Study of Ξ–Hadron Correlations in pp Collisions at $\sqrt{s}$ = 13 TeV Using the ALICE Detector

JONATAN ADOLFSSON
DEPARTMENT OF PHYSICS | LUND UNIVERSITY | 2020

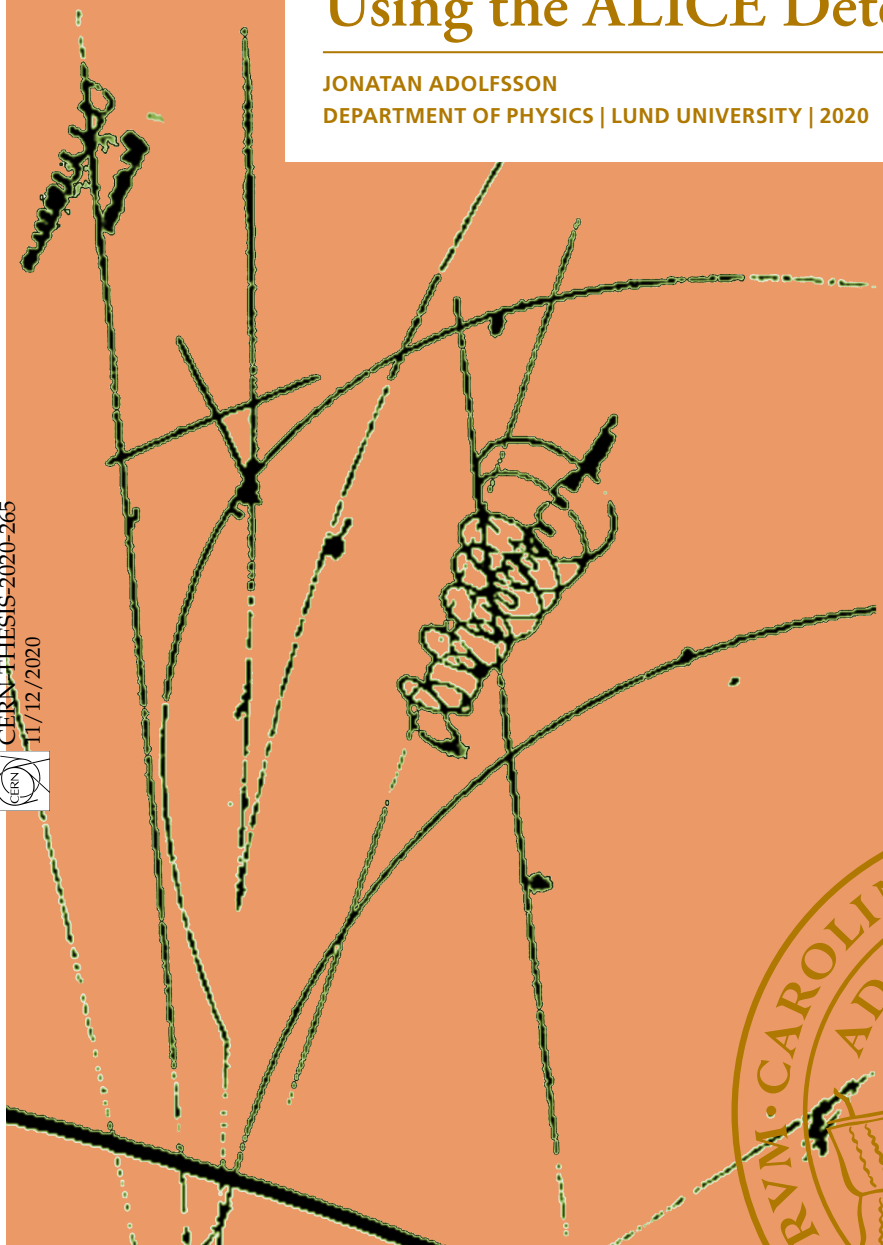

CERN-THESIS-2020-265
11/12/2020

Printed by Media-Tryck, Lund 2020 NORDIC SWAN ECOLABEL 3041 0903

Department of Physics
Particle Physics

ISBN 978-91-7895-604-3

Study of $\Xi-$Hadron Correlations in pp Collisions at $\sqrt{s} = 13$ TeV Using the ALICE Detector

by Jonatan Adolfsson

LUND
UNIVERSITY

Thesis for the degree of Doctor of Philosophy
Thesis advisors: Dr. David Silvermyr, Prof. Peter Christiansen,
Prof. Anders Oskarsson, Dr. Alice Ohlson
Faculty opponent: Dr. Prof. Nu Xu

To be presented, with the permission of the Faculty of Science of Lund University, for public criticism in the Rydberg lecture hall (Rydbergsalen) at the Department of Physics on Friday, the 11th of December 2020 at 13:15.

Organization
LUND UNIVERSITY

Department of Physics
Box 118
S–221 00 LUND
Sweden

Author(s)
Jonatan Adolfsson

Document name
DOCTORAL DISSERTATION

Date of disputation
2020-12-11

Sponsoring organization

Title and subtitle
Study of $\Xi-$Hadron Correlations in pp Collisions at $\sqrt{s} = 13$ TeV Using the ALICE Detector

Abstract
By colliding heavy nuclei at high energies, which is done at RHIC and the LHC, a strongly interacting Quark Gluon Plasma (QGP) is created. This manifests itself through several different signatures, which until recently was thought to uniquely probe the QGP. Recently, however, similar signatures have been observed also in small systems, such as pp collisions with high charged-particle multiplicity, which is quite puzzling since a QGP is not expected to be formed in such dilute systems with short lifetimes. One such observable is the enhanced relative yields of multistrange baryons, such as the Ξ baryon, which has been observed in e.g. Pb–Pb collisions. More recently, this yield enhancement has been observed to scale smoothly with multiplicity also in pp collisions.

The main analysis presented in this thesis aims at understanding the production mechanism of strange quarks in pp collisions at $\sqrt{s} = 13\,\text{TeV}$, and in this way reach an explanation of the origin of the observed strangeness enhancement. This is done by studying angular $\Xi - h$ correlations, where h is either of π, K, p, Λ, or Ξ hadrons, by using data from the ALICE detector. The results are compared with four phenomenological models; three flavours of the QCD inspired PYTHIA8 which is based on colour strings, and the core-corona model EPOS LHC. The PYTHIA tunes are the Monash tune, the Junction Mode 0 tune, and a yet unofficial tune with rope hadronisation, which is a proposed mechanism for the observed strangeness enhancement. In EPOS, this is modelled by an increasing fraction of a core that behaves like a medium.

The results show that the $\Xi - \pi$ correlation function is dominated by a narrow near-side peak. This is not present in any of the other correlations, which on the other hand have a wide extension in rapidity. This means that pions decouple later in the evolution from the Ξ baryon compared to the other species, likely within the jet, which was concluded to be due to charge balance, whereas the other correlations are attributed to strangeness and baryon decoupling. In all PYTHIA flavours, strong correlations within the jet are present for all combinations except $\Xi -$ p correlations, meaning that strangeness and baryon number are produced earlier in the evolution in data than in PYTHIA. For EPOS, on the other hand, the correlation function is very dilute for most species, which was concluded to be due to local conservation of quantum numbers not being properly accounted for. Therefore, it is not yet possible to use this measurement to test the underlying mechanism provided by this model. Based on this, the observations in data indicate that the strangeness production mechanism is likely either due to a core-corona like state, or some hybrid mechanism where string interactions are also important.

Key words
LHC, ALICE, quark gluon plasma, small systems, pp collisions, strangeness enhancement, correlations

Classification system and/or index terms (if any)

Supplementary bibliographical information

Language
English

ISSN and key title

ISBN
978-91-7895-604-3 (print)
978-91-7895-605-0 (pdf)

Recipient's notes

Number of pages
259

Price

Security classification

I, the undersigned, being the copyright owner of the abstract of the above-mentioned dissertation, hereby grant to all reference sources the permission to publish and disseminate the abstract of the above-mentioned dissertation.

Signature

Date 2020-11-06

DOKUMENTDATABLAD enl SIS 61 41 21

Study of $\Xi-$Hadron Correlations in pp Collisions at $\sqrt{s} = 13$ TeV Using the ALICE Detector

by Jonatan Adolfsson

LUND
UNIVERSITY

Cover illustration front: Decay of a Ξ^- baryon imaged by a bubble chamber. Original image credits: CERN.

© Jonatan Adolfsson 2020

Faculty of Science, Department of Physics

ISBN: 978-91-7895-604-3 (print)
ISBN: 978-91-7895-605-0 (pdf)

Printed in Sweden by Media-Tryck, Lund University, Lund 2020

Media-Tryck is a Nordic Swan Ecolabel certified provider of printed material. Read more about our environmental work at www.mediatryck.lu.se

MADE IN SWEDEN

Abstract

By colliding heavy nuclei at high energies, which is done at RHIC and the LHC, a strongly interacting Quark Gluon Plasma (QGP) is created. This manifests itself through several different signatures, which until recently was thought to uniquely probe the QGP. Recently, however, similar signatures have been observed also in small systems, such as pp collisions with high charged-particle multiplicity, which is quite puzzling since a QGP is not expected to be formed in such dilute systems with short lifetimes. One such observable is the enhanced relative yields of multistrange baryons, such as the Ξ baryon, which has been observed in e.g. Pb–Pb collisions. More recently, this yield enhancement has been observed to scale smoothly with multiplicity also in pp collisions.

The main analysis presented in this thesis aims at understanding the production mechanism of strange quarks in pp collisions at $\sqrt{s} = 13\,\text{TeV}$, and in this way reach an explanation of the origin of the strangeness enhancement observed there. This is done by studying angular $\Xi{-}h$ correlations, where h is either of π, K, p, Λ, or Ξ hadrons, by using data from the ALICE detector. The results are compared with four phenomenological models; three flavours of the QCD inspired PYTHIA8 which is based on colour strings, and the core-corona model EPOS LHC. The PYTHIA tunes are the Monash tune, the Junction Mode 0 tune, which has an additional mechanism for baryon formation, and a yet unofficial tune including rope hadronisation, which is a proposed mechanism for the observed strangeness enhancement. In EPOS, the enhanced strangeness is modelled by an increasing fraction of a core that behaves like a medium.

The results show that the $\Xi - \pi$ correlation function is dominated by a narrow near-side peak. This is not present in any of the other correlations, which on the other hand have a wide extension in rapidity. This means that pions decouple later in the evolution from the Ξ baryon compared to the other species, likely within the jet, which was concluded to be due to charge balance, whereas the other correlations are attributed to strangeness and baryon decoupling. In all PYTHIA flavours, strong correlations within the jet are present for all combinations except $\Xi - \text{p}$ correlations, meaning that strangeness and baryon number are produced earlier in the evolution in data than in PYTHIA. The junction model however gave a description of the $\Xi{-}$baryon correlation that was closer to data than the Monash tune, indicating that the additional baryon mechanism included there is more likely to be correct. For EPOS, on the other hand, the correlation function is very dilute for most species, which was concluded to be due to local conservation of quantum numbers not being properly accounted for. Therefore, it is not yet possible to use this measurement to test the underlying mechanism provided by this model. Based on this, the observations in data indicate that the strangeness production mechanism is likely either due to a core-corona like state, or some hybrid mechanism where string interactions are also important.

Correlations were also measured as a function of multiplicity, yielding very similar results across multiplicity classes. Therefore it was concluded that the strangeness and baryon production mechanisms in pp collisions are likely the same regardless of multiplicity.

Populärvetenskaplig sammanfattning

Att vatten övergår till ånga när man hettar upp det är ett välkänt fenomen, men även materia som bygger upp atomkärnor kommer att genomgå en fasövergång om den hettas upp tillräckligt mycket. Det som då bildas är ett kvark-gluonplasma (QGP), en exotisk form av materia där atomkärnans minsta beståndsdelar – kvarkar och gluoner – mer eller mindre rör sig fritt, vilket radikalt skiljer sig från vanlig materia. Dess egenskaper kan sammanfattas som ett mycket hett och tätt medium, som i det närmaste beter sig som en vätska. Detta tillstånd tros ha existerat under den första hundratusendelen av en sekund efter Big Bang, så genom att studera detta kan vi lära oss mer om hur universums byggstenar en gång bildades.

Genom att accelerera tunga joner, t.ex. blykärnor, till mycket höga energier och sedan kollidera dem, kan man skapa ett QGP i labbet. Detta görs bl.a. vid partikelfysiklaboratoriet CERN utanför Genève. Det QGP som då bildas existerar endast under ett extremt kort ögonblick, varefter det sönderfaller och en skur av partiklar bildas. Genom att studera dessa partiklar, kan vi lära oss om kvark-gluonplasmats egenskaper. Det har gjorts teoretiska förutsägelser om hur kvark-gluonplasmat förväntas påverka partiklarnas sammansättning och fördelning i rummet, vilka har visat sig stämma bra med det som observeras. På senare tid har det däremot gjorts en upptäckt som inte passar så bra in i bilden, då liknande effekter har observerats också i kollisioner mellan protoner, vilka är 200 gånger mindre än en blykärna. Enligt gällande modeller så borde inte ett QGP kunna bildas i så små kollisionssystem – det går helt enkelt för kort tid innan kollisionen är över. Som forskare vill vi förstås veta vad som egentligen händer i dessa kollisioner, vilket är syftet med den här avhandlingen. För att förstå vad jag har studerat, behöver man först förstå hur en atomkärna är uppbyggd.

En atomkärna består av protoner och neutroner, men dessa är inte elementarpartiklar, utan är i sin tur uppbyggda av kvarkar och gluoner. En proton består av tre kvarkar – två uppkvarkar och en nerkvark, vilka har olika laddning – som är sammanbundna av gluoner. En neutron består istället av en uppkvark och två nerkvarkar. Det finns också tyngre kvarkar, vilka inte är vanligt förekommande i universum då de snabbt sönderfaller till upp- eller nerkvarkar, men det hindrar inte dem från att bildas vid högenergetiska partikelkollisioner. Den lättaste (och därmed vanligaste) av dessa tyngre kvarkar är särkvarken, vilken är som en tyngre version av nerkvarken.

En förutsägelse av kvark-gluonplasmat är att fler särkvarkar bildas i detta medium än vid partikelkollisioner där inget QGP bildas. Och mycket riktigt bildas fler särkvarkar i bly-blykollisioner än vid proton-protonkollisioner, men det sker också en tydlig ökning från protonkollisioner där få partiklar har deltagit till dem där många partiklar har gjort det. Flera teoretiska modeller har utvecklats för att förklara detta, vilka kan delas in i två kategorier som sinsemellan är väldigt olika, så andra typer av mätningar måste göras för att testa vilken som stämmer bäst överens med verkligheten.

För att få en uppfattning om vad som händer i dessa kollisioner har jag studerat Ξ-baryonen (Xi), vilken består av två särkvarkar och en nerkvark. Denna rör sig några centimeter innan den sönderfaller i ett väldigt distinkt mönster, vilket relativt lätt kan detekteras i ALICE-detektorn, vars data har använts till den här studien. Genom att mäta avståndet i detektorn från Ξ-baryonen till andra partiklar som bildas i kollisionen (eller mer specifikt fördelningen av avståndet), och i synnerhet till sådana som innehåller

särkvarkar, kan vi få en uppfattning om när och var särkvarken har bildats under kollisionen och på så sätt testa modellerna.

Två modeller har testats (plus ytterligare några variationer av den ena), en från vardera kategorin. Den ena modellen är en så kallad kärn-koronamodell, vilken förutsäger att ett QGP bildas även i protonkollisioner, men att den relativa volymen ökar när fler partiklar deltar. Den andra modellen är en så kallad strängmodell, där observationerna förklaras genom att det sker andra processer som liknar det man skulle förvänta sig från ett QGP, utan att ett sådant har bildats. Resultaten från studien visar att ingen av modellerna beskriver verkligheten särskilt väl. Särkvarkarna verkar bildas tidigare i kollisionen än vad som förutsägs av strängmodellen, vilket tyder på att den underläggande fysiken i den modellen kan vara fel, även om det fortfarande är för tidigt att säga. För kärn-koronamodellen är problemet ett annat, då den i nuläget saknar viktiga bevarandelagar och därför ger helt felaktiga förutsägelser. Så för att testa om den modellen i grunden stämmer, behövs mer jobb från utvecklarnas sida.

Popular Summary

The fact that water transforms into water vapour when heated is a well-known phenomenon, but also the matter building up atomic nuclei will undergo a phase transition if enough heated. What then is formed is a Quark Gluon Plasma (QGP), an exotic state of matter where the smallest constituents of the atomic nucleus – quarks and gluons – more or less can move freely, which is radically different from normal matter. Its properties can be summarised as a very hot and dense medium, which behaves like a fluid. This state of matter is believed to have existed during the first one hundred thousandth of a second after Big Bang, so by studying this we may learn more about how the building blocks of the Universe once were formed.

By accelerating heavy ions, e.g. lead nuclei, to very high energy, and then collide them, we may create a QGP in the laboratory. This is for instance done at the particle physics laboratory CERN outside Geneva. The QGP that is then formed only exists for an extremely short instance of time, after which it decays into a shower of particles. By studying these particles, we may learn about the properties of the QGP. Theoretical predictions have been made of how the QGP is expected to affect the abundances and spatial distribution of the created particles, which have turned out to agree with what is observed. More recently, however, there has been a discovery that does not fit so well into this picture, which is that similar effects have been observed also in collisions between protons, despite being 200 times smaller than a lead nucleus. According to current models, a QGP is not expected to form in such small collision systems – the time frame of the collision is simply too short. As researchers, we naturally want to understand what is really going on in these collisions, which is the purpose of this thesis. To understand what I have studied, though, one first needs to understand how an atomic nucleus is composed.

An atomic nucleus consists of protons and neutrons, but these are not elementary particles, but are in turn composed of quarks and gluons. A proton consists of three quarks – two up quarks and one down quark, which have different electric charge – which are bound together by gluons. A neutron is instead composed of one up quark and two down quarks. There also exist heavier quarks, which are not very abundant in the Universe since they quickly decay into up or down quarks, but that does not prevent them from being created in high-energy particle collisions. The lightest (and hence most common) of these heavier quark is the strange quark, which is like a heavier version of the down quark.

One prediction of the QGP is that more strange quarks are produced in this medium than what would be the case for particle collisions without QGP formation. And as expected, more strange quarks are produced in lead-lead collisions than in proton-proton collisions, but there is also a clear enhancement in proton collisions with many produced particles, compared to those where only a few particles are detected. Several theoretical models have been developed aiming to explain this, which can be divided into two categories where the physics is very different, so other types of measurements are required to test which model agrees best with reality.

To get an idea of what happens in these collisions, I have studied the Ξ baryon ("Xi"), which consists of two strange quarks and one down quark. This typically moves a few centimetres before it decays through a very distinct pattern, which makes it

relatively easy to detect in the ALICE detector, where the data used in this study has been collected. By measuring the distance between the Ξ baryon and other particles that are produced in the collision (or more specifically the distribution of this distance), and in particular particles that contain strange quarks, we may gain some understanding of the time and position of the formation of strange quarks during the collision, and in this way test the models.

Two models have been tested (plus a few variations of one of them), one from either category. One of the models is a so-called core-corona model, which predicts that a QGP is formed also in proton collisions, but that the relative volume of it increases when more particles participate. The other model is a so-called string model, where the observations are explained by other processes that mimic what one expects from a QGP, but without actually forming such a state. The results from the study show that neither of the models describes the observations particularly well. The strange quarks seem to be formed earlier in the collision than what is predicted by the string model, which indicates that the underlying physics in that model may be incorrect, although it is still too early to tell for sure. The core-corona model has a different problem, as currently important conservation laws are not taken into account, resulting in predictions very far from reality. So in order to test whether the idea behind this model is correct, more work is required from the developers.

Personal Acknowledgements

These past months have been insane, and I would never have been able to reach this far without the aid of all people that have provided help and support along the way.

First, however, I would like to quote my personal goals for the PhD thesis, which I wrote down just at the start of my PhD in 2016: I wanted to "develop a new analysis method, preferably with the goal of gaining more information w.r.t. some parameter [...] in small systems". With this thesis, I have achieved exactly that, but I would never have come up with this research idea on my own. So thank you Peter for your great research proposal, and I hope my analysis and results met your expectations – they definitely met my own. I would also like to thank you for inspiring me into the field of particle physics seven years ago, and for your continuous supervision of the physics part of this work.

David, you have my gratitude for involving me into the chip testing campaign, tirelessly answering to my questions, to read through and provide feedback to every single section in this thesis, and to help with all complicated administrative stuff throughout the years. I am sure I have not been the easiest student to deal with all the time, but I am glad you managed.

Anders, thank you for all your great explanations, for giving me inspiration for the noise analysis, and for coming up with the idea for the delay measurement. All your feedback on the instrumentation part was really valuable as well.

