre Auger collaboration (A. Abdul Halim *et al.*). 'Testing hadronic-model predictions of depth of maximum of air-shower profiles and ground-particle signals using hybrid data of the Pierre Auger Observatory'. In: *Phys. Rev. D* 109, 102001 (2024). doi: 10.1103/PhysRevD.109.102001.
- [69] S. Baur *et al.* 'Core-corona effect in hadron collisions and muon production in air showers'. In: *Phys. Rev. D* 107, 094031 (2023). doi: 10.1103/PhysRevD.107.094031.
- [70] ALICE collaboration (J. Adam *et al.*). 'Enhanced production of multi-strange hadrons in high-multiplicity proton–proton collisions'. In: *Nat. Phys.* 13, 535 (2017). doi: 10.1038/nphys4111.
- [71] NA35 collaboration (J. Baechler *et al.*). 'Strangeness enhancement in central S + S collisions at 200 GeV/nucleon'. In: *Nucl. Phys. A* 525, 221c (1991). doi: 10.1016/0375-9474(91)90328-4.
- [72] WA97 collaboration (E. Andersen *et al.*). 'Strangeness enhancement at mid-rapidity in Pb–Pb collisions at 158 A GeV/c'. In: *Phys. Lett. B* 449, 401 (1999). doi: 10.1016/S0370-2693(99)00140-9.

- [73] E.V. Shuryak. ‘Quantum chromodynamics and the theory of superdense matter’. In: *Phys. Rep.* 61, 71 (1980). doi: 10.1016/0370-1573(80)90105-2.
- [74] P. Koch, B. Müller and J. Rafelski. ‘Strangeness in relativistic heavy ion collisions’. In: *Phys. Rep.* 142, 167 (1986). doi: 10.1016/0370-1573(86)90096-7.
- [75] ALICE collaboration (B. Abelev *et al.*). ‘Multiplicity dependence of pion, kaon, proton and Lambda production in p -Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV’. In: *Phys. Lett. B* 728, 25 (2014). doi: 10.1016/j.physletb.2013.11.020.
- [76] ALICE collaboration (J. Adam *et al.*). ‘Multi-strange baryon production in p -Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV’. In: *Phys. Lett. B* 758, 389 (2016). doi: 10.1016/j.physletb.2016.05.027.
- [77] ALICE collaboration (B. Abelev *et al.*). ‘Multi-strange baryon production at mid-rapidity in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV’. In: *Phys. Lett. B* 728, 216 (2014). doi: 10.1016/j.physletb.2013.11.048.
- [78] M. Vasileiou (for the ALICE collaboration). ‘Strangeness production with ALICE at the LHC’. In: *Physica Scripta* 95, 064007 (2020). doi: 10.1088/1402-4896/ab85fc.
- [79] LHCb collaboration (R. Aaij *et al.*). ‘Observation of strangeness enhancement with charmed mesons in high-multiplicity p Pb collisions at $\sqrt{s_{NN}} = 8.16$ TeV’. In: *Phys. Rev. D* 110, L031105 (2024). doi: 10.1103/PhysRevD.110.L031105.
- [80] LHCb collaboration (R. Aaij *et al.*). ‘Evidence for modification of b quark hadronization in high-multiplicity pp collisions at $\sqrt{s} = 13$ TeV’. In: *Phys. Rev. Lett.* 131, 061901 (2023). doi: 10.1103/PhysRevLett.131.061901.
- [81] B. Andersson *et al.* ‘Parton fragmentation and string dynamics’. In: *Phys. Rep.* 97, 31 (1983). doi: 10.1016/0370-1573(83)90080-7.
- [82] B. Andersson. *The Lund model*. Cambridge University Press, 1998.
- [83] A. Ortiz Velasquez *et al.* ‘Color reconnection and flowlike patterns in pp collisions’. In: *Phys. Rev. Lett.* 111, 042001 (2013). doi: 10.1103/PhysRevLett.111.042001.
- [84] C. Bierlich and J.R. Christiansen. ‘Effects of color reconnection on hadron flavor observables’. In: *Phys. Rev. D* 92, 094010 (2015). doi: 10.1103/PhysRevD.92.094010.
- [85] C. Bierlich *et al.* ‘Effects of overlapping strings in pp collisions’. In: *J. High Energy Phys.* 03, 148 (2015). doi: 10.1007/JHEP03(2015)148.