Alice, you are probably one of the best teachers in physics analysis that I have met, and you always come up with the best ideas on how to improve things in my analysis. My results would never have become this good without your help. Thank you!

I would also like to thank the PhD students that have been part of the ALICE group throughout these years, you have all been very good friends. Special thanks to Vytautas for your well-written thesis and Adrian for your work with the trees. I would never have finished on time without your help! Thanks also to Oliver for preparing the China trip, Omar for keeping the working spirit up in the office, and Martin for the board game nights during my first year.

And thanks to Tuva for inpiring me to start a PhD in ALICE. Your thesis was really inspiring reading when I started!

I would also like to express my gratitude to Christian, as well as the rest of the PYTHIA team – Torbjörn, Gösta, Leif, and others – for providing such nice models to compare to. I might not believe they are correct, but very useful nonetheless.

Florido, thank you for keeping my computer in a good condition, and your fast response whenever I run into problems with it; Bozena, for helping with all the paper work and making sure that everything is in order with my employment; Else, for having a look at my thesis and taking care of some of the guidance for the ALICE PhD students; and Lennart for running the IT server.

I would like to thank all the ATLAS PhD students for being my friends throughout these years and hopefully many years to come: Sasha, Katja, Trine, Eric, Eva, Caterina, Eleni, Nathan, and Alex. We have had so much fun during after-works, division excursions, PhD defence parties, and hikes. I would also like to thank the ALICE Master students for bringing life to the office: Madeleine, Adam, Rikard, Maria, other Madeleine, Anna, Martin, and Lisa. Some of you have been really good friends as well!

Most of you will probably not read this, but I would also like to thank the members of FISK for giving me something other than work to enjoy and get absorbed with at times.

Last but not least, I would like to thank my parents for hosting me during the extraordinary times last spring and for always giving me a home to return to during breaks, so I can escape the city life. Mum, thanks for all your support throughout the years, and for believing in me!

Contents

I Analysis Motivation **1**

1 Introduction **3**
1.1 General Background . . . 3
1.2 Research Goals . . . 5
1.3 Outline of This Thesis . . . 6
1.4 Own Contributions . . . 7

2 Fundamental Theory **9**
2.1 The Standard Model of Particle Physics . . . 9
2.2 The Electroweak Interaction and Feynman Diagrams . . . 11
2.3 Quantum Chromodynamics . . . 14
2.4 The QCD Phase Diagram . . . 17
2.5 Properties of the Quark-Gluon Plasma . . . 19

3 Experimental Signatures of the QGP **21**
3.1 The Evolution of a Heavy-Ion Collision . . . 21
3.2 Collision Geometry . . . 23
3.2.1 The Glauber Model . . . 24
3.3 Measuring the Temperature: Thermal Photons . . . 25
3.4 Jet Quenching . . . 27
3.5 Heavy-Quarkonia Melting . . . 31
3.6 Strangeness Enhancement . . . 33
3.7 Collective Flow . . . 35
3.7.1 Radial Flow . . . 35
3.7.2 Anisotropic Flow . . . 37
3.8 Phenomenological Models . . . 42
3.8.1 Overview . . . 42
3.8.2 The Lund String Model . . . 43
3.8.3 The Angantyr Model . . . 46
3.8.4 The EPOS Model . . . 49
3.8.5 Successful Predictions by PYTHIA, Angantyr, and EPOS . . . 50

II Experimental Setup and Development **57**

4 The ALICE Detector **59**
4.1 Detector Overview . . . 59
4.2 Inner Tracking System . . . 61

4.3 V0 Detector . . . 64
4.4 Time Projection Chamber . . . 64
4.5 Time-Of-Flight Detector . . . 69
4.6 Tracking in the Central Barrel . . . 70
4.7 Triggering and Data Aquisition . . . 70

5 Chip Testing for the TPC **75**
5.1 SAMPA: a New Readout Chip Design . . . 75
5.2 Testing Overview . . . 78
5.3 Cross-Talk Measurements . . . 82
5.4 Noise Measurements . . . 84
5.5 Rise Time Measurements . . . 85
5.6 Gain Measurements . . . 88
5.7 Bit Issues: the Odd-Even Effect . . . 89

6 Noise Measurements for the MCH **91**
6.1 Introduction . . . 91
6.2 Correlation Between Measured Baseline and Noise Levels . . . 92
6.3 Models for Discretisation . . . 93
6.3.1 Ideal Model . . . 93
6.3.2 Taking Into Account the Odd-Even Effect . . . 94
6.4 Inverse . . . 97
6.5 Application to the Data Set . . . 99
6.5.1 Hypotheses . . . 99
6.5.2 Error estimate . . . 100
6.5.3 Validation . . . 100
6.6 Results . . . 100
6.7 Application to the Automatic Testing . . . 105
6.7.1 Motivation . . . 105
6.7.2 Measurement Details . . . 106
6.7.3 Results . . . 106
6.8 Discussion . . . 114
6.9 Conclusions . . . 116

III Main Analyses **117**

7 Measurement of Flow in Xe–Xe Collisions **119**
7.1 Introduction . . . 119
7.2 Analysis Methods . . . 119
7.3 Results . . . 120
7.4 Discussion . . . 123
7.5 Conclusions . . . 124

8 Analysis of Ξ-Hadron Correlations **125**
8.1 Introduction and Overview . . . 125
8.2 Rapidity Versus Pseudorapidity . . . 126
8.3 Measuring the Correlation Function . . . 127
8.4 Event Selection . . . 129

8.5 Particle Identification . . . 130
8.5.1 Identification of Long-Lived Particles (Direct Detection) . . . 130
8.5.2 Identification of V^0 Particles and Cascades . . . 133
8.6 Efficiency Corrections . . . 141
8.7 The Linear Algebra Method . . . 143
8.8 Feed-down corrections . . . 147
8.9 Monte Carlo Simulations . . . 150
8.9.1 Monte Carlo Closure Test . . . 150
8.9.2 Model Comparisons . . . 154
8.10 Systematic Uncertainties . . . 155

9 Correlation Results **163**
9.1 ALICE Results . . . 163
9.1.1 Multiplicity Dependent Results . . . 165
9.2 Simulation Results . . . 171
9.3 Comparisons Between Models and Data . . . 177

10 Discussion of Correlations and Conclusions **185**
10.1 Correlations Originating from the Underlying Event . . . 186
10.2 Correlations Originating from Ξ Baryon Interactions . . . 190
10.2.1 $\Xi - \pi$ Correlations . . . 191
10.2.2 $\Xi -$ K Correlations . . . 193
10.2.3 $\Xi -$ p Correlations . . . 195
10.2.4 $\Xi - \Lambda$ and $\Xi - \Xi$ Correlations . . . 197
10.3 Multiplicity Dependent Results . . . 198
10.4 Conclusions . . . 199
10.5 Outlook . . . 199

IV Appendices **203**

A Mathematical Derivations **205**
A.1 Derivation of Existence of a Local Inverse of the Distribution of the ADC Output . . . 205
A.2 Calculation of Λ Sideband Coefficients . . . 206

B List of Acronyms **209**

C List of Common Hadrons **213**

D Complementary Figures **215**
D.1 MCH Noise Measurements . . . 215
D.2 Ξ−Hadron Analysis Description . . . 220
D.3 Complementary Ξ-Hadron Correlation Results . . . 233

Bibliography **237**

Part I

Analysis Motivation

Chapter 1

Introduction

1.1 General Background

If colliding heavy nuclei which are accelerated to extremely high energies, it is now widely accepted that the nuclear matter will undergo a phase transition to a very hot (more than 100 000 times the temperature at the core of the Sun!) and dense state of matter, called the Quark Gluon Plasma (QGP). This has many exotic properties, since it is strongly interacting and yet behaves as the most perfect fluid that is known. In this state, the smallest substituents of nuclei, called quarks, which are normally confined into composite particles called hadrons[1] due to the strongly interacting gluons holding them together, are liberated – the quarks and gluons are *deconfined*. The QGP is not only produced in heavy-ion collisions though, but is thought to have existed during the first 10 µs or so after the Big Bang, and does likely exist in the cores of neutron stars. The latter have however much higher baryonic densities, so the matter created in heavy-ion collisions is more similar to the conditions shortly after Big Bang. And since this matter eventually evolved into the atoms and nuclei present in the Universe today, we can perhaps learn more about how it was formed by studying the QGP in the laboratory. This does however only exist for about 10^{-23} s, which is much more short-lived than anything we can detect, so how can we be so sure that a QGP indeed is formed in these collisions?

The answer to that question is that theoretical models have made many predictions of what the footprint of a QGP will look like in the detector, which will be described in detail in Chapter 3 of this thesis. For many years there had not been any conclusive evidence for finding such signatures, until the heavy-ion programme started at the Relativistic Heavy-Ion Collider (RHIC) at Brookhaven outside New York in 2000. This has used a rich collection of heavy-ion systems, with Au–Au collisions at 200 GeV per nucleon pair being the most dominant one. The data was collected at four detectors, but only STAR is still operational. While STAR has a main focus on collecting as many hadrons as possible, the other major experiment, PHENIX, had a larger focus on rare probes, making them complementary. PHENIX is currently being replaced by a new detector called sPHENIX, which will enable much larger detection rates.

[1] These include protons and neutrons, but there are many more.

From the RHIC data, many properties of the QGP were determined to a good precision, but several questions still remained. Some of these began to get their answer in 2010, when the heavy-ion programme started at the Large Hadron Collider (LHC), which enabled collisions at a much higher energy – it is currently at 5 TeV per nucleon pair. But instead new discoveries were made, which raised many new questions, which I will return to in a short while.

First, however, I will give a short introduction to the LHC. The LHC is located at CERN outside Geneva, and with a circumference of 27 km, it is not only the largest particle accelerator in the world, but the largest machine overall. It is located in a tunnel about 100 m below ground, and is outlined in the aerial photo shown in Fig. 1.1. For most of the year, LHC collides protons, but during a short period each year, typically either lead ions are collided or protons are collided with lead ions. There are four major experiments at the LHC – ATLAS, CMS, ALICE, and LHCb. Of these, ATLAS and CMS are general-purpose detectors aiming at finding rare processes, and in this way pursuing the frontiers of fundamental physics, whereas LHCb is aimed at flavour physics, searching for exotic states of hadrons. While all experiments participate in the heavy-ion programme, ALICE is the only dedicated heavy-ion detector. This is the experiment where I have been working.

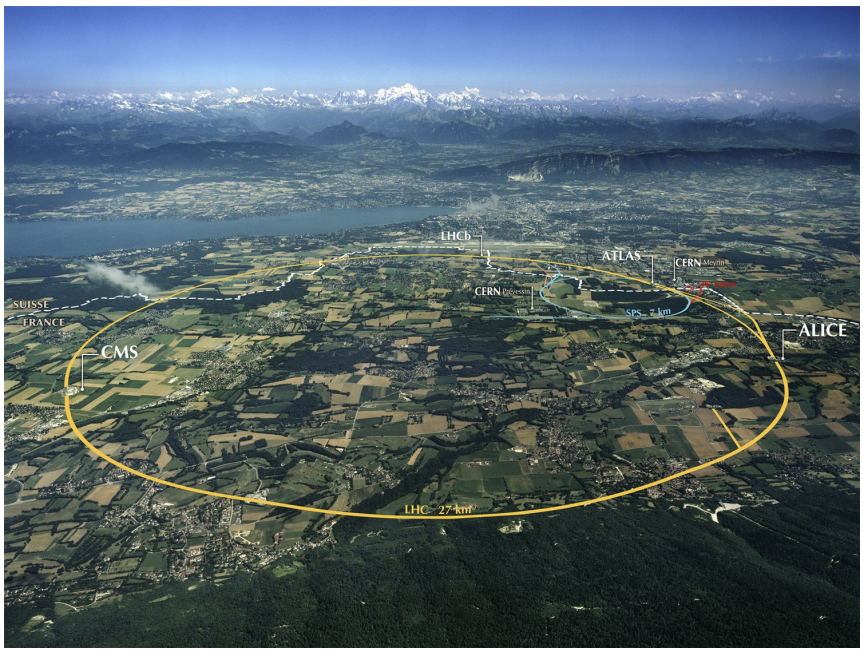

Figure 1.1 – Aerial view of CERN, where the major accelerators SPS and the LHC are outlined. The main offices are at the Meyrin site. Figure taken from Ref. [1].

The main tracking device of the ALICE detector is the Time Projection Chamber (TPC), which while its detection rates are quite limited, it is able to record most of the charged particles traversing it to a very precise position, and identify them – which is not possible at this scale using any other detector technology. This is possible even for

a heavy-ion collision with a huge amount of tracks, as is shown in the event display in Fig. 1.2. To understand what happens in a heavy-ion collision, we need to collect as much information about the event as possible, and for this the TPC is ideal. Therefore a TPC is also used by STAR, and will be used in sPHENIX as well. Moreover, the ALICE TPC is currently being upgraded to enable much higher detection rates, which will make it possible to study also rarer processes associated with the QGP in great detail. The ALICE TPC, along with most of the other subdetectors in ALICE, are described in Chapter 4.

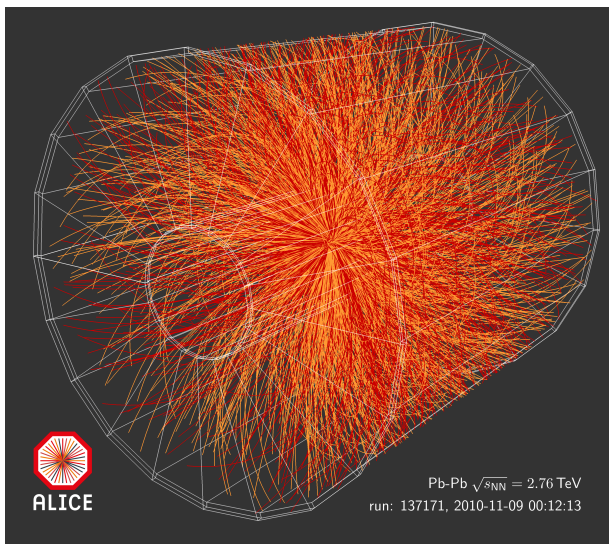

Figure 1.2 – Example of an event display in a central Pb–Pb collision recorded by the ALICE Collaboration. A total of about 2400 tracks are produced. Figure taken from Ref. [2].

While the LHC clarified many open questions about QGP physics, the most remarkable discovery – which has also been hinted at at RHIC – is that most of the observables thought to be associated with a QGP have now also been found in proton–proton collisions. This is quite puzzling, since the lifetime of the system formed in these collisions should be too short to reach thermal equilibrium – a requirement for a phase transition. Now several phenomenological models are being developed aiming at understanding this – of which some are described in Section 3.8. These are vastly different from each other, so more experimental input is needed to give insight into what happens in these collisions, which is the aim of this thesis.

1.2 Research Goals

One of the signatures previously thought to be associated with QGP formation, is the enhanced yields of baryons containing more than one strange quark (cf. Section 2.1), which is known as *strangeness enhancement*. One such baryon is the Ξ baryon, containing two strange quarks. Now it has been discovered that there is a smooth scaling

in relative yield of these hadrons when going from collisions where a low number of particles have been produced to larger collision multiplicities. This does not seem to depend on which collision system is involved, but the scaling is most prominent in pp collisions. Therefore, the strangeness production mechanism may be the key to understand what happens in small systems. To get a better understanding of this, particle correlations with the Ξ baryon are used to determine where the strange quarks are produced in the system. By comparing this with phenomenological models, this may give understanding of the mechanisms involved.

This research is closely related to the CLASH project (although not officially part of it), which is a collaboration between part of the development team behind the PYTHIA event generator (cf. Section 3.8.2) and the Lund ALICE group. The goal of the project is to develop both experimental observables that can give a large distinction power between different models, and give input to the theorists on where further development is needed.

1.3 Outline of This Thesis

I will start with introducing the Standard Model of particle physics, which is the current theory for describing all known particles and interactions (except gravity) at a fundamental level. From this I will move towards a description of the QGP and describe all major observables associated with it in heavy-ion collisions. For each observable, I will assess the current experimental status also in small systems. After introducing the observables, I will describe some of the phenomenological models aiming at finding a unified picture of the physics in small and large systems.

In the next part, I will describe the ALICE detector, with a focus on the subsystems involved in the research presented in this thesis. Then I will move on to my service task, which every PhD student in the collaboration is required to do. For me, the service task was to help out with the prototype testing of the chips used for the TPC upgrade. The detection rates enabled following the upgrade will make it possible to extend the main analysis topic to rarer processes such as Ω production, so indirectly this work will be very useful for continuing the research I have started in this thesis. The chips will however not only be used in the TPC, but also in the muon tracking chambers, where the specification for the tolerable noise limit was set below the resolution of the chip itself. I developed a novel way to access this noise level, which is described and tested in this thesis. This could for instance be of use for quality testing (of electrical shielding etc.) if not having sensitive enough equipment.

Following the detector description and upgrade, I will describe the main analysis topics. I will start with a description of a measurement of flow in Xe–Xe collisions that I was involved in. These ions were collided in the LHC during a single day in 2017. This analysis aimed at testing initial-state models by comparing flow across multiple collision systems. Since more experienced members of the analysis group did most of the contributions to these results, this analysis is described only briefly. Instead the focus in the thesis is on the main analysis, which is to measure correlations between the Ξ baryon and five different hadronic species (including Ξ itself) to get a better understanding of the strangeness production mechanism in pp collisions. The results

are compared with the phenomenological models described in Section 3.8. Finally, I will suggest several analyses for extending the work presented here.

1.4 Own Contributions

The ALICE Collaboration has about 1800 members, which all get their names on the papers published by the collaboration. Therefore, no one is directly involved in all studies carried out by the collaboration, and it is therefore only meaningful to include analyses that I have been directly involved in. These are the publications that I have contributed to (given in chronological order):

- J. Adolfsson *et al.* SAMPA Chip: the New 32 Channels ASIC for the ALICE TPC and MCH Upgrades. *JINST* **12** C04008, 2017.
  Paper summarising the results from the testing of the V2 SAMPA prototype. The final phase of this testing was done by our group in Lund, which I participated in. A longer paper on the TPC upgrade, including this work, will be submitted to JINST soon.

- ALICE Collaboration. Anisotropic flow in Xe–Xe collisions at $\sqrt{s_{\mathrm{NN}}} = 5.44\,\mathrm{TeV}$. *Phys. Lett.* **B 784**, 82-95, 2018.
  This is the first report of elliptic flow in Xe–Xe collisions. I was a part of the paper committee and participated in the writing, although my results never made into the final publication.

- J. Adolfsson. Measurements of Anisotropic Flow in Xe–Xe Collisions at $\sqrt{s_{\mathrm{NN}}} = 5.44\,\mathrm{TeV}$ Using the ALICE Detector. *MDPI Proc.* **10**, 41, 2019.
  Conference proceedings to the Hot Quarks Conference *2018, mainly summarising the above publication.*

- J. Adolfsson *et al.* QCD Challenges from pp to A-A collisions. *arXiv*: 2003.10997, 2020. Accepted for publication by *Eur. Phys. J. A*.
  Proceedings to the 3rd International Workshop on QCD Challenges from pp to A–A*, 2019, summarising all ideas which were discussed during the workshop. I participated actively in the workshop and provided some figures for the final document, which I also helped reviewing.*

- J. Adolfsson. Studying particle production in small systems through correlation measurements in ALICE. *arXiv*: 2005.14675, 2020. Submitted to *Acta Phys. Pol. B Proc. Suppl.*
  Conference proceedings to Excited QCD *2020, summarising recent correlation measurement in ALICE, including preliminary results for the main analysis presented in this thesis.*

The main work in this thesis has not yet made it into a publication. The analysis and early results have already been approved as ALICE Preliminary results. A publication is planned for the near future, although it may take a while for it to pass all the approval steps in the ALICE Collaboration.

Chapter 2

Fundamental Theory

In this chapter, the Standard Model of particle physics, which is the theory describing everything we currently understand about interactions at microscopic scales, will be introduced. Then, some elementary quantum field theory will be introduced starting from Quantum Electrodynamics (QED). Having laid the foundations, the main focus will be on Quantum Chromodynamics (QCD), which describes the strong force, the main topic of this thesis. In particular, the focus will be on how QCD gives rise to a quark-gluon plasma (QGP) at large temperatures and what properties it predicts for this QGP.

2.1 The Standard Model of Particle Physics

The Standard Model of particle physics describes all known fundamental particles, as well as three of the four fundamental forces in nature, namely the *strong*, *electromagnetic*, and *weak* interactions. The last one, *gravity*, is too weak to be possible to study at microscopic levels with current technology[1] and the theory governing it, general relativity, has proven to be very difficult to combine with the Standard Model. But that is a research field on its own and will not be discussed further here.

The strong interaction is what holds nuclear matter together. Effectively, this is by far the strongest force, but its range is very short, only a few fm. The electromagnetic interaction is what acts on charged particles. It has infinite range, is the cause of phenomena such as electricity and magnetism, and is the reason why atoms hold together, but is also important for the interactions at the LHC. The weak force is so weak that it is not experienced directly in everyday life, but it plays a major role in radioactive decay and decays of the particles studied in this thesis. Therefore, all of these interactions will be explained in detail here.

A summary of all particles in the Standard Model is given in Fig. 2.1. These are usually divided into *fermions*, which are subdivided into quarks and leptons, and *bosons*, which are subdivided into gauge bosons and scalar bosons. The fundamental difference between the two categories is that fermions have half-integer values (1/2 for fundamental

[1]The reason why we experience gravity is that it is always attractive and has infinite range, so it will have a very large impact at macroscopic scales.

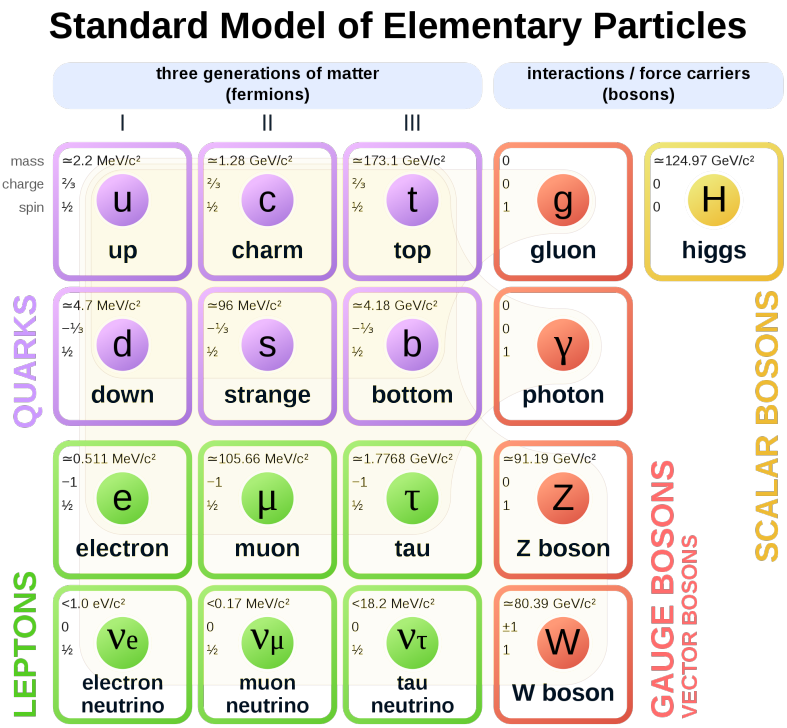

Figure 2.1 – Summary table of all particles included in the Standard Model. See details in the text. Figure taken from Ref. [3].

particles) of the intrinsic angular momentum, called *spin*, while bosons have an integer spin: 1 for gauge bosons and 0 for scalar bosons. This gives them fundamentally different properties, with the most important one being that two fermions cannot share the same quantum state, known as the Pauli principle, while bosons can. This means that the energy levels of fermionic matter are quantised, which limits the number of different (composite) particles that exist below a certain mass. For bosons, this is not the case, meaning that it is possible to put an infinite number of bosons at the same energy, with the consequence that the total number of bosons increases indefinitely with decreasing energy scale. Moreover, the total net fermion number is conserved (they can only change from one type of fermion into another, in cases where such mechanisms are allowed, or be annihilated by their anti-particles), but the boson number is not.

Fermions are matter particles, whereas gauge bosons carry the forces between them. The fermions further come in three generations, where the first generation build up (most of) the matter around us and the other two generations consist of heavier versions of the particles in the first generation. The quarks are the particles which are subject to the strong force, and thus these are the ones that build up nuclear matter. More specifically, a proton for instance consists of two up quarks and a down quark – called valence quarks – but due to quantum fluctuations there are pairs of quark–antiquarks as

well, called sea quarks, which also include heavier flavours, in particular strange quarks. When describing bound quark states, known as *hadrons*, it is usually the valence quarks that are listed.

The quarks have colour charge, which the massless gluon force mediators couple to, and come in three colours called red (r), green (g), and blue (b), and corresponding anticolours for antiquarks. All particles seen in nature are colour neutral, which either can be formed through an rgb triplet, or by a colour-anticolour pair. Consequently, there are two kinds of hadrons[2]: baryons, which consist of three quarks, and mesons, which consist of a quark–antiquark ($q\overline{q}$) pair. The hadrons most important for this work are summarised in Appendix C.

The electromagnetic interaction is mediated by massless photons, and the weak force by massive W and Z bosons. These bosons couple together into a unified SU(2) group[3] at high enough energy, but due to electroweak symmetry breaking, they have decoupled into three massive (W^+, W^-, and Z^0) and one massless state [5]. While photons couple to all charged particles, including W bosons, the massive gauge bosons couple to all fermions, including neutrinos as well as the W and Z bosons themselves. The weak interaction can further be divided into the *charged current*, mediated by the W boson, and the *neutral current*, mediated by the Z boson. The neutral current is involved in scattering processes, but unlike the charged current, it cannot change flavour. The charge current on the other hand, can change a quark or lepton into another, but only between species of different charge. This is the process through which heavier fermions decay into lighter ones, and will be described more closely in Section 2.2.

Finally, there exists one scalar boson within the Standard Model, namely the Higgs boson. This is an excitation of the Brout-Englert-Higgs field, which provides the mechanism for giving fundamental particles (except neutrinos) their mass[4], and was predicted by Peter Higgs in 1964 [6]. This was discovered through a combined effort by ATLAS and CMS in 2012, making it the most recent fundamental particle to have been discovered [7, 8].

2.2 The Electroweak Interaction and Feynman Diagrams

Quantum Electrodynamics is the theory describing the electroweak interaction, but before going into this, I will introduce an important tool for calculating and visualising interactions within the Standard Model, namely the *Feynman diagram*, which was introduced by Richard Feynman in 1948 [9]. In these, the interaction vertices are combined graphically, as shown in Fig. 2.2, where each line combining two vertices is being known as a propagator, which is a representation of the virtual field propagating the interac-

[2]Excluding more exotic state, such as the pentaquark states $P_c(4450)$ and $P_c(4380)$ discovered by LHCb in 2015 [4].

[3]Special unitary group of two dimensions.

[4]The masses of hadrons are typically much higher than the combined mass of their constituent quarks – which is a consequence of that fermions only can couple into discrete energy levels and that the strength of the strong interaction results in quite large confinement energies – but if quarks and leptons were massless, they would never be able to combine into atoms or even nuclei.

tion. Fermions are represented by arrows – which have to be connected – photons, W, and Z bosons by wavy lines, gluons by curly lines, and scalar bosons by dashed lines.

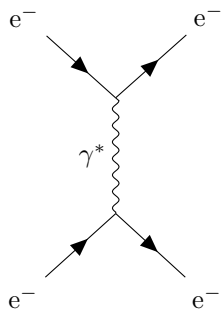

Figure 2.2 – A Feyman diagram showing electron–electron scattering. The interaction is driven by propagation of a virtual photon (hence the asterisk) which connects to the fermion fields at the nodes in the diagram. Time is propagated from left to right.

From this, the amplitude of the process – which is a part of the matrix element $\mathcal{M}$ – can be calculated by assigning the interaction strength $\sqrt{\alpha}$ to each vertex – entering as factors in the calculation – and a factor to the propagator, which in lowest-order perturbation theory is proportional to [10, p. 21]

$$\frac{1}{q^2 - M^2c^2},$$

where q is the four-momentum transfer of the interaction, M the mass of the mediator, and c is the speed of light. The full calculation for this is quite complicated and will not be done here, but the important results for the time being are that $\alpha_{\mathrm{EM}} = e^2/(4\pi\varepsilon_0\hbar c) \simeq 1/137$ for electromagnetic interactions at low energy scales, and that the electromagnetic force decreases as $1/r^2$, where r is the distance between the two charged particles. The total cross section is then proportional to $|\mathcal{M}|^2$, and consequently to $(\sqrt{\alpha})^4 = \alpha^2$ since there are two vertices entering into the calculation of the matrix element. For the weak interaction, the calculation is somewhat different, yielding the interaction strength $\alpha_{\mathrm{W}} \simeq 1/236$.

This is comparable to the electromagnetic interaction, so why is the weak interaction so weak? The answer to this lies in the propagator, since the mass term reduces the interaction probability significantly. In most practical situations, $q^2 \ll M^2$, so the cross section falls as $1/M^4$. With the masses of the W and Z bosons being 80.4 GeV/c^2 and 91.2 GeV/c^2, respectively [11], the effective range gets reduced to the order of $\hbar c/M \sim 2 \cdot 10^{-18}$ m, resulting in a very small interaction probability. Therefore, it is preferred to use the effective interaction strength $G_F = \sqrt{2} \cdot 4\pi\hbar\alpha_{\mathrm{W}}/(M_{\mathrm{W}}c) = 1.166 \cdot 10^{-5}\,\mathrm{GeV}^{-2}$ instead of α_{W} [10, pp. 36-38]. As a consequence of the low effective interaction strength, weakly decaying particles are relatively long-lived, which can be derived from the decay width Γ, which again is proportional to $|\mathcal{M}|^2$. The decay time does however have strong energy dependence. From dimensional arguments alone, one can expect that $\Gamma \propto G_F^2 Q^5$, where Q is the energy transfer, which holds as long as one does not consider inhibiting factors such as mixed states (see below) and changes in spin.

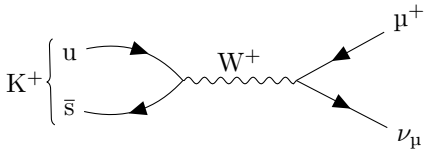

Figure 2.3 – Feyman diagram describing K^+ decay into a muon and a muon neutrino.

An example of a weak process is charged-kaon decay, which is shown in Fig. 2.3. This particle has a lifetime of $1.24 \cdot 10^{-8}$ s, which can be compared to typical strong interactions with lifetimes of $\sim 10^{-24}$ s and electromagnetic interactions where it is $\sim 10^{-15}$ s. Here the W boson changes the flavour of the strange (s) quark into an up (u) quark. The reason why this is possible at all is that flavour and mass states mix. Thus the weak decay will turn the s quark into a mixed state of $(\mathrm{u}, \mathrm{c}, \mathrm{t})$, but since the charm quark is too heavy, it can only change into an up quark. This mixing is described by the CKM matrix, which is a unitary matrix on the form

$$V_{\mathrm{CKM}} = \begin{pmatrix} V_{\mathrm{ud}} & V_{\mathrm{us}} & V_{\mathrm{ub}} \\ V_{\mathrm{cd}} & V_{\mathrm{cs}} & V_{\mathrm{cb}} \\ V_{\mathrm{td}} & V_{\mathrm{ts}} & V_{\mathrm{tb}} \end{pmatrix},$$

where $V_{\alpha\beta}$ is the relative coupling strength between quark states α and β in the weak interaction, or more precisely

$$|g_{\alpha\beta}|^2 = |V_{\alpha\beta}|^2 g_{\mathrm{W}}^2,$$

where $g_{\mathrm{W}}^2 = 4\pi\alpha_{\mathrm{W}}$, so the strength is reduced by $|V_{\alpha\beta}|^2$ compared to lepton interactions and the decay width by $|V_{\alpha\beta}|^4$. The matrix element describing the coupling between u and s has an absolute value of $|V_{\mathrm{us}}| = 0.2248$, meaning that the decay probability is suppressed by a factor of $2.6 \cdot 10^{-3}$ compared to pure quark states. Unitarity implies that $|V_{\mathrm{us}}|^2 + |V_{\mathrm{cs}}|^2 + |V_{\mathrm{ts}}|^2 = 1$, further implying that the coupling to the charm quark is much greater (the coupling to the top quark is negligible), but not energetically possible given the higher c quark mass. This explains why the lifetime of kaons is only about a factor of two shorter than that of pions, despite the much larger mass (494 MeV/c^2 compared to 140 MeV/c^2).

Screening

So far, I have just mentioned Feynman diagrams at tree level, but higher-order processes also enter into the calculation of the total cross section, such as the ones shown in Fig. 2.4. These are called loop diagrams. In the case of QED, next-to-leading order (NLO) diagrams contain processes where a dilepton or diquark pair is formed. In Fig. 2.4a, the dielectron pair is formed within the virtual photon field surrounding the electron, and is known as a quantum fluctuation. This increases the effective charge of the electron, effectively increasing the coupling strength at short distances, which is known as *screening*. In quantum field theory, the increased strength comes from NLO diagrams such as the one shown in Fig. 2.4b. In practice, also processes involving quarks and other leptons contribute to the screening.

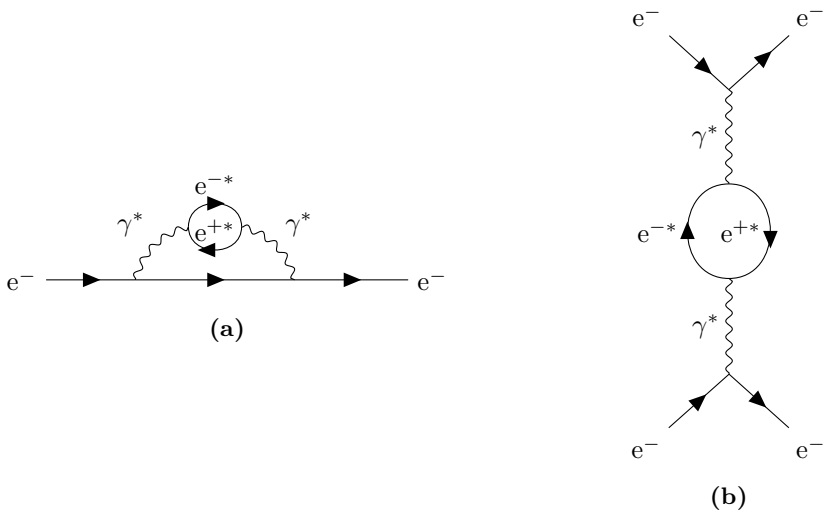

Figure 2.4 – Example of NLO Feynman diagrams of **(a)** quantum fluctuations and **(b)** electron–electron scattering (cf. 2.2), both involving the creation of a virtual dielectron pair.

In quantum mechanics, shorter distances are equivalent to higher energy, and thus the electromagnetic interaction strength increases with energy, i.e. α_{EM} is energy dependent. Doing the full calculation leads to the result [13, pp. 234-235]

$$\alpha_{\mathrm{EM}}(q^2) = \frac{\alpha_{\mathrm{EM}}(\mu^2)}{1 + k\alpha_{\mathrm{EM}}(\mu^2)/(3\pi)\ln(\mu^2/(-q^2)}, \tag{2.1}$$

where q^2 is the energy scale where we want to measure the coupling strength and $-\mu^2$ is some reference energy scale, preferably the one where α_{EM} is normally defined. The factor k comes from summing the contributions from all fermions, and is

$$k = n_{\mathrm{l}} + \frac{4}{3}n_{\mathrm{u}} + \frac{1}{3}n_{\mathrm{d}},$$

where n_{l} is the number of lepton flavours, and n_{u} and n_{d} are the numbers of quark flavours with charge $+2/3$ and $-1/3$, respectively, available at the energy scale μ^2. If $|q^2| > M_{\mathrm{W}}^2$, this would also enter into the scaling factor. Coupling strengths which vary with energy are known as *running couplings*, and are even more important in QCD, as will be shown in the next section.

2.3 Quantum Chromodynamics

Two results from QCD make it very different from QED, and these are that there are three colour charge states, and that the gluon itself has colour charge. Some implications of the former is that one requires three quarks or an quark-antiquark pair to form a colour neutral object, as discussed in Section 2.1, and that there exist baryons where

all valence quarks have the same flavour and occupy the same quantum state, which would otherwise violate the Pauli principle[5]. The second result does however have far more dramatic consequences, which will be described here.

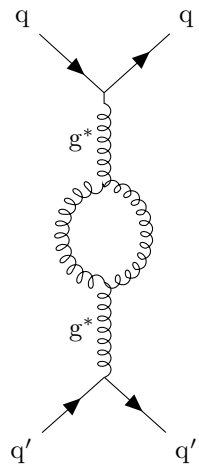

Figure 2.5 – NLO diagram of scattering between two quarks $(\mathrm{q}, \mathrm{q}')$ in QCD, where a gluon loop is formed due to self-interaction.

One important implication of the gluon self-interaction is that the screening diagram in Fig. 2.4b is complemented by diagrams such as the one shown in Fig. 2.5. While similar diagrams as the QED process also exist (with diquark pairs instead of dilepton pairs and gluons instead of photons), they turn out to be less important. If increasing the distance, the number of possible gluon loops increases, leading to an *increased* interaction strength with distance, i.e. the screening has been replaced by an anti-screening effect. The running coupling constant in QED defined in Eq. (2.1) then has the following analogy in QCD [13, pp. 236-239]:

$$\alpha_{\mathrm{s}}(q^2) = \frac{\alpha_{\mathrm{s}}(\mu^2)}{1 + (33 - 2n_f)\alpha_{\mathrm{s}}(\mu^2)/(12\pi)\ln(-q^2/\mu^2)}, \tag{2.2}$$

where α_{s} is the strong coupling constant and n_f is the number of flavours accessible at the measured q^2 (this term is what enters from the $\mathrm{q\overline{q}}$ screening mentioned above). A consequence of the reversal of the factor within the logarithm as compared to QED, is that α_{s} gets weaker with increasing energy scales. This theory agrees very well with measured values of α_{s} at different energies, as is shown in Fig. 2.6. The standard reference for α_{s} is measured at $|q| = m_{\mathrm{Z}}$, where $\alpha_{\mathrm{s}} = 0.1182 \pm 0.0012$ [11].

This is only about an order of magnitude greater than α_{EM}, so why is the strong force so strong? Here one needs to remember that the energy scale associated with the hadrons observed in the Universe is much lower than this. In particular, there is an energy scale where the denominator in Eq. (2.2) approaches zero, meaning that the

[5]The discovery of one such state, the Ω baryon, gave strong evidence that the quark model is correct [12].

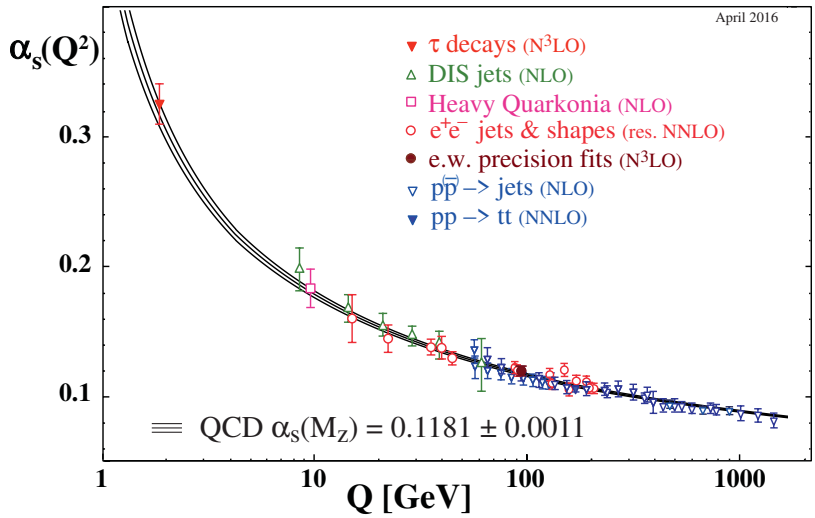

Figure 2.6 – Measured values of α_s as a function of energy scale Q (q in text), along with predictions based on Eq. (2.2). Figure taken from Ref. [11].

strong coupling constant essentially gets very strong. Solving for the measured value of $\alpha_s(m_Z)$, where $n_f = 5$, leads to a divergence at $q \sim 90\,\text{MeV}$, but due to additional higher-order processes, the stated model for the strong coupling constant does not really hold in this regime. Instead, one defines an energy scale Λ_{QCD} at $\alpha_s \simeq 1$, which in some sense is a limit for when the strong coupling gets strong. Again using Eq. (2.2) leads to $\Lambda_{QCD} \sim 200\,\text{MeV}$. Well above this value, QCD becomes perturbative, making calculations much easier (or in many cases even possible).