- [86] C. Bierlich, G. Gustafson and L. Lönnblad. ‘Collectivity without plasma in hadronic collisions’. In: *Phys. Lett. B* 779, 58 (2018). doi: 10.1016/j.physletb.2018.01.069.
- [87] C. Bierlich *et al.* ‘Strangeness enhancement across collision systems without a plasma’. In: *Phys. Lett. B* 835, 137571 (2022). doi: 10.1016/j.physletb.2022.137571.
- [88] CMS collaboration (A.M. Sirunyan *et al.*). ‘Measurement of the average very forward energy as a function of the track multiplicity at central pseudorapidities in proton-proton collisions at $\sqrt{s} = 13$ TeV’. In: *Eur. Phys. J. C* 79, 893 (2019). doi: 10.1140/epjc/s10052-019-7402-3.
- [89] L. Evans and P. Bryant. ‘LHC machine’. In: *J. Instrum.* 3, S08001 (2008). doi: 10.1088/1748-0221/3/08/S08001.
- [90] O. Brüning *et al.* ‘70 years at the high-energy frontier with the CERN accelerator complex’. In: *Nat. Rev. Phys.* 6, 628 (2024). doi: 10.1038/s42254-024-00758-5.
- [91] ALICE collaboration (K. Aamodt *et al.*). ‘The ALICE experiment at the CERN LHC’. In: *J. Instrum.* 3, S08002 (2008). doi: 10.1088/1748-0221/3/08/S08002.
- [92] ALICE collaboration (B. Abelev *et al.*). ‘Performance of the ALICE experiment at the CERN LHC’. In: *Int. J. Mod. Phys. A* 29, 1430044 (2014). doi: 10.1142/S0217751X14300440.
- [93] ATLAS collaboration (G. Aad *et al.*). ‘The ATLAS experiment at the CERN Large Hadron Collider’. In: *J. Instrum.* 3, S08003 (2008). doi: 10.1088/1748-0221/3/08/S08003.
- [94] ATLAS collaboration (G. Aad *et al.*). ‘Performance of the ATLAS detector using first collision data’. In: *J. High Energy Phys.* 09, 056 (2010). doi: 10.1007/JHEP09(2010)056.
- [95] CMS collaboration (S. Chatrchyan *et al.*). ‘The CMS experiment at the CERN LHC’. In: *J. Instrum.* 3, S08004 (2008). doi: 10.1088/1748-0221/3/08/S08004.
- [96] CMS collaboration (V. Khachatryan *et al.*). ‘CMS tracking performance results from early LHC operation’. In: *Eur. Phys. J. C* 70, 1165 (2010). doi: 10.1140/epjc/s10052-010-1491-3.
- [97] ALICE collaboration (S. Acharya *et al.*). ‘Transverse momentum spectra and nuclear modification factors of charged particles in Xe–Xe collisions at $\sqrt{s_{NN}} = 5.44$ TeV’. In: *Phys. Lett. B* 788, 166 (2019). doi: 10.1016/j.physletb.2018.10.052.