The energy scale of quarks confined into a hadron is below Λ_{QCD}, and thus QCD is non-perturbative and dominated by higher-order diagrams. Consequently, the force between two quarks in a hadron does not decrease if trying to separate them, which can be seen by looking at the strong potential, which in the non-relativistic case[6] can be approximated as

$$V(r) \approx -\frac{\alpha_s}{r} + \kappa r, \tag{2.3}$$

where the constant $\kappa \sim 1\,\text{GeV/fm}$ [10, p. 182]. The force is obtained by taking the derivative of $V(r)$, so the first term is the same as for electromagnetic interactions. The second term, however, stabilises at a very large value at large r. As a consequence, all quarks in the Universe are confined into hadrons, which is known as *confinement*. If

[6]Strictly speaking, this is only applicable to heavy-flavour quarks, but as will be described in Section 3.8.2, this potential can be used for modelling hadron fragmentation also at relativistic energies.

trying to separate two quarks, e.g. in a collider experiment, a shower of new hadrons is formed, which is known as a *jet*.

Lattice QCD

In general, the approach with Feynman diagrams is successful when the interaction strength, α, is small since then higher order diagrams become decreasingly important. For the strong interactions this is in general not the case – the theory is non-perturbative. Therefore, for instance confinement has never been shown analytically. One can however get some insights by a technique called *lattice QCD*, where QCD processes are quantised and calculated by numerically solving path-integrals over a fixed lattice in space [14]. This is very resource demanding, at least if aiming at achieving reasonably high precision. Here, one has to introduce some cut-off scale $1/a$ in the spacing, and to obtain meaningful results, one needs to extrapolate this to the limit $a \to 0$, which is not trivial.

Despite its limitations, lattice QCD has led to some remarkable results. One of the results is the strong potential defined in Eq. (2.3), but another one, which is particularly relevant for the research presented in this thesis, is that lattice QCD has given predictions of the QCD phase diagram, which will be discussed next.

2.4 The QCD Phase Diagram

As has already been argued, due to the large value of α_s at the non-perturbative-QCD energy scale, quarks only exist in confined states. There is however nothing that prevents us from increasing the energy scale *within* a hadron or some other QCD medium, which is what happens during a collision in a collider such as the LHC. This will cause α_s to decrease, and if the energy is high enough, it becomes so weak that quarks start moving freely within the medium. This is known as *asymptotic freedom*, and if reaching such a high energy scale, we are able to study QCD processes experimentally.

The idea of asymptotic freedom at high QCD energies was introduced already in the 1970s, and in 1975 it was predicted that at high enough temperature or net baryon density, hadronic matter will transform into another phase of matter, which a few years later was dubbed the Quark-Gluon Plasma (QGP) [15, 16]. More recently, more advanced computations using lattice QCD have shown that there is a clear transition in equation of state in the temperature range $T_c \sim 140 - 190\,\text{MeV}$ (in natural units, $\hbar = c = k_B = 1$), as is shown in Fig. 2.7 [17, 18]. This should be attributed to the just mentioned phase transition. Along with other theoretical models, this has provided input to the full QCD phase diagram, which is shown in Fig. 2.8.

As seen, there are two ways to create a quark-gluon plasma: either one can increase the baryon density by compressing nuclear matter, which happens naturally in neutron stars, or one can increase the temperature, which would happen if one could go backwards in time to the first few moments after Big Bang. In the former, a QGP is expected to exist within the neutron star core, with the possibility of a state of matter behaving like a colour superconductor[7]. Two planned experiments – BM@N at NICA,

[7]It is predicted that if the baryonic density is high enough, colour symmetry will be broken, so the

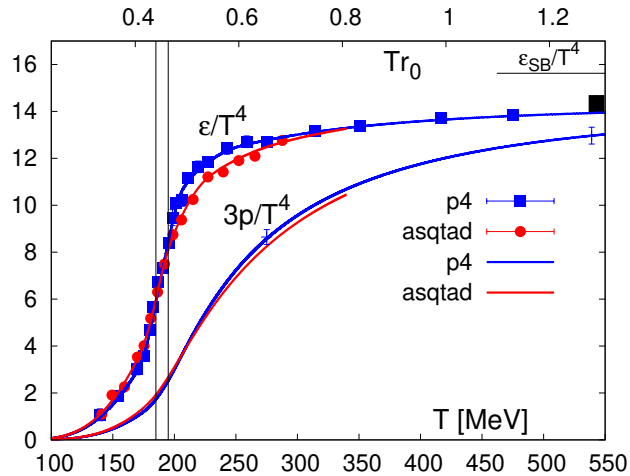

Figure 2.7 – Energy density and pressure (scaled by T^{-4}) as a function of temperature in a hadron gas, calculated using Lattice QCD using two different actions. As seen, there is a clear transition starting just above 150 MeV. Figure taken from Ref. [17].

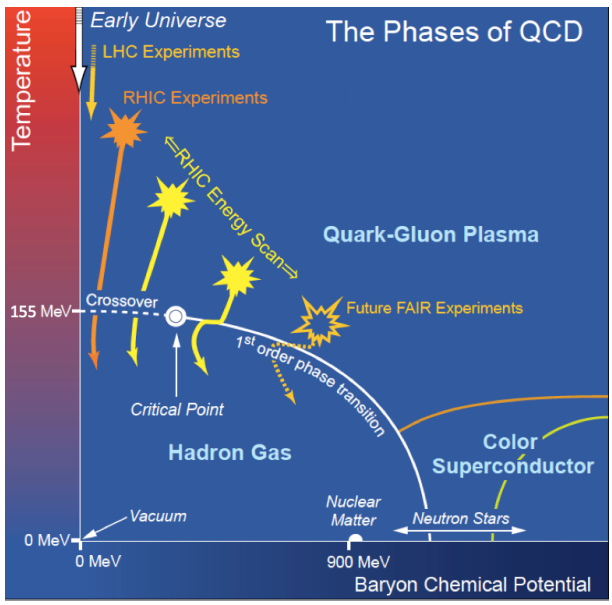

Figure 2.8 – Current understanding of the QCD phase diagram as a function of temperature and baryon chemical potential (= net baryon density). Conditions for normal nuclear matter, neutron stars, and hypothesised conditions just after Big Bang are indicated in the figure, along with a few trajectories mapped by a selection of experiments. Figure adapted from Ref. [19].

potential of two colours will be different than for the third one [20]. This is analogous to what happens in an electric superconductor.

where the commissioning has already started, and CBM at FAIR – are going to study this regime, which will hopefully shed some more light on the equation of state. In the latter regime, a QGP is expected to have existed during the first $\sim 10^{-5}$ s after Big Bang [21]. This region is being probed by heavy-ion collisions at the RHIC and LHC experiments – and in the latter case in particular by ALICE, as will be described later in this thesis. In this way, the goal is to gain more understanding both of QCD and the conditions in the early Universe.

An important difference between the high-temperature phase transition and the one at high baryon density, is that in the former case, there is a 2nd order phase transition or cross-over, meaning that there is a gradual transition from one phase to the other (this is the reason for the rather smooth behaviour in Fig. 2.7), while in the latter case, there is expected to be a 1st order phase transition, i.e. a sharp transition (like the one between ice and water). If this picture is correct, there has to exist a critical point somewhere in between, which is being searched for by the Beam Energy Scans at RHIC. Close to the critical point, there are expected to be large fluctuations in several observables, and there is a large ongoing effort in trying to discover these and rule other explanations out.

2.5 Properties of the Quark-Gluon Plasma

In this section, a few of the key features of the QGP in the high-temperature regime will be described. At first, it should be noted that the picture of asymptotic freedom may not be entirely correct, since the medium is still strongly interacting (while α_s is small it is still being far from zero) and the quarks and gluons interact with each other to a large degree. Instead, it is preferred to use the term *deconfinement*, i.e. the quarks are no longer in a confined state. As a consequence, hadrons are unlikely to form in this medium, and due to the strong interactions, it will interact with any hadronic matter traversing it.

To access some of the medium properties in the QGP, one can make use of a quite remarkable duality, called the Anti-de-Sitter/Conformal Field Theory (AdS/CFT) conjecture, which states that any conformal field theory, such as QCD, can be transformed into anti-de-Sitter space, which is a quantum-gravity formalism for black holes [22]. In the latter regime, which is weakly coupled, some problems which are non-perturbative in the standard formalism can be solved exactly or perturbatively, which has led to predictions of the transport properties and viscosity of the QGP. One should keep in mind however that the transformation is somewhat idealised, assuming a few unphysical properties of the QGP such as an infinite number of colours, but nevertheless the calculations are useful. This has revealed that the QGP is expected to behave like a near-perfect fluid, with a lower bound of the ratio of shear viscosity to volume entropy density at [23]

$$\eta/s \geq \frac{\hbar}{4\pi k_B} \approx \frac{0.08\hbar}{k_B}.$$

As will be discussed in Chapter 3, experimental results have shown that η/s in the QGP formed in heavy-ion collisions is close to this value. In fact, there are no other known fluids in the Universe with a lower η/s. As a comparison, superfluid helium has

a minimum at $\eta/s \approx 0.7\hbar/k_B$, and in graphene – which also has proven to behave like a superfluid – it has been measured to be $\eta/s \approx 0.2\hbar/k_B$ [24].

These properties give several predictions which can be tested experimentally. These will be discussed in detail in Chapter 3, along with experimental results.

Chapter 3

Experimental Signatures of the Quark Gluon Plasma

The QGP is studied through heavy-ion collisions, where a QGP with a lifetime of the order of 10 fm/c ($\sim 4 \cdot 10^{-23}$ s) can be created [25]. Despite this short timescale, we can learn much about the QGP from heavy-ion data. At the LHC, this has mostly been limited to Pb–Pb collisions, whereas a larger variety of collision systems have been available at RHIC, although most of the results produced there are from Au–Au collisions. Since the QGP has properties of a medium, it needs to have reached some degree of equilibrium, and therefore a certain system size should be required for it to form. Therefore, a QGP is not expected to be formed in small systems such as pp and p–Pb collisions. In this picture, these systems were thought as references to probe hadron- and cold-nuclear-matter effects[1], respectively. More recently, however, many of the signatures expected to be associated with a QGP have been observed also in these small systems, which is quite puzzling.

In this chapter, I will go through some of the most important of these signatures, along with their experimental status both in large in small systems. At the end of the chapter, I will describe some of the phenomenological models which aim at finding a unified picture for what happens in both large and small systems.

3.1 The Evolution of a Heavy-Ion Collision

The evolution of a heavy-ion collision can be characterised by four different phases: the pre-equilibrium phase when the QGP is formed, the hydrodynamic expansion, the chemical freeze-out when hadrons are formed, and the kinetic freeze-out. These are summarised in the space-time diagram shown in Fig. 3.1. This picture is used to illustrate that fast-moving hadrons are produced later in the collision. Thus, the hadron

[1]Since the colliding proton will only interact with a few of the nucleons in the lead nucleus, data from p–Pb collisions will reveal effects of interactions between the particles produced in the collision and the (cold) nuclear medium of non-participating partons.

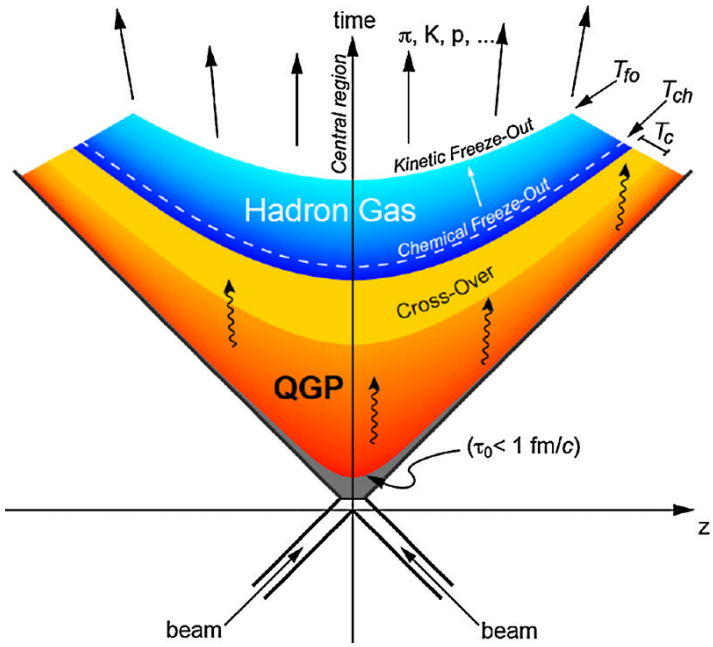

Figure 3.1 – Evolution of a heavy-ion collision, shown as expansion along the beam axis versus time. The different phases discussed in the text are labelled in the figure, except the pre-equilibrium phase, which is marked by dark grey. The wavy lines mark thermal photons radiated by the QGP and the straight arrows hadrons released during kinetic freeze-out. The different temperatures associated with each phase are marked in the upper right corner. Figure created by Boris Hippolyte; this version is retrieved from Ref. [26].

formation time scales with the longitudinal proper time $\tau = \sqrt{t^2 - z^2}$, which is the reason for the behaviour in the figure [27, pp. 73-83].

Several models have been developed aiming at describing the formation of a QGP in heavy-ion collisions; of these the Color Glass Condensate (CGC) model has been the most successful. This picture is inspired by the rapidly increasing gluon density measured at low momentum-fraction x of the proton by the HERA experiment [28]. At some point, the density would be so high that for any further changes (going to smaller x at the same energy level) the number of extra gluons from splitting and gluons lost to absorption cancel so that the density saturates. Before the collision, the ions will be strongly Lorentz contracted, resembling pancakes, resulting in a very high transverse gluon density in the collision frame. This high density reduces the interaction strength, hence these low-x gluons move slowly and are trapped in a "glass". The words "condensate" and "color" are used due to the high number density and the coloured nature of gluons, respectively[2]. During the collision, the two gluonic sheets will pass through each other, creating strong colour-electric and -magnetic fields in

[2]The American spelling of colour/color is preferred when describing the CGC, since it has seen most of its development in the American community.

the longitudinal direction, not completely different from the string picture described in Section 3.8.2. The matter created between the two CGC sheets is called a *glasma*, which is a highly coherent state [29]. Such a state of matter is not stable, and would in a short while ($< 1\,\mathrm{fm}/c$) decay into a QGP.

There are many different models for the process of going from a glasma-like state into a thermalised medium, with some of the most common involving parton (i.e. quark, gluon) rescattering by the built-up colour field [30]. To access the initial state experimentally, one needs to tie it to some other observable, e.g. anisotropic flow (cf. Section 3.7), and compare the result to a model prediction. This is far from trivial, since it involves propagation through the later phases of the collision evolution, and therefore requires other models describing this.

As the QGP behaves like a nearly perfect fluid (cf. Section 2.5), the hydrodynamic evolution is modelled by relativistic hydrodynamics [27, p. 80]. The plasma created is highly collective, meaning that its expansion is dependent on the overall medium rather than individual partons, which is also the case for liquids in our everyday life. It is also very hot, and thus it emits black-body radiation in the form of high-energy photons.

As the plasma expands, it cools down. Eventually it reaches the critical temperature $T_c \sim 165\,\mathrm{MeV}$ where the phase transition (cross-over) into hadronic matter occurs (cf. Section 2.4). Then the plasma will transform into a Hadron Resonance Gas (HRG), which is a gas that consists of hadrons, i.e. hadronisation has occurred and the system is no longer deconfined. The HRG behaves like a weakly coupled QCD gas, i.e. the hadrons still interact with each other, which can change their momenta.

Finally, during the kinetic freeze-out, the system has become so diluted that the hadrons stop interacting with each other. These hadrons – or their decay products – are the particles that are observed in the detector.

3.2 Collision Geometry

Since nuclei, as well as protons, are composite objects, they have an extension in space; for a Pb nucleus the radius is about 6.6 fm [31]. As such, not all nucleons will participate in the collision and contribute to QGP formation. Nucleons (or partons) that take part

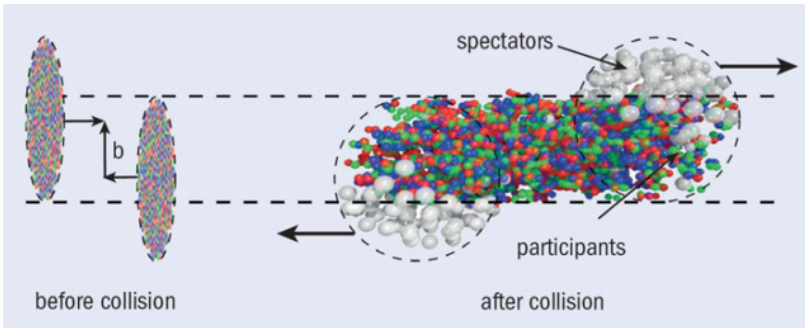

Figure 3.2 – Definitions of impact parameter b, spectators, and participants in a nuclear collision, where the impact of Lorentz contraction has been reduced for clarity. Figure taken from Ref. [32].

in the collision are called *participants*, whereas the non-participating nucleons are called *spectators*. The number of participants depend on the impact parameter b, which is a measure of the distance between the centres of the nuclei during collision. Therefore the number of participants scale with decreasing b. These concepts are summarised in Fig. 3.2.

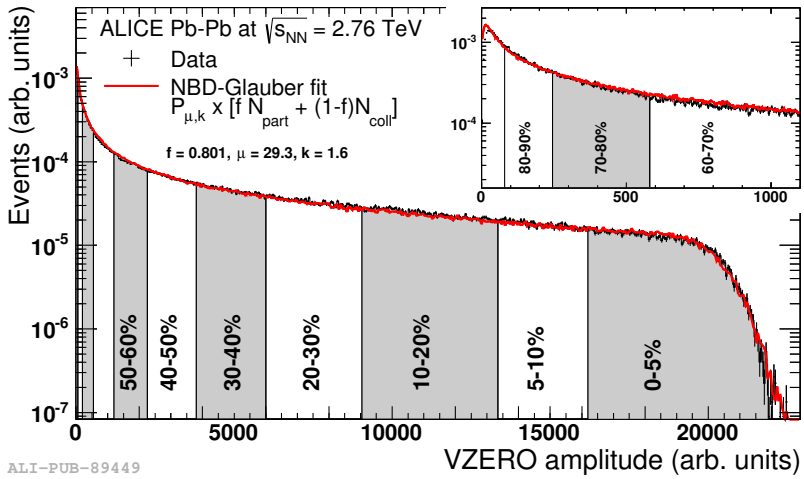

Figure 3.3 – Distribution of number of events versus V0 amplitude (the measure used by ALICE for estimating collision multiplicity) measured for Pb–Pb collisions at 2.76 TeV/nucleon and divided into centrality classes corresponding to a certain fraction of the total number of events. The inlay shows the most peripheral region. The red line is a fit to the data using a Glauber model. Figure taken from Ref. [33].

The system size, and consequently the number of produced hadrons, scale with the number of binary collisions; in a central Pb–Pb collision, this is about 1900 [33]. In practice, it is convenient to divide the events into centrality classes depending on the number of produced hadrons, and indirectly the number of binary collisions. Here 0% means a head-on collision and 100% an ultra-peripheral collision. The distribution is then fitted to the expected distribution from the Glauber model (see below). The result of such a fit is shown in Fig. 3.3. Then each event is assigned to a centrality class based on the measured event activity; for ALICE the procedure is described in Section 4.3.

Similar multiplicity classes can be defined also for small systems, but here the concept of centrality is less straightforward, given that there are only a few participating nucleons. Moreover, the events are less isotropic due to larger relative fluctuations, making a multiplicity measurement more difficult for detectors that have a limited acceptance, such as ALICE.

3.2.1 The Glauber Model

The Glauber model is a geometrical model used for simulating the number of participants in a heavy-ion collision [34]. The idea is to simulate the number of interacting nucleons based on the nucleon densities of the colliding nuclei, which in a spherically

symmetric nucleus can be modelled by the Woods-Saxon potential

$$\rho(r) = \frac{\rho_0(1 + wr^2/R^2)}{1 + \exp((r - R)/a)} \tag{3.1}$$

where ρ_0 is the density at the centre of the nucleus, a is the nuclear thickness parameter, $w \approx 0$ is a parameter for characterising deviations from spherical shape, R is the nuclear radius, and r is the radial coordinate. By assuming that nucleons move independently in the nucleus and are unaffected by each other, one can use this distribution to randomly sample the positions of all nucleons. The interaction probability in the collision is given by the cross-section σ_{NN}, which needs to be measured experimentally, e.g. through a scattering experiment. Then the total cross-section in an AA collision depending on the impact parameter is

$$\frac{\mathrm{d}\sigma^{\mathrm{AA}}(b)}{\mathrm{d}^2\mathbf{b}} = 1 - \prod_{i,j} \int \mathrm{d}^2\mathbf{r}_i \mathrm{d}^2\mathbf{r}_j \left(1 - \frac{\mathrm{d}\sigma^{\mathrm{NN}}(b_{\mathrm{ij}}}{\mathrm{d}^2\mathbf{b}}\right) \rho(\mathbf{r}_i)\rho(\mathbf{r}_j),$$

where $b_{ij} = |\mathbf{b} + \mathbf{r}_i - \mathbf{r}_j|$, and the local positions $\mathbf{r}_i$ and $\mathbf{r}_j$ and the global impact parameter $\mathbf{b}$ are given by the MC simulation. In practice, the total cross-section needs to be separated into an elastic and an inelastic part in order to determine how many nucleons are participating in the collision.

For a more accurate description at high energies, one needs to include the nuclear substructure in the simulation, i.e. using nucleons as sources in the MC generation is not enough. This has been done by using various number of quarks as sources, with 3-5 sources per nucleon yielding the best description of the collision geometry at LHC energies [35].

3.3 Measuring the Temperature: Thermal Photons

As mentioned, the QGP emits thermal radiation in the form of photons, but can we measure it? In fact we can, but the measurement is quite challenging. Photons are abundant, but there is a large background from electromagnetic decays (in particular from π^0), so one has to measure this precisely. Nevertheless, this has successfully been done at ALICE and the RHIC experiments.

What has been done is to carefully sum the photon spectra from all decay processes – requiring extensive use of MC generators – into a combined spectrum called the 'cocktail'. Then the thermal spectrum is obtained by taking the ratio between the measured spectrum γ_{inc} and the cocktail γ_{decay}, or in practice the double ratio

$$R_\gamma = \frac{\gamma_{\mathrm{inc}}/\pi^0}{\gamma_{\mathrm{decay}}/\pi^0}, \tag{3.2}$$

which is done in order to cancel out systematic uncertainties [36]. The result is shown for central Pb–Pb collisions at $\sqrt{s_{\mathrm{NN}}} = 2.76\,\mathrm{TeV}$ [3] in Fig. 3.4. As seen, there is an excess both at low and high transverse momentum (p_{T}) but the high-p_{T} component

[3] Centre-of-mass energy per nucleon pair.

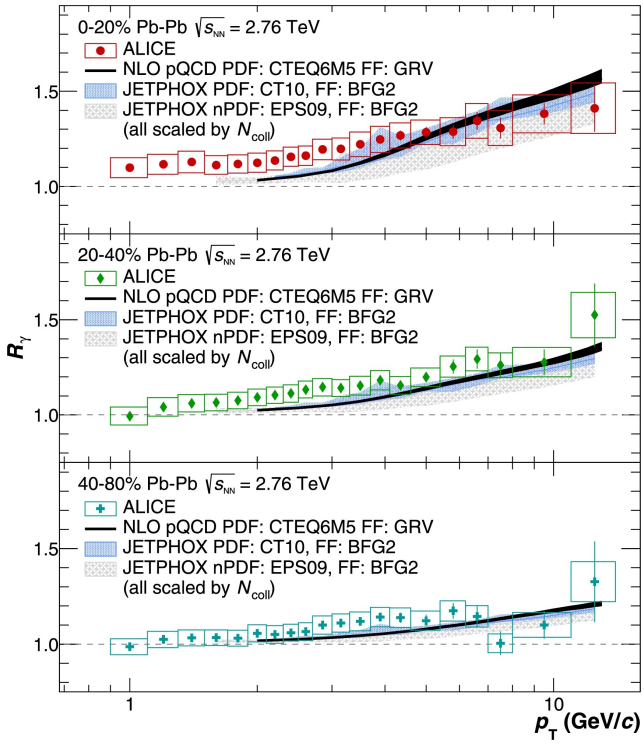

Figure 3.4 – Double ratio R_γ, defined in Eq. (3.2), for Pb–Pb collisions at $\sqrt{s_{\mathrm{NN}}} = 2.76\,\mathrm{TeV}$, divided into three different centrality classes starting from the most central events in the top panel. The results are compared with expectations from three different perturbative-QCD calculations. Figure taken from Ref. [36].

can be understood from NLO processes not taken into account when measuring the cocktail. The low-p_{T} component on the other hand, is due to thermal radiation. This can be fitted to an exponential curve of the form

$$A \exp(-p_{\mathrm{T}}/T_{\mathrm{eff}}),$$

corresponding to the tail of the black-body radiation spectrum, to access the effective temperature T_{eff} averaged over the lifetime of the medium. This yields $T_{\mathrm{eff}} = 297 \pm 14\,(\mathrm{stat.}) \pm 41\,(\mathrm{syst.})$ MeV in the most central bin, or $\sim 3.4 \cdot 10^{12}\,\mathrm{K}$. The peak temperature will of course be higher than this, although the exact value is model dependent and this measurement has too large uncertainties to distinguish between models. Nevertheless this result has made it into the Guinness World Records as the highest achieved artificial temperature [37].

In small systems on the other hand, as well as peripheral Pb–Pb collisions, the error bars are still too large to find any deviation from the cocktail not consistent with the NLO calculation, as is shown in Fig. 3.5 and the lower panel of Fig. 3.4. With more statistics becoming available in ALICE in Run 3 and beyond, the situation might change, but until then, this observable cannot be used to find evidence for a QGP in

small systems.

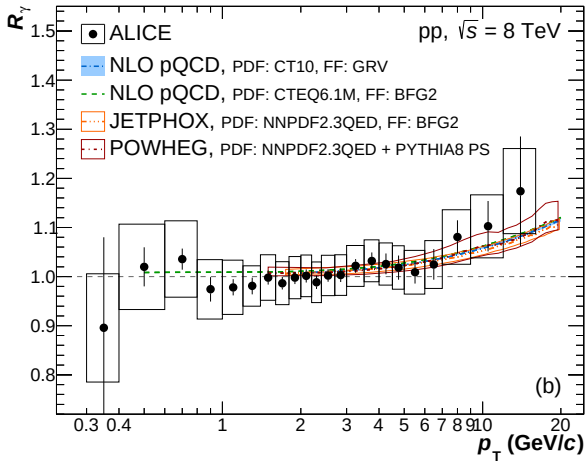

Figure 3.5 – Measurement of R_γ (cf. Eq. (3.2)) in pp collisions at $\sqrt{s} = 8\,\text{TeV}$, along with perturbative-QCD model predictions. Figure taken from Ref. [38].

3.4 Jet Quenching

Perhaps the most striking consequence of the QGP is that due to interactions with the medium, jets – i.e. showers of hadrons originating from separation of coloured particles – passing through it start to dissolve. This is known as *jet quenching*. It should be pointed out that this phase happens well before hadronisation, so it is the parton(s) and the colour field itself, which later gives rise to a jet, that are dissolving.

Jets originate from hard (i.e. high-energy) processes where scattering occurs between two quarks or gluons with large momenta, which results in the build-up of a colour string (cf. Section 3.8.2). Along this string, energy will be radiated in the form of gluons, resulting in formation of a collimated hadron shower. If the string is passing through a strongly-interacting medium, however, it will absorb these gluons and hence the jet loses energy – it gets quenched. Therefore, jets going back-to-back, which would have been of similar energy in a pp collision, have a strong momentum imbalance, as shown in Fig. 3.6.

A simple way to quantify this experimentally, is to make use of the nuclear modification factor

$$R_{\text{AA}} = \frac{\mathrm{d}^2 N_{\text{AA}}/\mathrm{d}p_{\text{T}}\mathrm{d}\eta}{N_{\text{coll}}\mathrm{d}^2 N_{\text{pp}}/\mathrm{d}p_{\text{T}}\mathrm{d}\eta}, \tag{3.3}$$

where N_{AA} and N_{pp} are the number of produced particles of interest measured in A–A and pp collisions, respectively, and η is the pseudorapidity (cf. Section 8.2). This is a measure of how much the spectra are changed in a heavy-ion collision compared to the pp reference, i.e. it measures the impact of the QGP medium. This is scaled by the number of binary collisions N_{coll}, since the number of produced particles trivially scales

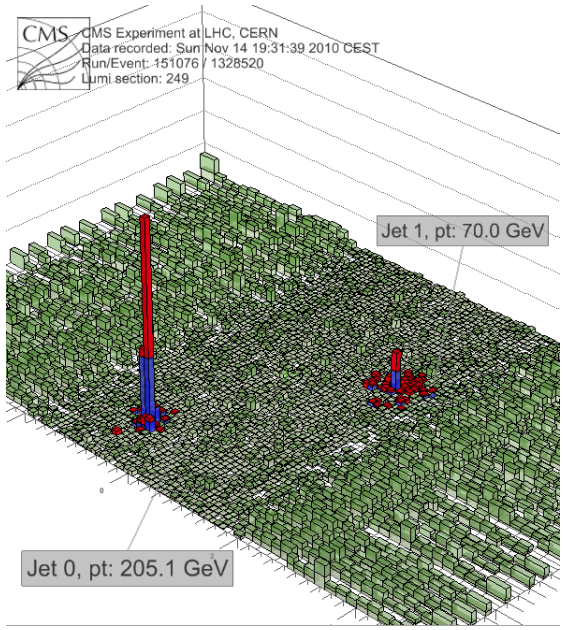

Figure 3.6 – CMS event display of a Pb–Pb dijet event with a collision energy of $\sqrt{s_{\mathrm{NN}}} = 2.76\,\mathrm{TeV}$, draw as a Lego plot with collected energy on the z axis, pseudorapidity going from top left to bottom right, and azimuthal angle in the other direction. Due to jet quenching, there is a large momentum imbalance between the jets. Figure taken from Ref. [39].

with this number if there are no in-medium effects. While Eq. (3.3) could be modified to quantify the impact of many different parameters, measuring the spectra is useful since most of the energy from the fragmenting parton goes into the leading hadron (i.e. the one with the highest p_{T}). The reverse also holds, i.e. most high-p_{T} hadrons come from the leading jet, so jet quenching will manifest itself as a drop in the high-p_{T} part of the inclusive hadron spectrum. A measurement of R_{AA} for inclusive charged hadrons, as well as separated into pions, kaons, and protons, in Pb–Pb collisions is shown in Fig. 3.7. The jet quenching effect is very prominent in central collisions and decreases as one is going towards more peripheral collisions, as expected when the size of the QGP decreases.

Phenomenologically, jet quenching is far from trivial to model, in particular since the underlying physics is non-perturbative. There are however a few useful quantities that can be used to characterise this, with the most universal one being the *transport coefficient*

$$\hat{q} = \frac{(\hbar c)^2}{r_{\mathrm{D}}^2 \Lambda},$$

where Λ is the mean free path of gluons in the medium, and r_{D} is the Debye length, or screening length associated with the QCD (anti)screening described in Sections 2.2–2.3[4].

[4]In a QGP, r_{D} is smaller than the separation between quarks, which is another way to explain

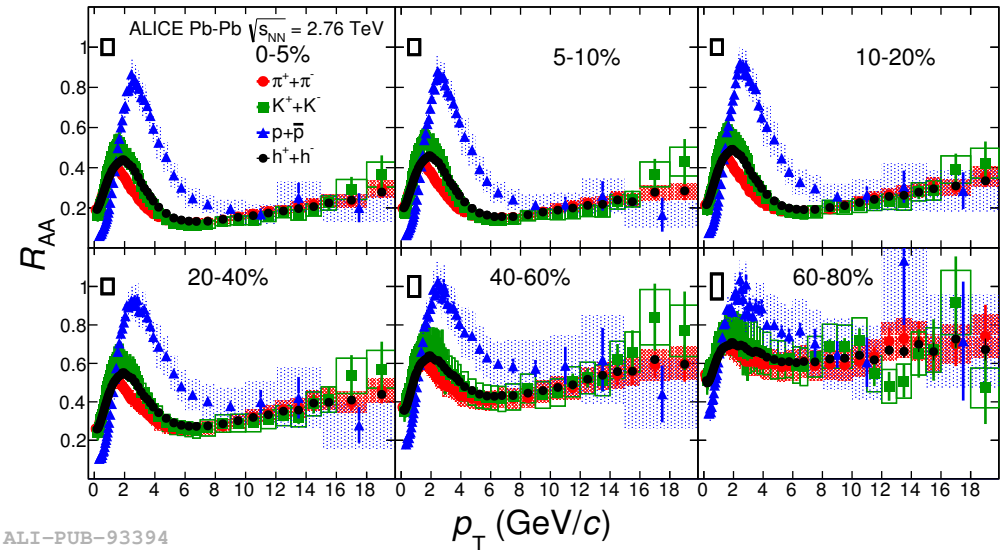

Figure 3.7 – Nuclear modification factor of inclusive charged hadrons as well as pions, kaons, and protons, measured in Pb–Pb collisions at $\sqrt{s_{\mathrm{NN}}} = 2.76\,\mathrm{TeV}$ for different centrality classes. Figure taken from Ref. [40].

This is a particularly useful quantity, since it has been shown that the jet energy loss approximately follows [41]

$$\Delta E^{\mathrm{jet}} \propto \hat{q} L^2, \tag{3.4}$$

where L is the path length traversed by the jet in the medium. To actually extract the $\hat{q}$ parameter is much more complicated (Eq. (3.4) only gives the proportionality), but can be done by fitting R_{AA} measurements to jet quenching models. The result is [42]

$$\frac{\hat{q}}{T^3} \approx \begin{cases} 4.6 \pm 1.2 & \text{at RHIC,} \\ 3.7 \pm 1.4 & \text{at the LHC,} \end{cases}$$

in natural units. Inserting approximate values for the temperatures yield

$$\frac{\hat{q}}{T^3} \approx \begin{cases} 1.2 \pm 0.3\,\mathrm{GeV}^2/\mathrm{fm}, & T = 370\,\mathrm{MeV}, \\ 1.9 \pm 0.7\,\mathrm{GeV}^2/\mathrm{fm}, & T = 470\,\mathrm{MeV}. \end{cases}$$

To check how well these values agree with QCD, $\hat{q}$ cannot be calculated from QCD directly, but just as for η/s, the AdS/CFT conjecture can give some insight (cf. Section 2.5). Even though this makes calculations simpler, one needs to go to *at least* NLO to obtain a result which agrees reasonably with data. Keeping in mind that the solution is obtained from an idealised QGP, one obtains $2.27 \lesssim \hat{q}/T^3 \lesssim 3.64$, which is close to the fit to experimental results.

deconfinement.

As R_{AA} uses the pp measurement as a reference, Eq. (3.3) cannot be used to measure jet quenching in pp collisions. Moreover, following Eq. (3.4), the jet energy loss should scale as $N_{\mathrm{coll}}^{2/3}$ (assuming a spherically symmetric QGP) if the mechanism is the same as in large systems, meaning that it is a very small effect.

Thus, jet quenching is very challenging to measure in small systems. In order to measure this in pp, ALICE has made use of another prediction, namely that the number of recoil jets at moderately high p_{T} will get reduced if there is jet quenching. To quantify this, the parameter

$$\Delta_{\mathrm{recoil}}(p_{\mathrm{T}}^{\mathrm{reco}}, \Delta\varphi) = \frac{1}{N_{\mathrm{trig}}} \frac{\mathrm{d}^2 N_{\mathrm{jet}}}{\mathrm{d}p_{\mathrm{T}}^{\mathrm{reco}} \mathrm{d}\Delta\varphi}\bigg|_{p_{\mathrm{T}}^{\mathrm{trig}} \in \mathrm{TT}_{\mathrm{sig}}} - c_{\mathrm{ref}} \cdot \frac{1}{N_{\mathrm{trig}}} \frac{\mathrm{d}^2 N_{\mathrm{jet}}}{\mathrm{d}p_{\mathrm{T}}^{\mathrm{reco}} \mathrm{d}\Delta\varphi}\bigg|_{p_{\mathrm{T}}^{\mathrm{trig}} \in \mathrm{TT}_{\mathrm{ref}}} \tag{3.5}$$

has been defined, where N_{trig} and N_{jet} are the number of trigger and recoil jets, respectively, $p_{\mathrm{T}}^{\mathrm{trig}}$ and $p_{\mathrm{T}}^{\mathrm{reco}}$ are the transverse momentum of the trigger and recoil jet, respectively, and $\Delta\varphi$ is the relative distance in azimuthal angle. $\mathrm{TT}_{\mathrm{sig}}$ and $\mathrm{TT}_{\mathrm{ref}}$ are

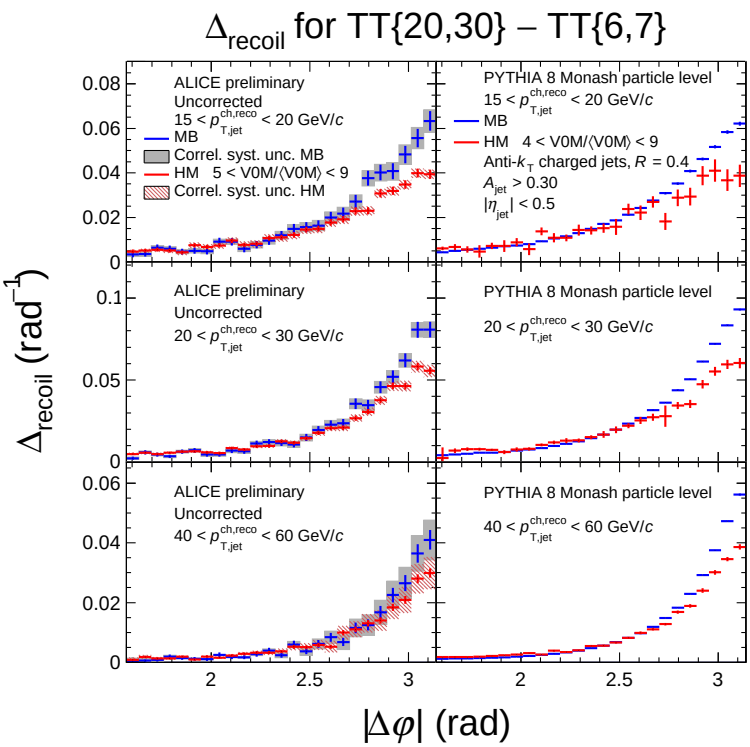

Figure 3.8 – Measured Δ_{recoil}, defined by Eq. (3.5), for MB and HM events in different $p_{\mathrm{T}}^{\mathrm{reco}}$ intervals for **left**: ALICE data and **right**: simulations by PYTHIA8 Monash. Figure taken from Ref. [43].

the signal and reference regions, respectively, where $\mathrm{TT}_{\mathrm{ref}}$ is put in a very different p_{T} interval to estimate the combinatorial background; thus a normalisation factor c_{ref} has been applied to make the two regions compatible. The ALICE analysis was done in the jet-p_{T} intervals $20 < p_{\mathrm{T}}^{\mathrm{trig}} < 30\,\mathrm{GeV}/c$ for the signal and $6 < p_{\mathrm{T}}^{\mathrm{trig}} < 7\,\mathrm{GeV}/c$ for the background (denoted $\mathrm{TT}\{20, 30\}$ and $TT\{7, 6\}$, respectively), where a comparison was made between a high-multiplicity (HM) sample and minimum-bias (MB) events [43]. The results show that there is indeed a suppression of back-to-back jets in the HM sample with respect to the MB sample, as shown in Fig. 3.8, which indicates that jet quenching indeed may have occurred. However, when comparing to PYTHIA using a similar sample (the HM sample is here defined differently, so a quantitative comparison is not possible), the same effect is observed also here. PYTHIA is a pure QCD event generator with no medium effects (cf. Section 3.8.2), so this is clearly due to something else. Therefore, the ALICE results are likely due to some measurement bias, and there is no conclusive evidence for jet quenching in pp collisions.

So far, jet quenching has not been observed in any other small systems either, but given the weak signal and challenge to obtain an unbiased result, this is perhaps not too surprising and so far does definitely not disprove its existence.

3.5 Heavy-Quarkonia Melting

Another expected signature of the QGP is that quarkonium states that exist in the medium will start to dissolve – or melt – in the QGP, which in particular is expected for states with low binding energy. This occurs because the long-range confining part of the QCD potential in Eq. (2.3) is screened in the QGP. Such states can form e.g. through hard scatterings of partons, i.e. the same process which forms jets. This melting is particularly evident when studying bottomonium states ($\mathrm{b\overline{b}}$) in Pb–Pb collisions at LHC energies, as seen in Fig. 3.9. The $\Upsilon(1S)$ state has a higher binding energy than the excited states and will thus not be suppressed to the same degree as $\Upsilon(2S)$ and $\Upsilon(3S)$, when comparing to pp collisions at the same energy. Therefore, the suppression increases from 62% in $\Upsilon(1S)$ to at least 90% (95% confidence limit) for $\Upsilon(3S)$ [44].