-
- [98] B.A. Kniehl *et al.* 'Heavy quark pair production in gluon fusion at next-to-next-to-leading $O(\alpha_s^4)$ order: One-loop squared contributions'. In: *Phys. Rev. D* 78, 094013 (2008). doi: 10.1103/PhysRevD.78.094013.
- [99] LHCb collaboration (R. Antunes Nobrega *et al.*). *LHCb technical design report: Reoptimized detector design and performance*. 2003. CERN-LHCC-2003-030.
- [100] R. Aaij *et al.* 'A comprehensive real-time analysis model at the LHCb experiment'. In: *J. Instrum.* 14, P04006 (2019). doi: 10.1088/1748-0221/14/04/P04006.
- [101] LHCb collaboration (R. Aaij *et al.*). 'Design and performance of the LHCb trigger and full real-time reconstruction in Run 2 of the LHC'. In: *J. Instrum.* 14, P04013 (2019). doi: 10.1088/1748-0221/14/04/P04013.
- [102] LHCb collaboration (R. Aaij *et al.*). 'Precision luminosity measurements at LHCb'. In: *J. Instrum.* 9, P12005 (2014). doi: 10.1088/1748-0221/9/12/P12005.
- [103] LHCb collaboration (R. Aaij *et al.*). 'The LHCb Upgrade I'. In: *J. Instrum.* 19, P05065 (2024). doi: 10.1088/1748-0221/19/05/P05065.
- [104] LHCb collaboration. *HLT1 trigger efficiencies in 2024 data*. 2024. LHCb-FIGURE-2024-030.
- [105] LHCb collaboration (I. Bediaga *et al.*). *LHCb SMOG upgrade*. 2019. CERN-LHCC-2019-005.
- [106] R. Aaij *et al.* 'Performance of the LHCb Vertex Locator'. In: *J. Instrum.* 9, P09007 (2014). doi: 10.1088/1748-0221/9/09/P09007.
- [107] G.A. Cowan. 'Performance of the LHCb Silicon Tracker'. In: *Nucl. Instrum. Methods Phys. Res. A* 699, 156 (2013). doi: 10.1016/j.nima.2012.05.074.
- [108] P. d'Argent *et al.* 'Improved performance of the LHCb Outer Tracker in LHC Run 2'. In: *J. Instrum.* 12, P11016 (2017). doi: 10.1088/1748-0221/12/11/P11016.
- [109] M. De Cian *et al.* *Fast neural-net based fake track rejection in the LHCb reconstruction*. 2017. LHCb-PUB-2017-011.
- [110] M. Adinolfi *et al.* 'Performance of the LHCb RICH detector at the LHC'. In: *Eur. Phys. J. C* 73, 2431 (2013). doi: 10.1140/epjc/s10052-013-2431-9.
- [111] C. Abellán Beteta *et al.* *Calibration and performance of the LHCb calorimeters in Run 1 and 2 at the LHC*. 2020. arXiv: 2008.11556 [physics.ins-det].
- [112] A.A. Alves Jr. *et al.* 'Performance of the LHCb muon system'. In: *J. Instrum.* 8, P02022 (2013). doi: 10.1088/1748-0221/8/02/P02022.

- [113] F. Archilli *et al.* ‘Performance of the muon identification at LHCb’. In: *J. Instrum.* 8, P10020 (2013). doi: 10.1088/1748-0221/8/10/P10020.
- [114] TOTEM collaboration (G. Anelli *et al.*). ‘The TOTEM experiment at the CERN Large Hadron Collider’. In: *J. Instrum.* 3, S08007 (2008). doi: 10.1088/1748-0221/3/08/S08007.
- [115] CMS and TOTEM collaborations (S. Chatrchyan *et al.*). ‘Measurement of pseudorapidity distributions of charged particles in proton–proton collisions at $\sqrt{s} = 8$ TeV by the CMS and TOTEM experiments’. In: *Eur. Phys. J. C* 74, 3053 (2014). doi: 10.1140/epjc/s10052-014-3053-6.
- [116] ALICE collaboration (S. Acharya *et al.*). ‘Charged-particle multiplicity distributions over a wide pseudorapidity range in proton-proton collisions at $\sqrt{s} = 0.9, 7,$ and 8 TeV’. In: *Eur. Phys. J. C* 77, 852 (2017). doi: 10.1140/epjc/s10052-017-5412-6.
- [117] LHCb collaboration (R. Aaij *et al.*). ‘Measurement of the nuclear modification factor and prompt charged particle production in p -Pb and pp collisions at $\sqrt{s_{NN}} = 5$ TeV’. In: *Phys. Rev. Lett.* 128, 142004 (2022). doi: 10.1103/PhysRevLett.128.142004.
- [118] T. Sjöstrand, S. Mrenna and P. Skands. ‘Pythia 6.4 physics and manual’. In: *J. High Energy Phys.* 05, 026 (2006). doi: 10.1088/1126-6708/2006/05/026.
- [119] T. Sjöstrand, S. Mrenna and P. Skands. ‘A brief introduction to Pythia 8.1’. In: *Comput. Phys. Commun.* 178, 852 (2008). doi: 10.1016/j.cpc.2008.01.036.
- [120] Geant4 collaboration (S. Agostinelli *et al.*). ‘Geant4—A simulation toolkit’. In: *Nucl. Instrum. Methods Phys. Res. A* 506, 250 (2003). doi: 10.1016/S0168-9002(03)01368-8.
- [121] J. Allison *et al.* ‘Geant4 developments and applications’. In: *IEEE Trans. Nucl. Sci.* 53, 270 (2006). doi: 10.1109/TNS.2006.869826.
- [122] ALICE collaboration (S. Acharya *et al.*). *The ALICE definition of primary particles*. 2017. ALICE-PUBLIC-2017-005.
- [123] E. Rodrigues. *Dealing with clones in the tracking*. 2006. LHCb-2006-057.
- [124] M. Needham. *Clone track identification using the Kullback-Leibler distance*. 2008. LHCb-2008-002.
- [125] M. Alexander *et al.* ‘Mapping the material in the LHCb Vertex Locator using secondary hadronic interactions’. In: *J. Instrum.* 13, P06008 (2018). doi: 10.1088/1748-0221/13/06/P06008.
- [126] R. Barlow. *Systematic errors: Facts and fictions*. 2002. arXiv: hep-ex/0207026.