Given the scaling with binding energy, this melting can in principle be used as another probe for the QGP temperature. In practice, the situation is however more complicated. In particular, there might be other processes in the QGP affecting the production of hadrons in various ways. This is particularly evident when studying J/ψ ($\mathrm{c\overline{c}}$) production in heavy ion collisions, which was the first heavy-quarkonium state to be proposed as a probe for the QGP [45]. This has a lower binding energy than the excited bottomonium states, so naïvely one would expect a stronger suppression. As seen in Fig. 3.10a, this is not really the case, with a higher J/ψ R_{AA} at the LHC than at RHIC [46, 47]. These results can be explained as a recombination effect – due to the higher temperature at the LHC, more charm quarks are produced in the QGP, and thus there is a large probability for charm quarks produced in independent processes to recombine into J/ψ during freeze-out. The p_{T} dependence of R_{AA} (Fig. 3.10b) observed at LHC energies fits into this picture, with a smaller suppression of J/ψ at low p_{T}, where most of the recombination is expected to occur. At RHIC energies, the suppression is uniform, as expected from the lower number of charm quarks produced in the QGP at this temperature.

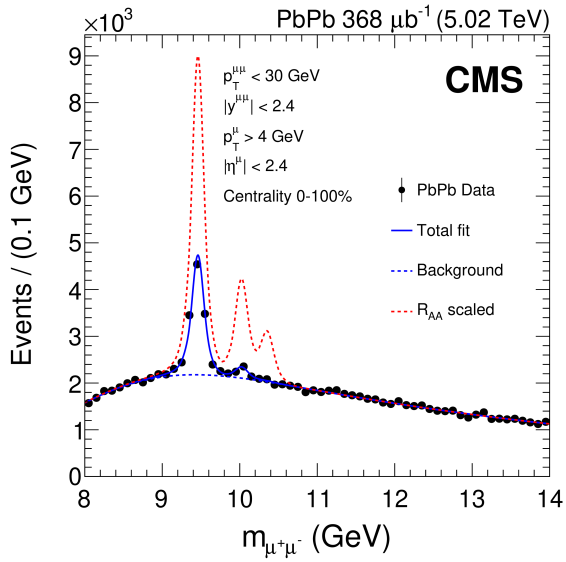

Figure 3.9 – Invariant-mass spectrum in the dimuon channel in the vicinity of the three $\Upsilon(n\mathrm{S})$ states, measured in Pb–Pb collisions at $\sqrt{s_{\mathrm{NN}}} = 5.02\,\mathrm{TeV}$. The dashed line is a scaling of the spectrum with R_{AA}, i.e. what would be expected from pp collisions with the same background statistics. Figure taken from Ref. [44].

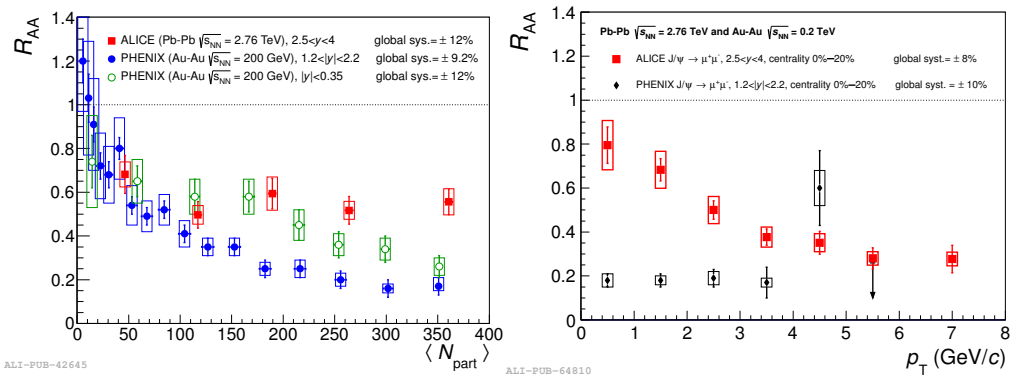

(a) **(b)**

Figure 3.10 – Comparisons of J/ψ R_{AA} (defined in Eq. (3.3)) between Pb–Pb collisions at $\sqrt{s_{\mathrm{NN}}} = 2.76\,\mathrm{TeV}$ measured by ALICE and Au–Au collisions at 200 GeV measured by PHENIX, as a function of **(a)** charged multiplicity and **(b)** p_{T}. Figures taken from Refs. [46] and [47].

In small systems, yields of J/ψ and the excited state $\psi(2\mathrm{S})$ have been measured as a function of multiplicity in p–Pb collisions, as is shown in Fig. 3.11. While the J/ψ results are reasonably well understood from scattering effects in the cold nuclear

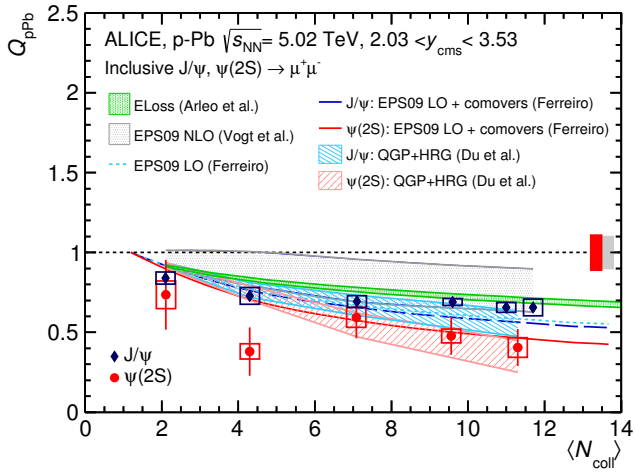

Figure 3.11 – Q_{pPb} (p–Pb analogy to R_{AA} defined in Eq. (3.3)) of J/ψ and $\psi(2\mathrm{S})$, measured as a function of multiplicity at $\sqrt{s_{\mathrm{NN}}} = 5.02\,\mathrm{TeV}$, along with a few model predictions. Figure taken from Ref. [48].

matter, the same cannot be said about the suppression of $\psi(2\mathrm{S})$, which would require some additional mechanism [48]. One possibility which matches the results is that a QGP has formed, and the suppression is due to interactions with it and the subsequent HRG following kinetic freeze-out.

Measurements of J/ψ yields have also been carried out in pp collisions, but since these results are particularly relevant when comparing with predictions by the models discussed in Section 3.8.5, they are presented there.

3.6 Strangeness Enhancement

Just as the number of charm quarks increases in the nuclear medium as compared to baryonic matter, the same effect happens also for strange quarks (although the underlying mechanism is different), where it is even more prominent. This effect is known as *strangeness enhancement*, and manifests itself experimentally through an increase in the relative yields of strange hadrons in heavy-ion collisions, as compared to the pp reference. This feature is particularly prominent for multistrange baryons, i.e. Ξ (two strange quarks) and Ω (three strange quarks).

Strangeness enhancement was proposed in 1982 as a signature for detecting a QGP [49], with the explanation that the strange-quark production is enhanced in the QGP since this only depends on the masses of the constituent quarks (as opposed to a hadron gas, where the production would be suppressed since one would need to take into account the confinement mass of kaons as well[5]). One can also imagine that this production may

[5]With the charged-kaon mass being 494 MeV, and the pion mass being 140 MeV, one can estimate the confined-strange quark mass as $\sim 420\mathrm{MeV}/c^2$, while the constituent-strange quark mass is measured to $\sim 96\,\mathrm{MeV}/c^2$ [11].

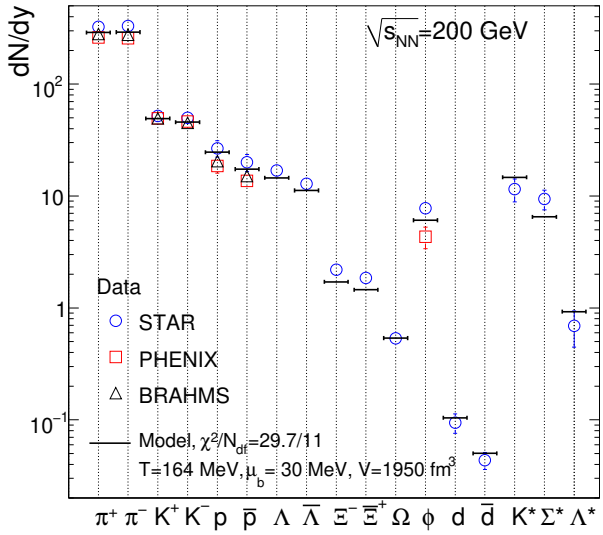

Figure 3.12 – Measured abundances of a selection of light hadrons and resonances (of which most are listed in Appendix C) at RHIC, along with predictions from a thermal model, where the freeze-out temperature ($T_{\mathrm{ch}} = 164\,\mathrm{MeV}$) has been fitted to data. Note the logarithmic scale. Figure taken from Ref. [50].

be altered by the hadronisation phase, but one result indicates that this should play a minor role for the observed strangeness enhancement, namely that the yields of light hadrons (i.e. hadrons consisting of u, d, and s quarks) measured in heavy-ion collisions at the SPS and RHIC are accurately predicted over several orders of magnitude by the thermal model [50], as is shown in Fig. 3.12 for RHIC data. This strongly supports that an equilibrated QGP is formed in these collisions. Another success of this model is that the yields of multistrange baryons saturate at high multiplicity, which is an indication that the medium has become fully thermalised.

What perhaps does not fit so well in the thermal picture, is that a similar enhancement is observed also in p–Pb and high-multiplicity pp collisions. Comparing the different systems, there is even a smooth scaling with multiplicity, which does not seem to depend on which system is involved [51], as is shown in Fig. 3.13. This suggests a common origin, and several attempts have been done to develop models for explaining this, as will be discussed in Section 3.8.5. In a pure thermal picture a similar scaling is however also possible from canonical suppression, which is a mechanism that suppresses strangeness production in smaller QGP volumes.

The purpose of the main analysis in this thesis is to use the charged Ξ baryon (selected since it is much more abundant than $\Omega^{\pm}$) to pinpoint the production mechanism of strange quarks and hadrons in small systems, and hence the origin of strangeness enhancement. This is further described in Chapter 8 and onwards.

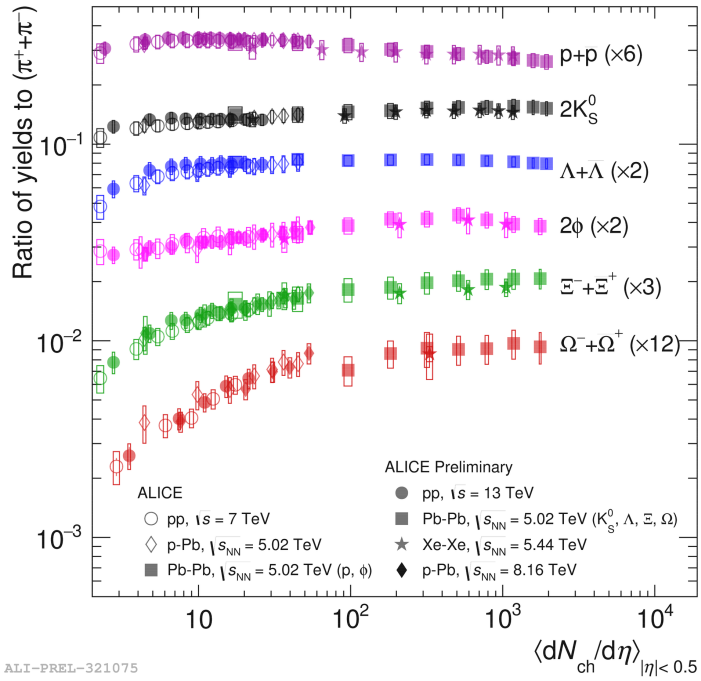

Figure 3.13 – Yield ratios between various hadron species and pions as a function of multiplicity, for pp, p–Pb, Pb–Pb, and Xe–Xe collisions, measured by the ALICE detector. Figure taken from Ref. [52].

3.7 Collective Flow

So far, I have mainly discussed how the chemical composition and distribution of hadrons are affected by the QGP medium, but these quantities will also be affected by the expansion of the medium itself. Such collective effects are known as *flow* and have both radial and azimuthal components. I will start with going through the signatures of radial flow, before moving on to the azimuthal behaviour.

3.7.1 Radial Flow

An important effect of the radial expansion of the medium is that this will give the created hadrons a boost in radial velocity. As the total momentum is

$$\mathbf{p} = \frac{m\mathbf{v}}{\sqrt{1 - |\mathbf{v}|^2/c^2}},$$

where $\mathbf{v}$ is the velocity and m the mass, this will result in a scaling in p_{T} with mass. Thus there is a depletion of heavier particles at low momentum but an enhancement at intermediate p_{T}, as can be very clearly seen when comparing the proton R_{AA} to other particles in Fig. 3.7.

Another consequence of radial flow is that the *(charge) balance function* gets narrower with increasing radial flow [53]. This is defined as

$$\mathbb{B}(\Delta y, \Delta\varphi) = \frac{1}{2}\left(\mathbb{C}_{(+,-)}(\Delta y, \Delta\varphi) + \mathbb{C}_{(-,+)}(\Delta y, \Delta\varphi) - \mathbb{C}_{(+,+)}(\Delta y, \Delta\varphi) - \mathbb{C}_{(-,-)}(\Delta y, \Delta\varphi)\right), \tag{3.6}$$

where $\mathbb{C}(\Delta y, \Delta\varphi)$ is the two-particle correlation function, i.e. the distribution of particle pairs in relative rapidity (cf. Section 8.2) and azimuthal angle, and the subscripts denote the signs of the particles; see details on how to measure this in Section 8.3. This is a measure of the distribution of balancing charges in the event (the subtraction is done to remove non-charge dependent correlations), which is expected to be a localised phenomenon due to local charge conservation. Therefore, as the medium expands, the balancing happens at a larger distance from the collision, resulting in a reduced solid angle of the balance function.

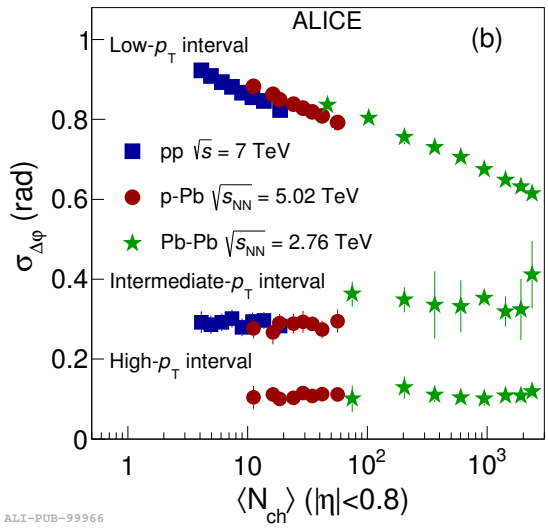

Figure 3.14 – Balance function width, projected in $\Delta\varphi$, for Pb–Pb, p–Pb, and pp collisions measured by ALICE for pairs of low- ($0.2 < p_{\mathrm{T}}^{\mathrm{assoc}} < p_{\mathrm{T}}^{\mathrm{trig}} < 2.0\,\mathrm{GeV}/c$), intermediate- ($2.0 < p_{\mathrm{T}}^{\mathrm{assoc}} < 3.0 < p_{\mathrm{T}}^{\mathrm{trig}} < 4.0\,\mathrm{GeV}/c$), and high-$p_{\mathrm{T}}$ ($3.0 < p_{\mathrm{T}}^{\mathrm{assoc}} < 8.0 < p_{\mathrm{T}}^{\mathrm{trig}} < 15.0\,\mathrm{GeV}/c$) particles, respectively. Figure taken from Ref. [54].

Experimentally, a narrowing of the balance function with increased multiplicity has been observed both in Pb–Pb, p–Pb, and pp collisions, as is shown in Fig. 3.14, but only between pairs of low-p_{T} hadrons [54], which indicates that this only affects particles in the underlying event. When going to high p_{T}, the particle spectra will be dominated by hard processes, i.e. jets, and then the charge balance will occur within the jet cone and should not be affected by the expansion of the medium. This suggests that the observed narrowing is indeed due to radial flow, and that there is flow also in small systems.

3.7.2 Anisotropic Flow

Any imbalance in the azimuthal distribution of produced particles which is due to flow of the expanding medium, is called *anisotropic flow*. This is an inevitable effect of an expanding medium starting from a non-uniform initial state, i.e. what will happen in a non-central collision. The system will then expand along the pressure gradient, which is strongest in the reaction plane, i.e. the plane cutting through the centre of the colliding nuclei along the beam direction, as is shown in Fig. 3.15. As the medium expands, the pressure gradient will continue to flow along the reaction plane, leading to an elliptic distribution of the final-state hadrons. This flow component is known as *elliptic flow*, but for a complete description of the expanding medium, one can make a Fourier decomposition of the particle distribution in azimuthal angle with respect to the symmetry planes Ψ_n:

$$\frac{\mathrm{d}N}{\mathrm{d}\varphi} = \frac{1}{2\pi}\left(1 + \sum_{\mathrm{n}=1}^{\infty} 2v_\mathrm{n}\cos\left(\mathrm{n}\left(\varphi - \Psi_\mathrm{n}\right)\right)\right), \tag{3.7}$$

where v_n are Fourier coefficients.

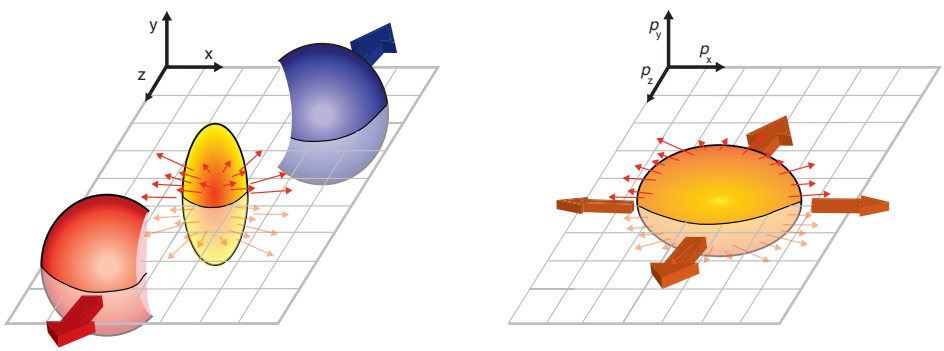

Figure 3.15 – Illustration of the reaction plane and elliptic flow. The pressure gradient is the strongest in the reaction plane, and thus the medium expands along this plane. The left panel shows the initial state, whereas the right panel shows the system when it has expanded for a while. Figures created by Boris Hippolyte, based on originals by Masashi Kaneta.

Here, elliptic flow is described by v_2, but there are also other components. The lowest-order component, v_1 is called directed flow, but since this is not a signature of an expanding medium, it will not be discussed here. Of the higher-order components, the dominant one is v_3, called *triangular flow*, which along with higher v_n is due to more complicated variations in the pressure gradient, as is illustrated in Fig. 3.16. All $v_\mathrm{n}, \mathrm{n} \geq 2$, are also affected by fluctuations in the initial state, and thus there is flow also in central heavy-ion collisions.

The anisotropic flow is related to the initial geometry through the relation [56]

$$v_\mathrm{n} \approx \kappa_\mathrm{n}\varepsilon_\mathrm{n}, \tag{3.8}$$

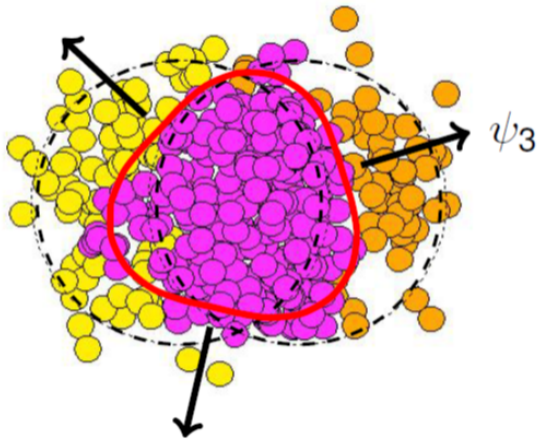

Figure 3.16 – Illustration of how anisotropies in the initial geometry can give rise to higher-order v_n (here only the triangular geometry component is shown), using input from a Glauber model. As the medium expands, these flow components still remain. Participants are drawn in magenta, whereas yellow and orange are used for the spectators. Figure adapted by Burak Alver, based on the original in Ref. [55].

where ε_n is the eccentricity, and the proportionality constant κ_n for $\mathrm{n} = 2$ and $\mathrm{n} = 3$ is expected to scale with the transverse charged-particle density,

$$\frac{1}{S}\frac{\mathrm{d}N_\mathrm{ch}}{\mathrm{d}\eta},$$

where $S = 4\pi\sigma_x\sigma_y$ is the transverse area and $\mathrm{d}N_\mathrm{ch}/\mathrm{d}\eta$ is the charged particle density of the system. The quantities σ_x and σ_y are the standard deviations of the source (parton or nucleus, depending on which model is used) distribution in the plane transverse to the reaction plane. These are highly dependent on the initial-state model, and therefore measurements of anisotropic flow can be used to test this, as is done in Chapter 7.