-
- [127] F.N. Fritsch and R.E. Carlson. 'Monotone piecewise cubic interpolation'. In: *SIAM J. Numer. Anal.* 17, 238 (1980). doi: 10.1137/0717021.
- [128] M. Needham and T. Ruf. *Estimation of the material budget of the LHCb detector*. 2007. LHCb-2007-025.
- [129] LHCb collaboration (R. Aaij *et al.*). 'Measurement of the track reconstruction efficiency at LHCb'. In: *J. Instrum.* 10, P02007 (2015). doi: 10.1088/1748-0221/10/02/P02007.
- [130] LHCb collaboration (R. Aaij *et al.*). 'Measurement of prompt hadron production ratios in pp collisions at $\sqrt{s} = 0.9$ and 7 TeV'. In: *Eur. Phys. J. C* 72, 2168 (2012). doi: 10.1140/epjc/s10052-012-2168-x.
- [131] H.P. Dembinski and M. Roth. 'An algorithm for automatic unfolding of one-dimensional data distributions'. In: *Nucl. Instrum. Methods Phys. Res. A* 729, 410 (2013). doi: 10.1016/j.nima.2013.07.060.
- [132] R. Ulrich, T. Pierog and C. Baus. *Cosmic ray Monte Carlo package, CRMC*. Zenodo, 2021. doi: 10.5281/zenodo.5243165.
- [133] ATLAS collaboration (M. Aaboud *et al.*). 'Measurement of the inelastic proton-proton cross section at $\sqrt{s} = 13$ TeV with the ATLAS detector at the LHC'. In: *Phys. Rev. Lett.* 117, 182002 (2016). doi: 10.1103/PhysRevLett.117.182002.
- [134] LHCb collaboration (R. Aaij *et al.*). 'Measurement of the inelastic pp cross-section at a centre-of-mass energy of 13 TeV'. In: *J. High Energy Phys.* 06, 100 (2018). doi: 10.1007/JHEP06(2018)100.
- [135] TOTEM collaboration (G. Antchev *et al.*). 'First measurement of elastic, inelastic and total cross-section at $\sqrt{s} = 13$ TeV by TOTEM and overview of cross-section data at LHC energies'. In: *Eur. Phys. J. C* 79, 103 (2019). doi: 10.1140/epjc/s10052-019-6567-0.
- [136] R. Barlow and C. Beeston. 'Fitting using finite Monte Carlo samples'. In: *Comput. Phys. Commun.* 77, 219 (1993). doi: 10.1016/0010-4655(93)90005-w.
- [137] H. Dembinski and A. Abdeltoteleb. 'A new maximum-likelihood method for template fits'. In: *Eur. Phys. J. C* 82, 1043 (2022). doi: 10.1140/epjc/s10052-022-11019-z.
- [138] LHCb collaboration (R. Aaij *et al.*). 'Measurement of antiproton production in p -He collisions at $\sqrt{s_{NN}} = 110$ GeV'. In: *Phys. Rev. Lett.* 121, 222001 (2018). doi: 10.1103/PhysRevLett.121.222001.

- [139] LHCb collaboration (R. Aaij *et al.*). ‘Measurement of antiproton production from antihyperon decays in p He collisions at $\sqrt{s_{\text{NN}}} = 110$ GeV’. In: *Eur. Phys. J. C* 83, 543 (2023). doi: 10.1140/epjc/s10052-023-11673-x.
- [140] L. Anderlini *et al.* *The PIDCalib package*. 2016. LHCb-PUB-2016-021.
- [141] G. Graziani *et al.* ‘A neural-network-defined Gaussian mixture model for particle identification applied to the LHCb fixed-target programme’. In: *J. Instrum.* 17, P02018 (2022). doi: 10.1088/1748-0221/17/02/P02018.
- [142] T. Chen and C. Guestrin. *XGBoost: A scalable tree boosting system*. 2016. arXiv: 1603.02754 [cs.LG].
- [143] R. Barlow. ‘Event classification using weighting methods’. In: *J. Comput. Phys.* 72, 202 (1987). doi: 10.1016/0021-9991(87)90078-7.
- [144] M. Pivk and F.R. Le Diberder. ‘ $sPlot$: A statistical tool to unfold data distributions’. In: *Nucl. Instrum. Methods Phys. Res. A* 555, 356 (2005). doi: 10.1016/j.nima.2005.08.106.
- [145] H. Dembinski *et al.* ‘Custom orthogonal weight functions (COWs) for event classification’. In: *Nucl. Instrum. Methods Phys. Res. A* 1040, 167270 (2022). doi: 10.1016/j.nima.2022.167270.
- [146] T.B. Berrett and R.J. Samworth. ‘USP: An independence test that improves on Pearson’s chi-squared and the G -test’. In: *Proc. Royal Soc. A* 477, 20210549 (2021). doi: 10.1098/rspa.2021.0549.
- [147] M. Abadi *et al.* *TensorFlow: Large-scale machine learning on heterogeneous distributed systems*. 2016. arXiv: 1603.04467 [cs.DC].
- [148] S. Baker and R.D. Cousins. ‘Clarification of the use of chi-square and likelihood functions in fits to histograms’. In: *Nucl. Instrum. Methods Phys. Res.* 221, 437 (1984). doi: 10.1016/0167-5087(84)90016-4.
- [149] ATLAS collaboration (M. Aaboud *et al.*). ‘Measurement of the total cross section from elastic scattering in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector’. In: *Phys. Lett. B* 761, 158 (2016). doi: 10.1016/j.physletb.2016.08.020.
- [150] TOTEM collaboration (G. Antchev *et al.*). ‘Luminosity-independent measurement of the proton-proton total cross section at $\sqrt{s} = 8$ TeV’. In: *Phys. Rev. Lett.* 111, 012001 (2013). doi: 10.1103/PhysRevLett.111.012001.
- [151] J.L. Albacete *et al.* ‘Predictions for cold nuclear matter effects in $p + \text{Pb}$ collisions at $\sqrt{s_{\text{NN}}} = 8.16$ TeV’. In: *Nucl. Phys. A* 972, 18 (2018). doi: 10.1016/j.nuclphysa.2017.11.015.

-
- [152] J. Altmann *et al.* ‘QCD challenges from pp to AA collisions: 4th edition’. In: *Eur. Phys. J. C* 84, 421 (2024). doi: 10.1140/epjc/s10052-024-12650-8.
- [153] N. Armesto. ‘Nuclear shadowing’. In: *J. Phys. G* 32, R367 (2006). doi: 10.1088/0954-3899/32/11/R01.
- [154] J.W. Cronin *et al.* ‘Production of hadrons at large transverse momentum at 200, 300, and 400 GeV’. In: *Phys. Rev. D* 11, 3105 (1975). doi: 10.1103/PhysRevD.11.3105.
- [155] J.L. Nagle and W.A. Zajc. ‘Small system collectivity in relativistic hadronic and nuclear collisions’. In: *Annu. Rev. Nucl. Part. Sci.* 68, 211 (2018). doi: 10.1146/annurev-nucl-101916-123209.
- [156] E. Maguire, L. Heinrich and G. Watt. ‘HEPData: A repository for high energy physics data’. In: *J. Phys. Conf. Ser.* 898, 102006 (2017). doi: 10.1088/1742-6596/898/10/102006.
- [157] C. Bierlich *et al.* ‘Robust independent validation of experiment and theory: RIVET Version 3’. In: *SciPost Phys.* 8, 026 (2020). doi: 10.21468/SciPostPhys.8.2.026.
- [158] C. Bierlich *et al.* ‘Robust independent validation of experiment and theory: RIVET Version 4 release note’. In: *SciPost Phys. Codeb.* 36 (2024). doi: 10.21468/SciPostPhysCodeb.36.
- [159] LHCb collaboration (R. Aaij *et al.*). ‘Measurement of the b -quark production cross section in 7 and 13 TeV pp collisions’. In: *Phys. Rev. Lett.* 118, 052002 (2017). doi: 10.1103/PhysRevLett.118.052002.
- [160] LHCb collaboration (R. Aaij *et al.*). ‘Measurement of b hadron fractions in 13 TeV pp collisions’. In: *Phys. Rev. D* 100, 031102(R) (2019). doi: 10.1103/PhysRevD.100.031102.
- [161] LHCb collaboration (R. Aaij *et al.*). ‘Measurement of B^+ , B^0 and Λ_b^0 production in p Pb collisions at $\sqrt{s_{NN}} = 8.16$ TeV’. In: *Phys. Rev. D* 99, 052011 (2019). doi: 10.1103/PhysRevD.99.052011.
- [162] LHCb collaboration (R. Aaij *et al.*). ‘Measurement of the prompt D^0 nuclear modification factor in p -Pb collisions at $\sqrt{s_{NN}} = 8.16$ TeV’. In: *Phys. Rev. Lett.* 131, 102301 (2023). doi: 10.1103/PhysRevLett.131.102301.
- [163] LHCb collaboration (R. Aaij *et al.*). ‘Study of J/ψ production and cold nuclear matter effects in p Pb collisions at $\sqrt{s_{NN}} = 5$ TeV’. In: *J. High Energy Phys.* 02, 072 (2014). doi: 10.1007/JHEP02(2014)072.
- [164] LHCb collaboration (R. Aaij *et al.*). *Measurement of b hadron fractions in 13 TeV pp collisions*. HEPData, 2024. doi: 10.17182/hepdata.141536.