While it is possible to calculate the symmetry plane, this has a large risk of introducing a measurement bias. To avoid this, the flow coefficients are calculated through correlation factors between all combinations of two or more particles in the events. One can show that if the azimuthal distribution of particles follows the geometry of the medium at chemical freeze-out, then it holds that [57]

$$v_\mathrm{n}^2 = \langle\langle\cos(\mathrm{n}(\varphi_1 - \varphi_2))\rangle\rangle\,, \tag{3.9}$$

where the average is taken over all combinations of track pairs in an event, which in turn is averaged over all events. Evaluating this yields the 2-particle correlator,

$$\langle 2\rangle \equiv \langle\cos(\mathrm{n}(\varphi_1 - \varphi_2))\rangle = \frac{1}{M(M-1)}\sum_{i=1}^{M}\sum_{j=1,j\neq i}^{M}\cos(\mathrm{n}(\varphi_i - \varphi_j)), \tag{3.10}$$

where M is the particle multiplicity in the event. Defining the 2-particle *cumulant* as

$$c_\mathrm{n}\{2\} = \langle\langle 2\rangle\rangle\,, \tag{3.11}$$

then gives the expression

$$v_{\mathrm{n}}\{2\} = \sqrt{c_{\mathrm{n}}\{2\}}, \tag{3.12}$$

where the "2" is used to emphasise that this v_{n} is calculated from the 2-particle cumulant.

There is however one problem with this approach, and that is the assumptions leading to Eq. (3.9). In a real event, contributions from jet fragmentation and weak decays will also enter into the calculation of $\langle 2 \rangle$. This is known as *non-flow*, and needs to be suppressed in order to use cumulants for flow measurements. There are two ways to do this: either one can introduce a pseudorapidity gap and only use pairs between particles at either side of the gap in the calculation of $v_{\mathrm{n}}\{2\}$, which is known as the *scalar product method* [58], or one can include more particles when calculating the flow coefficients. In this way, the 4-particle correlator is defined as

$$\begin{aligned}\langle 4 \rangle \equiv \langle \cos(\mathrm{n}(\varphi_1 + \varphi_2 - \varphi_3 - \varphi_4)) \rangle = \frac{1}{M(M-1)(M-2)(M-3)} \\ \cdot \sum_{i=1}^{M} \sum_{j=1, j\neq i}^{M} \sum_{k=1, k\neq i,j}^{M} \sum_{l=1, l\neq i,j,k}^{M} \cos(\mathrm{n}(\varphi_i + \varphi_j - \varphi_k - \varphi_l)).\end{aligned} \tag{3.13}$$

Now defining the 4-particle cumulant as

$$c_{\mathrm{n}}\{4\} = \langle \langle 4 \rangle \rangle - 2 \cdot \langle \langle 2 \rangle \rangle^2, \tag{3.14}$$

leads to the estimator [59]

$$v_{\mathrm{n}}\{4\} = (-c_{\mathrm{n}}\{4\})^{1/4}, \tag{3.15}$$

which is less sensitive to non-flow than $v_{\mathrm{n}}\{2\}$, even with a rapidity gap.

Trying to evaluate Eq. (3.13) directly would be extremely inefficient, so in practice this is calculated by decomposing the summation into a combination of flow vectors

$$Q_{\mathrm{n}} = \sum_{k=1}^{M} \exp(\mathrm{i}\mathrm{n}\varphi_k),$$

which are calculated iteratively through the *generic framework*, which also allows extensions to higher-order cumulants [60], which sometimes can be useful when comparing results to model predictions. This results for the lowest-order cumulants are given in Section 7.2.

While $v_{\mathrm{n}}\{4\}$ and the scalar product-$v_{\mathrm{n}}\{2\}$ in their own give non-flow suppressed estimates of v_{n} – where the latter should be used only if the statistics are low – combining the two estimators can be used to estimate the contribution to v_n from flow fluctuations, $\sigma_{v_{\mathrm{n}}}$. It has been shown that [61]

$$\begin{aligned} v_{\mathrm{n}}\{2\} &\approx \langle v_{\mathrm{n}} \rangle + \frac{\sigma^2_{v_{\mathrm{n}}}}{2 \langle v_{\mathrm{n}} \rangle}, \\ v_{\mathrm{n}}\{4\} &\approx \langle v_{\mathrm{n}} \rangle - \frac{\sigma^2_{v_{\mathrm{n}}}}{2 \langle v_{\mathrm{n}} \rangle}, \end{aligned} \tag{3.16}$$

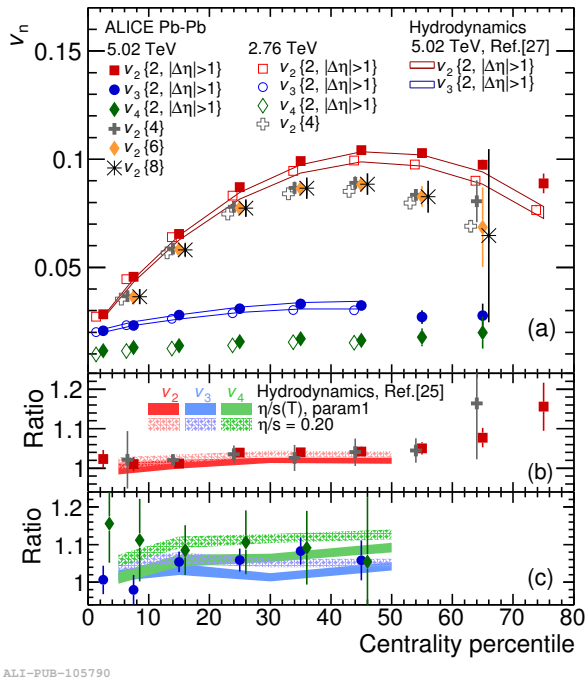

Figure 3.17 – Anisotropic flow v_n measured by ALICE in Pb–Pb collisions at $\sqrt{s_{NN}} = 5.02$ and 2.76 TeV, respectively. The lower panels show ratios between results from the two collision energies, along with hydrodynamical predictions from Refs. [63] and [64] (note that the references in the figure panels do not follow the enumeration used in this thesis). Figure taken from Ref. [62].

so

$$\sigma_{v_n} \approx \sqrt{\frac{(v_n\{2\})^2 - (v_n\{4\})^2}{2}}.$$

Experimental results for $v_2\{m\}$, $m \leq 8$, $v_3\{2\}$, and $v_4\{2\}$ as a function of multiplicity are shown in Fig. 3.17 [62]. The results – including the energy dependence – agree with a hydrodynamical description of the QGP. The elliptic-flow component increases when going from central to semi-peripheral collisions, but decreases again in peripheral collisions. Both of these trends are expected from geometry considerations, and $v_2\{4\} < v_2\{2\}$ as expected from Eq. (3.16).

One can also measure flow for identified particles. This exhibits a division between baryons and mesons, with the flow of baryons being pushed toward higher p_T. This roughly scales with the number of constituent quarks, at least at intermediate p_T, as is shown in Fig. 3.18 [65]. This indicates that the flow follows the underlying parton distribution, rather than that of the hadrons, i.e. it is the partons that flow in the medium. This model is called the *quark coalescence* model [66]. In data there is however a deviation by about 20% from this scaling, so the measured flow is likely affected by other processes as well.

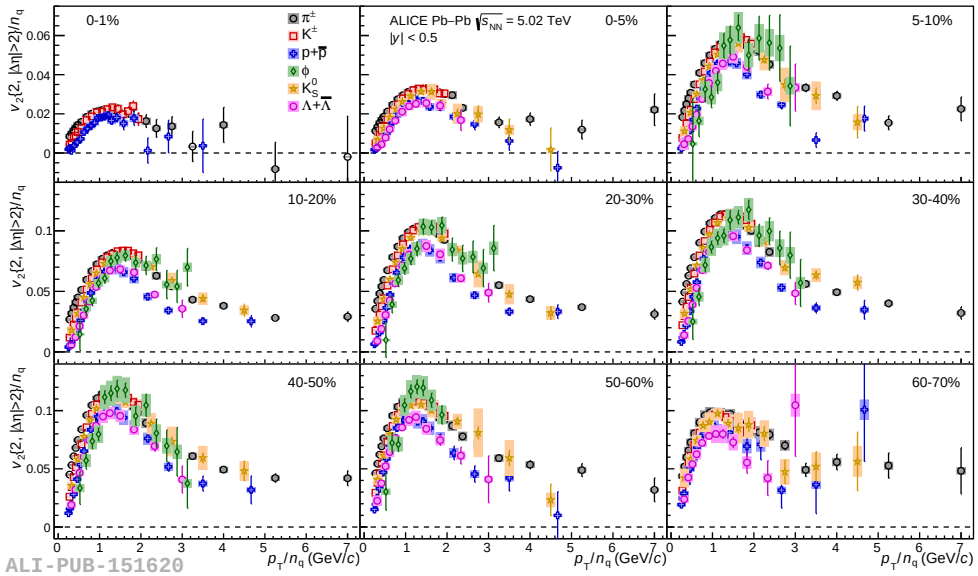

Figure 3.18 – Elliptic flow in Pb–Pb collisions at $\sqrt{s_{\mathrm{NN}}} = 5.02\,\mathrm{TeV}$, measured for identified particles as a function of $p_{\mathrm{T}}/n_{\mathrm{q}}$, where n_{q} is the number of constituent quarks, and divided into centrality classes. Figure taken from Ref. [65].

Anisotropic Flow in Small Systems

Just as radial flow, anisotropic flow has been observed in small systems, including collisions between protons – which are spherically symmetric – ruling out that this is caused by the geometry of the colliding nuclei. Moreover, the p_{T} dependence of v_2, v_3, and v_4 in both small and large systems is well reproduced by hydrodynamical models fitted to all system sizes *simultaneously* [67], as shown in Fig. 3.19. This strongly suggests that the origin is hydrodynamical, i.e. that a QGP droplet has been formed

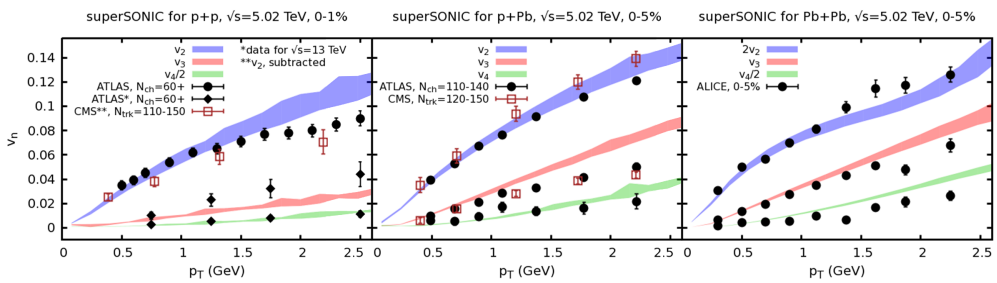

Figure 3.19 – Measured v_2, v_3, and v_4 at the LHC in **(left)** pp collisions, **(middle)** p–Pb collisions, and **(right)** Pb–Pb collisions at $\sqrt{s_{\mathrm{NN}}} = 5.02\,\mathrm{TeV}$, along with predictions by the superSONIC hydro model. Figure taken from Ref. [67].

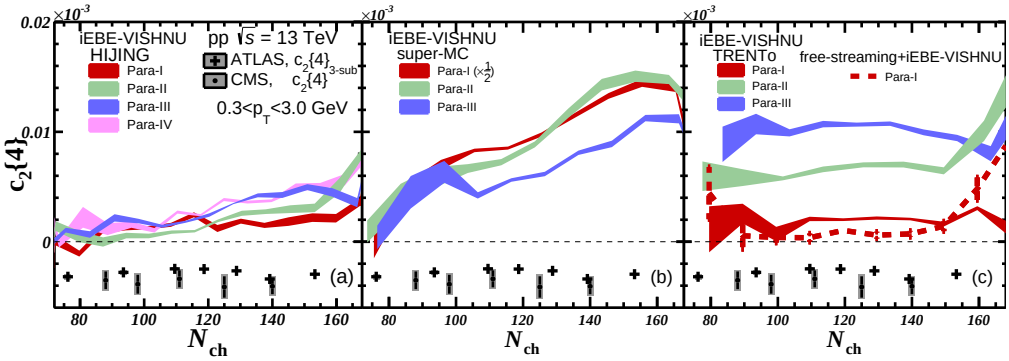

Figure 3.20 – Extracted $c_2\{4\}$ as a function of collision multiplicity in 13 TeV pp collisions for the iEBE-VISHNU hydrodynamic model, using fit parameters reproducing the observed v_n for three different initial-state models using several different parametrisations, compared with measurements from ATLAS and CMS. Figure taken from Ref. [68].

in small systems.

Recently however, this picture has been challenged. In particular it has been shown that a hydrodynamical model tuned such that it reproduces v_n, $2 \leq n \leq 4$, results in a positive $c_2\{4\}$ (Eq. (3.14)) in pp collisions, regardless of initial-state model [68], contrary to what is observed in data, as is shown in Fig. 3.20. In practice this means that the flow fluctuations are strongly overestimated, meaning that the mechanism generating flow in small systems in data might be different than according to these models.

3.8 Phenomenological Models

3.8.1 Overview

To summarise the results presented in Sections 3.3–3.7, modification of heavy-quarkonia yields, strangeness enhancement, and radial and anisotropic flow have been observed in both heavy-ion collisions and small systems, indicating that also the latter exhibit collective behaviour. Some of the results also indicate that there is a common origin, although this hypothesis is challenged by others. Moreover, there have neither been any signs of jet quenching nor thermal photons in small systems, although these signatures are clearly seen in heavy-ion collisions. These observables are however both challenging to measure and suffer from a loss in statistics, meaning that one cannot draw any final conclusions yet from the missing signatures in small systems.

Without enough evidence for either picture, it is not yet possible to conclude whether the origin of collective effects is the same or different in the two systems, or more specifically if the small-system effects are due to formation of a medium or if there are other effects which mimic this behaviour. In this thesis, I will focus on two approaches aiming to find a unified picture when going from small to large system, starting out from each of the two pictures. These are the *core-corona model* and the *Lund String Model*.

The core-corona model assumes that a dense core, behaving like a medium (i.e. a QGP) is formed in the centre of the interaction, and a more diluted state of QCD matter where thermal equilibrium has not been reached, called the corona, is formed in the peripheral regions. As the number of participants increases, the core fraction increases, resulting in a smoothly increasing collectiveness with system size and multiplicity. The core-corona model which currently has reached the latest stage in its development is the EPOS model, and therefore this is the one tested in this thesis. This is described in Section 3.8.4. The newest public release of EPOS is EPOS LHC [69], but a more sophisticated model – EPOS3 [70] – is under development.

The Lund String Model is a hadronisation model, i.e. the underlying physics is in the hadronic phase of QCD. In this model, the gluon field is represented by strings which break if reaching high enough energy – forming hadrons [71]. While there are other hadronisation models, such as Herwig (based on clusters) [72] and PHOJET (based on Pomeron fragmentation) [73], the Lund string model with the event generator PYTHIA [74] has been the most successful for reproducing many results in data. Until recently, PYTHIA could only describe pp collisions, but now there is ongoing work to extend this to heavy-ion collisions, which are described by the Angantyr extension [75]. The idea behind this framework is that the collective behaviour observed in both small systems and heavy-ion collisions is due to interactions between strings, i.e. there is no QGP. Thus, it also suggests a common origin for the collective effects. So far, Angantyr has not been very successful in describing most heavy-ion observables, but it has made some predictions in small systems which are being tested in this thesis. The Lund string model, PYTHIA, and Angantyr are described in Sections 3.8.2 and 3.8.3.

3.8.2 The Lund String Model

The basic idea behind the string model is that the proportionality constant κ in the linear term of the QCD potential described by Eq. (2.3), can be interpreted as a string tension between two quarks confined into a hadron [71]. Now, if the quarks start moving away from each other, e.g. as the result of a collision, their kinetic energy will be converted into potential energy. If the string reaches high enough energy, there is a chance that a $\mathrm{q\bar{q}}$ pair will form along the string using a mechanism similar to tunnelling – and the string breaks. In this way mesons are formed, but a diquark pair may also form, resulting in baryon formation. I will now briefly describe the string evolution in the centre-of-mass-frame. Although this can fairly easily be converted to other Lorentz frames, it makes the result less intuitive.

The usual way to describe the evolution of the string is to draw a space-time diagram as the one shown in Fig. 3.21a. At time $t = 0$, the quarks are close to each other, and have a total kinetic energy $E_\mathrm{k} = \sqrt{s}$ and a potential energy $E_\mathrm{p} = 0$. As the quarks move apart, the kinetic energy will eventually be fully converted into potential energy $E_\mathrm{p} = \kappa\Delta x$ at some distance $\Delta x \approx 2ct$ in the centre-of-mass frame, i.e. this happens at $t = \sqrt{s}/(2c\kappa)$. At that time, the quarks will again start moving towards each other, and at $t = \sqrt{s}/(c\kappa)$ they will meet again, but moving in opposite direction. Through the propagation of time, the quarks follow a yo-yo motion with period $\tau = 2\sqrt{s}/(c\kappa)$.

During this motion, when E_p is high enough that it is energetically possible, a $\mathrm{q\bar{q}}$ or $\mathrm{qq\bar{q}\bar{q}}$ pair may form along the string. Integrating over the full motion gives the

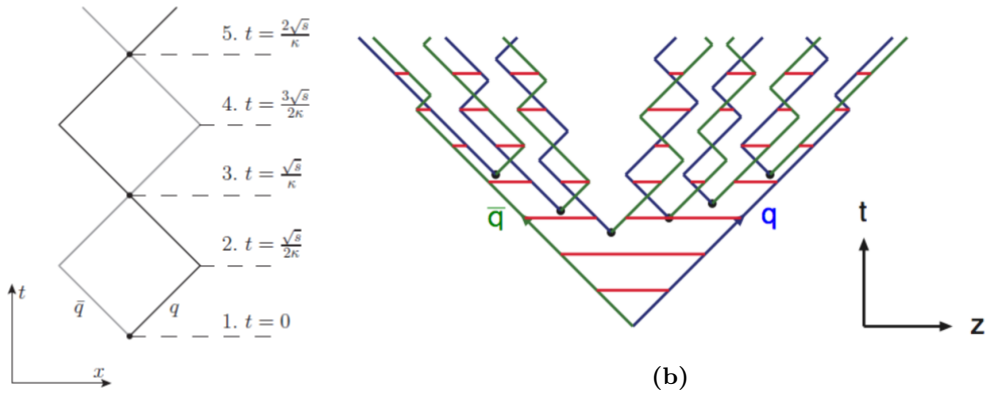

Figure 3.21 – **(a)** Space-time diagram of the evolution of a string formed from a $q\bar{q}$ pair in the centre-of-mass frame, with the quark drawn in black and the antiquark in grey. Different times are marked along the trajectory. **(b)** Space-time diagram of string fragmentation, where quarks are drawn in blue and antiquarks in green. Figure taken from Refs. [82] and [83].

following relation for the probability for meson formation:

$$\frac{\mathrm{d}P}{\mathrm{d}m_\mathrm{T}} \propto \exp\left(-\frac{\pi(m_\mathrm{T}^2}{\kappa}\right),$$

where $m_\mathrm{T} = \sqrt{m^2 + p_\mathrm{T}^2}$ is the transverse energy. This can be interpreted as a tunnelling probability, since the quarks need to be displaced by some distance $\Delta = 2m_\mathrm{T}c^2/\kappa$ in order for the string to reach enough potential energy. From this equation follows that the probability of forming a strange quark is suppressed by a factor

$$\rho = \exp\left(-\frac{\pi(m_\mathrm{s}^2 - m_\mathrm{u,d}^2)}{\kappa}\right), \tag{3.17}$$

where $m_\mathrm{u,d}$ is the typical mass of either an up or a down quark. It is however not clear which mass should be used for the strange quark (cf. Section 3.6), so this is one of many parameters in the string model that needs to be tuned to data, preferably from e^+e^- collisions to avoid contamination from interactions with other partons. A fit to LEP data yields $\rho \simeq 0.217$ [84]. For heavy quarks, such as charm, the probability of formation is essentially zero in this picture, and thus these are formed in other processes and mostly at high energy, i.e. in the perturbative regime of QCD.

If the initial energy is high enough, there may occur multiple breakings of the string, i.e. the string *fragments*, resulting in multiple string. Just as for the single-string picture, string fragmentation can conveniently be drawn in a space-time diagram, as is shown in Fig. 3.21b. In this picture, the area of the rectangle enclosed by each half-revolution of the yo-yo motion is equal to m_T^2, so this gives a description of the string energy.

Multiparton Interactions and Colour Reconnection

When colliding protons, the fragmentation picture becomes more complicated, since multiple partons from the colliding protons participate in the reaction. This leads to two phenomena called *multiparton interactions* and *colour reconnections*.