- [165] LHCb collaboration (R. Aaij *et al.*). *Measurement of B^+ , B^0 and Λ_b^0 production in pPb collisions at $\sqrt{s_{NN}} = 8.16$ TeV*. HEPData, 2024. doi: 10.17182/hepdata.153895.
- [166] LHCb collaboration (R. Aaij *et al.*). *Measurement of the prompt D^0 nuclear modification factor in p-Pb collisions at $\sqrt{s_{NN}} = 8.16$ TeV*. HEPData, 2025. doi: 10.17182/hepdata.153894.
- [167] J. Boelhaue, G. Daniels and L. Kolk. *RIVET plug-in for 'Measurement of the b-quark production cross section in 7 and 13 TeV pp collisions'*. 2025. LHCb-2016-I1504058.
- [168] J. Boelhaue, J. Friedman and L. Kolk. *RIVET plug-in for 'Measurement of B^+ , B^0 and Λ_b^0 production in pPb collisions at $\sqrt{s_{NN}} = 8.16$ TeV'*. 2025. LHCb-2019-I1720413.
- [169] J. Boelhaue, J. Friedman and L. Kolk. *RIVET plug-in for 'Measurement of the prompt D^0 nuclear modification factor in p-Pb collisions at $\sqrt{s_{NN}} = 8.16$ TeV'*. 2025. LHCb-2022-I2694685.
- [170] J. Boelhaue, G. Daniels and L. Kolk. *RIVET plug-in for 'Study of J/ψ production and cold nuclear matter effects in pPb collisions at $\sqrt{s_{NN}} = 5$ TeV'*. 2025. LHCb-2013-I1251899.
- [171] J. Boelhaue, G. Daniels and L. Kolk. *RIVET plug-in for 'Measurement of the nuclear modification factor and prompt charged particle production in p-Pb and pp collisions at $\sqrt{s_{NN}} = 5$ TeV'*. 2025. LHCb-2021-I1913240.
- [172] LHCb collaboration (R. Aaij *et al.*). *Measurement of prompt charged-particle production in pp collisions at $\sqrt{s} = 13$ TeV*. HEPData, 2023. doi: 10.17182/hepdata.136099.
- [173] L. Kolk. *RIVET plug-in for 'Measurement of prompt charged-particle production in pp collisions at $\sqrt{s} = 13$ TeV'*. 2023. LHCb-2021-I1889335.
- [174] D. Maurin *et al.* 'A cosmic-ray database update: CRDB v4.1'. In: *Eur. Phys. J. C* 83, 971 (2023). doi: 10.1140/epjc/s10052-023-12092-8.
- [175] J. Albrecht *et al.* *Global tuning of event generators with collider and astroparticle data. To be published.*
- [176] S. Baur *et al.* 'Combined analysis of accelerator and ultra-high energy cosmic ray data'. In: *Proc. Sci. ICRC2015*, 418 (2016). doi: 10.22323/1.236.0418.
- [177] A. Karneyeu *et al.* 'MCPlots: A particle physics resource based on volunteer computing'. In: *Eur. Phys. J. C* 74, 2714 (2014). doi: 10.1140/epjc/s10052-014-2714-9.

-
- [178] N. Korneeva, A. Karneyeu and P. Skands. ‘Event-generator validation with MCPlots and LHC@home’. In: *Eur. Phys. J. Plus* 139, 653 (2024). DOI: 10.1140/epjp/s13360-024-05353-2.
- [179] A. Buckley *et al.* ‘Systematic event generator tuning for the LHC’. In: *Eur. Phys. J. C* 65, 331 (2010). DOI: 10.1140/epjc/s10052-009-1196-7.
- [180] M. Krishnamoorthy *et al.* *Apprentice for event generator tuning*. 2021. arXiv: 2103.05748 [hep-ex].
- [181] O. Schulz *et al.* ‘BAT.jl: A Julia-based tool for Bayesian inference’. In: *Springer Nat. Comput. Sci.* 2, 210 (2021). DOI: 10.1007/s42979-021-00626-4.
- [182] S. La Cagnina *et al.* ‘A Bayesian tune of the Herwig Monte Carlo event generator’. In: *J. Instrum.* 18, P10033 (2023). DOI: 10.1088/1748-0221/18/10/P10033.
- [183] T. Huege (for the Pierre Auger collaboration). ‘The radio detector of the Pierre Auger Observatory – Status and expected performance’. In: *Eur. Phys. J. Web Conf.* 283, 06002 (2023). DOI: 10.1051/epjconf/202328306002.

Acknowledgements

Many people have accompanied and supported me during the work on the topics of my thesis. I would like to thank them in the following.

Firstly, my sincere thanks go to Prof. Johannes Albrecht for giving me the opportunity to work in his group. I had already been part of the group for my bachelor's and master's theses, in which I familiarised with rare decays at LHCb. I liked the idea of switching the topic to hadron production with an astroparticle connection. Johannes, you granted me a lot of freedom in my research, but nevertheless reminded me where to place the focus, many thanks! I am also grateful that I was able to stay at CERN to help with the assembly of the SciFi detector, which was a unique experience.

I am glad that Priv.-Doz. Dominik Elsässer agreed to serve as the second reviewer of my thesis, thanks! I enjoyed your lecture on astroparticle physics during my studies, which ends up fitting well to my thesis topics.

During a part of my research, I was funded by the SFB 1491 collaborative research centre of the German Research Foundation, which I appreciate much.

Then, I would like to thank Dr. Hans Dembinski, who was my direct supervisor. We were the precursors of the Astro-QCD subgroup and became a good team during the years. Many thanks, Hans, for your work on statistical methods, from which I was able to benefit in the analysis, and for all our discussions! I also learned a lot concerning programming from you.

It was a pleasure to collaborate with Prof. Michael Schmelling. I liked the refreshing exchange about analysis techniques in our weekly meetings; thank you, Michael! As you know, your insights about physics always fascinated me.

Two incidents overshadowed my time as a doctoral student, unfortunately. The COVID-19 pandemic forced us to stay in home office for several months and limited all communication to virtual means. The more severe occurrence was that Prof. Bernhard Spaan, who had led the group for many years, suddenly passed away end of 2021. I still remember his lecture on experimental particle physics, which further strengthened my interest in the field, and many of his funny anecdotes.

I am grateful that I was working in a group of many focused people. Throughout the years, I was sharing the office with various colleagues: Alex (Seuthe), Maik, Jannis, Nicole, Alex (Battig), Luca (Toscano), Nils, Michelle and Noah. I always enjoyed the atmosphere in our office; we were able to work in a quiet and concentrated way, but time for joking was nevertheless there. Thanks also to Vukan. It was

Acknowledgements

fun to work on HEPData and RIVET together with Lars although we needed to continue many discussions at lunch. Exchanging knowledge about LHCb and event generators with Felix was a pleasure as well.

I would like to thank everyone who took the time to read through chapters of my thesis and provided comments.

Special thanks go to Alex (Ratke) as former colleague and close friend of mine. We spent a large part of the bachelor's, master's and doctoral studies together and always motivated each other when much work had to be done. It really was a great time!

Lastly, but most importantly, I need to thank my family and friends, especially my parents. Without your immeasurable support, I would not have come this far. I am thankful that everyone was constantly understanding when my time was again short. And I liked those moments of explaining what I am actually doing in my work... *Danke!*