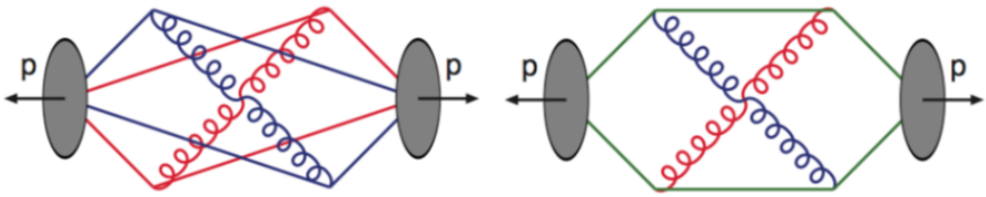

Figure 3.22 – Sketch of colour strings forming between partons in two colliding nuclei, with **(left)** CR disabled, and **(right)** CR enabled. In the latter case, fewer hadrons are created due to the (colour confined) strings being shorter. Figure taken from Ref. [85].

Multiparton interactions (MPIs) are simply processes where multiple partons participate. These are particularly important at small impact parameters (cf. Section 3.2), where they result in a larger number of produced hadrons, particularly at low p_{T}. This is now a standard feature included in several event generators, including PYTHIA [74], but if not including additional mechanisms, the p_{T} spectrum becomes too soft in high-multiplicity events when comparing to data. The reason is that following a collision, large colour fields will build up between partons in the beam remnants, as is schem-

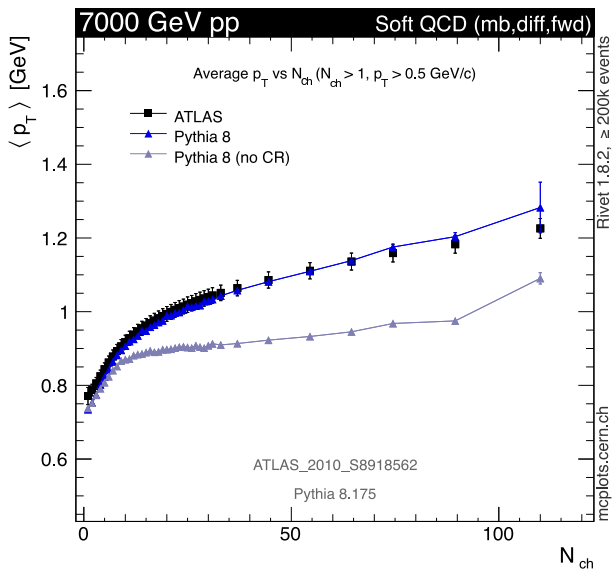

Figure 3.23 – Average p_{T} as a function of charged-particle multiplicity in pp collisions at $\sqrt{s} = 7$ TeV simulated by PYTHIA8 with and without colour reconnection enabled, compared with ATLAS data. Figure taken from Ref. [86].

atically shown in the left panel of Fig. 3.22. The number of hadrons created from this fragmentation scales almost linearly with string energy, resulting in an almost flat scaling of $\langle p_{\mathrm{T}} \rangle$ with multiplicity.

The solution to this is colour reconnections (CR), where additional partons from the beam remnants are allowed to reconnect with these strings, and in this way create colour neutral objects, as is schematically shown in the right panel of Fig. 3.22. In this way the strings become shorter and fewer charged particles are created per MPI, and as a consequence the average p_{T} per hadron increases in harder scatterings. With CR enabled, the average p_{T} as a function of multiplicity agrees very well with data [86], as is shown in Fig. 3.23. In PYTHIA8, colour reconnections are included in the default configuration.

String Junctions

In the standard configuration of PYTHIA, baryons are created through the formation of diquark pairs along the string, as is shown in Fig. 3.24a. In hadron–hadron collisions however, due to MPIs it is not unlikely that diquark states form from the beam remnants, i.e. there strings may form between pairs (or clusters) of quarks [87]. These strings will have two colours associated with them, allowing for more configurations during colour reconnections. One such possibility is that a junction is formed on the string. In this way, baryons may form by string breakings on either side of the junction, as is shown in Fig. 3.24b, i.e. this adds an additional mechanism for baryon formation, which is available as an extension to PYTHIA8.

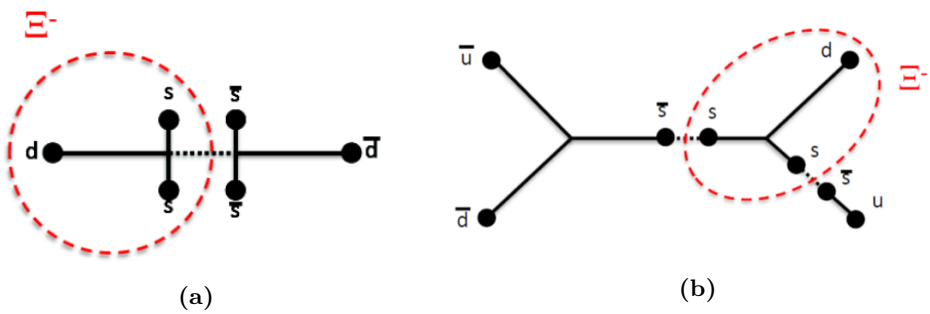

Figure 3.24 – **(a)** Mechanism for producing baryons in PYTHIA, using a pair of Ξ baryons as an example. The pair of solid lines is a string, which breaks through the formation of diquark pairs. **(b)** Additional mechanism added by the junction model. Here, the baryon can be formed from a junction between two strings, which breaks into a single $\mathrm{q\bar{q}}$ pair. Figures taken from Ref. [88].

3.8.3 The Angantyr Model

While PYTHIA does well in describing pp collisions, one piece is missing for it to be able to describe larger systems – even if assuming that the underlying physics is the same – namely a geometrical description of the collision. While it is possible to get a reasonable description of parameters such as parton density from the impact parameter

alone, this will not be enough for a large system such as a heavy-ion collision. If desiring to reproduce observables such as flow and strangeness enhancement, one also needs additional mechanisms.

All these requirements are included or being worked on in the Angantyr package [75], where the most important difference from PYTHIA is that strings now extend in space. Therefore, additional interactions are possible between overlapping strings. Two such mechanisms, which are currently implemented in Angantyr – and also enabled as extensions to PYTHIA8.2 – are *rope hadronisation* and *string shoving*. Moreover, the nuclear geometry is taken into account in Angantyr through a Glauber simulation (cf. Section 3.2.1), which describes the interactions with each nucleus.

Rope hadronisation is based on the same idea as the mechanism behind string junctions, namely that strings may form between parton clusters. Such a string is called a *colour rope* [78]. Depending on the colour configuration, the effective string tension then increases from κ to $\tilde{\kappa}$, which increases the production of strange quarks. Following Eq. (3.17), the new strangeness suppression factor then becomes

$$\tilde{\rho} = \rho^{\kappa/\tilde{\kappa}}. \tag{3.18}$$

A similar modification also occurs for the baryon yields.

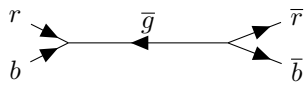

Figure 3.25 – Schematic drawing of how a string between two clusters of incoherent colours (here r and b) is formed when enabling colour ropes in PYTHIA. The two strings merge through the junction mechanism into a string with the combined colour, $rb = \overline{g}$.

In such a cluster of quarks, a subset of them may have the same colour, forming a coherent state. This would lead to an increased colour flow along the string, increasing the string tension. On the other hand, quarks having different colours combine into an anticolour (e.g. $r \oplus b = \overline{g}$), reducing the colour flow and hence the string tension. In PYTHIA, such a coupling is modelled through the string junction mechanism as shown in Fig. 3.25, and thus this mechanism is required to form ropes in PYTHIA. Consequently, only one string breaking may occur along this string. For generalised ropes, the number of breakings between clusters of n quarks is $\leq n$, with the maximum number occurring if they are all coherent.

The rope tension in terms of κ is given by the quadratic Casimir operator, leading to [78]

$$\frac{\tilde{\kappa}}{\kappa} = \frac{1}{4}\left(p^2 + q^2 + pq + 3(p+q)\right),$$

where p and q are the number of coherent colour and anti-colour states, respectively, in the cluster. For example, if forming a rope between a coherent $r \oplus r$ state and an $\overline{r} \oplus \overline{r}$ state, $p = 2$ and $q = 0$ (there are two coherent colours at each end of the rope, but we are only looking at one end), so $\tilde{\kappa} = 5/2\kappa$. This rope will break through two string breakings, reducing the tension to κ in the first breaking, so the *effective* tension of this breaking is $3/2\kappa$. Inserting this into Eq. (3.18) leads to a modification of the strangeness suppression factor from 0.217 to 0.361.

In Angantyr, ropes are formed when strings form close enough in space to overlap. Since PYTHIA works in the hadron gas regime of the QCD phase diagram, the effective string radius should be of the same order of magnitude as the confinement distance, i.e. $\sim 1\,\text{fm}$. An illustration of overlapping strings in a pp collision is shown in Fig. 3.26. When going from smaller to larger systems (or lower to higher multiplicities), the string density will increase, effectively increasing p and q, and hence the string tension. This is the reason for the predicted smooth increase in strange-baryon yields shown in Fig. 3.31.

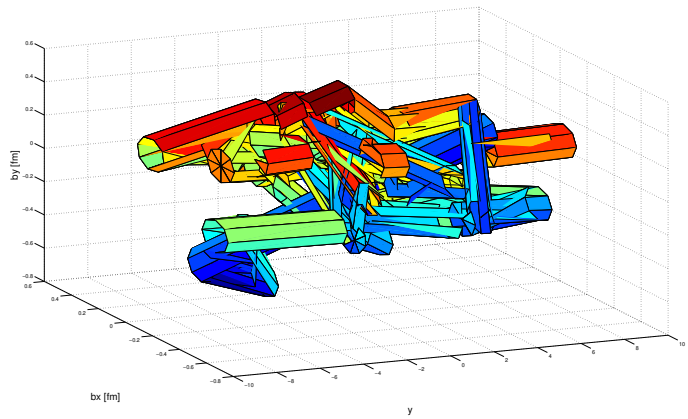

Figure 3.26 – Example of a snapshot of overlapping strings from a simulated pp event, where the colours represent the colour flow. Note that the string radius has been reduced to 0.1 fm for increased readability. Figure taken from Ref. [78].

While rope hadronisation results in enhanced strangeness yields, it is not able to add elliptic-flow like effects to the system. That requires string shoving, which is another effect of overlapping strings [89]. An effect of the enhanced string tension is that this increases the energy density and consequently the probability of radiating gluons. Such gluons may interact with the other strings in the overlap region, exerting a force on them. Thus, the strings will shove each other apart, as is shown in Fig. 3.27. If the shoving pressure is large enough, which depends on the string density, this will show up as flow among the produced hadrons.

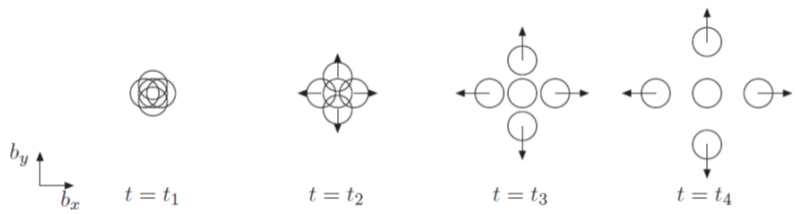

Figure 3.27 – Mechanism behind string shoving, showing snapshots of the evolution. Strings closely packed in space exert a pressure on each other, shoving them apart. Figure taken from Ref. [89].

3.8.4 The EPOS Model

The EPOS model is a core-corona model, meaning that the system formed is built up of a dense core part and a more dilute corona. The initial state is modelled by the formation of parton ladders (a representation of MPIs) which results in colour fields called *flux tubes* [69], schematically drawn in Fig. 3.28a, which are not very different from the strings in the Lund String Model. The core is formed in regions where the density of flux tubes is above a certain threshold, whereas the rest (typically confined to the peripheral regions of the collision system) make up the corona. The core part will then expand hydrodynamically and eventually hadronise through collective hadronisation, as is schematically shown in Fig. 3.28b, whereas the corona will fragment through string fragmentation.

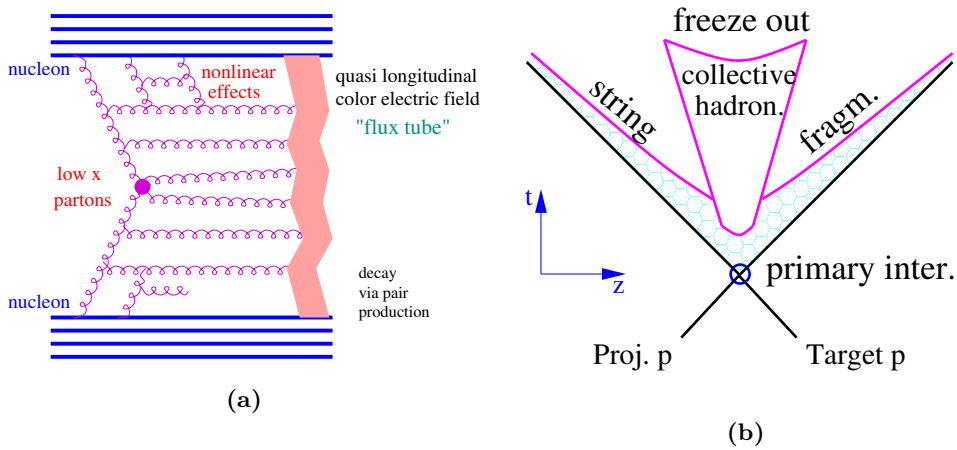

Figure 3.28 – **(a)** Schematical drawing of how a quasi-longitudinal colour field – called a flux tube – is formed between two interacting nucleons through MPIs. **(b)** Simplified space-time diagram of the evolution of a core-corona state. The central parts which form the core will expand and subsequently hadronise collectively, whereas the peripheral parts will decay through string fragmentation. Figures taken from Ref. [69].

In EPOS LHC [69], a full hydrodynamisation has not been carried out, unlike in EPOS3 (still in development) [70]. Instead a parametrisation is used to simulate flow, which works reasonably well for small systems, but should be used with care in heavy-ion collisions. Soft particles are strongly affected by the core, whereas harder particles survive to a larger extent and will fragment into jets. This is implemented in a way such that the angular separation from the leading particle very well reproduces what is seen in data, as is shown in Fig. 3.29, although there are other observables which are not equally well predicted by the model.

Hadronisation in EPOS LHC is implemented through a micro-canonical picture, where the phase space is divided into clusters which are sliced in η. In each cluster, the hadrons are drawn from a micro-canonical ensemble where the probability of drawing

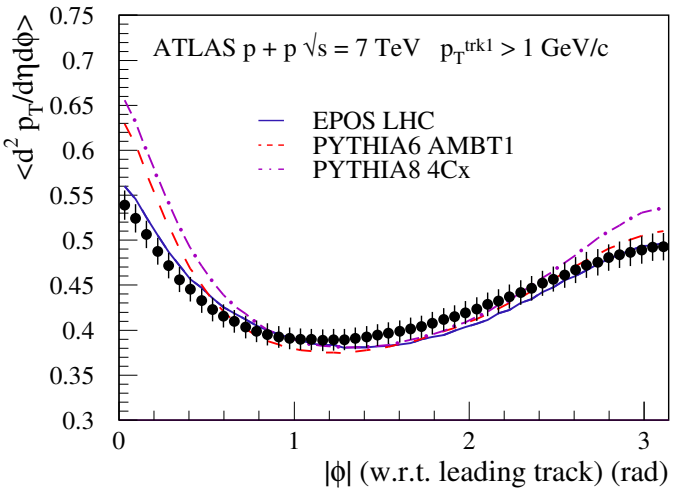

Figure 3.29 – Angular correlation in $|\Delta\varphi|$ between the leading hadron and other hadrons in the event in pp collisions at $\sqrt{s} = 7$ TeV for ATLAS data, compared with predictions by EPOS LHC, PYTHIA6, and PYTHIA8. Figure taken from Ref. [69].

n hadrons is [90]

$$\mathrm{d}P \propto \prod_{\alpha} \frac{1}{n_\alpha!} \prod_{i=1}^{n} \frac{\mathrm{d}^3 p_i g_i s_i}{(2\pi\hbar)^3 \cdot 2E_i} \frac{M}{\varepsilon_{\mathrm{FO}}} \delta\left(M - \sum_j E_j\right) \delta\left(\sum_j \mathbf{p}_j\right) \delta_{f,\sum_j f_j}, \tag{3.19}$$

where n_α is the number of hadrons of species α, $p_i = (E_i, \mathbf{p}_i)$ is the four-momentum of particle i, g_i its degeneracy, f_i its quark flavour content, s_i is a parameter describing strangeness production, $\varepsilon_{\mathrm{FO}}$ is the energy density at freeze-out (set to $0.22\,\mathrm{GeV/fm^3}$), and M is the cluster mass. Due to radial flow, in the core part M is scaled by a factor of

$$k = \frac{0.5 y_{\mathrm{rad}}^2}{y_{\mathrm{rad}} \sinh y_{\mathrm{rad}} - \cosh y_{\mathrm{rad}} + 1}$$

in order to conserve energy, where y_{rad} is the maximum radial rapidity, which is different for pp than for heavy-ion collisions. The δ terms (Dirac δ for continuous quantities and Kronecker δ for discrete ones) ensure that energy, momentum, charge, and flavour are conserved. In EPOS3, the micro-canonical picture is being replaced by a grand-canonical ensemble, which gives a full thermal picture for the hadronisation in the core part.

3.8.5 Successful Predictions by PYTHIA, Angantyr, and EPOS

In this section, I will present a few results implicating the partial success of either of PYTHIA (or in the case of rope hadronisation, also one of its predecessors), Angantyr, and EPOS, which are related to heavy-quarkonia, strangeness enhancement, and flow. To begin with, the relative yield of J/ψ increases with multiplicity in pp collisions,

as is shown in Fig. 3.30 [76]. While not necessarily an impication of a QGP in small systems, this is expected from several CGC models, i.e. the initial state preceding QGP formation. The reason is that the relative number of charged hadrons produced in these models decreases at high multiplicity, which is not the case for quarkonia. However, a similar effect is observed also in PYTHIA when MPIs are enabled, as well as in EPOS3, where the prediction becomes close to what is observed in data if enabling hydrodynamical evolution [77] (not shown in the figure).

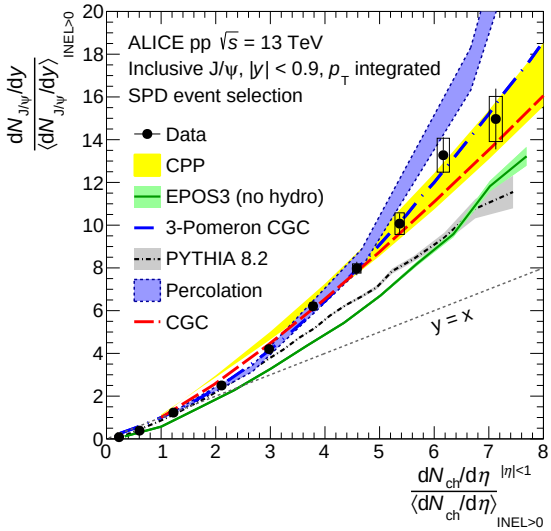

Figure 3.30 – Yield of J/ψ normalised to the event average, as a function of charged-particle multiplicity likewise normalised to its average, measured by ALICE in pp collisions at $\sqrt{s} = 13\,\text{TeV}$, and compared with predictions from several models, including two CGC models, PYTHIA8.2, and EPOS3. Figure taken from Ref. [76].

Another result – which is of more interest for this thesis – is that the strangeness enhancement results presented in Fig. 3.13 are reproduced qualitatively by both the DIPSY model and EPOS LHC, as is shown in Fig. 3.31, although not by PYTHIA8 [51]. Here, DIPSY was a model based on dipole cascades [79], but was abandoned since it did not reproduce many observables as well as PYTHIA, which is being developed by the same group. This was the first model where rope hadronisation was enabled [78]. The main idea is that the denser environment found in more central collisions increases the number of interactions between strange quarks, increasing the yields of multistrange hadrons. In EPOS on the other hand, the strangeness enhancement is due to an increase of the core fraction in the system, resulting in a larger size of the thermalised medium. In the main analysis of this thesis, described in Chapter 8, angular correlations are used to try to understand which mechanism is responsible for strangeness production in pp collisions, and hence the origin of strangeness enhancement.

The narrowing of the balance function in pp collisions with increasing multiplicity shown in Fig. 3.14 is quite well described by the CR mechanism [54], as is shown in Fig. 3.32a. A more recent result is however clashing with the PYTHIA prediction. If

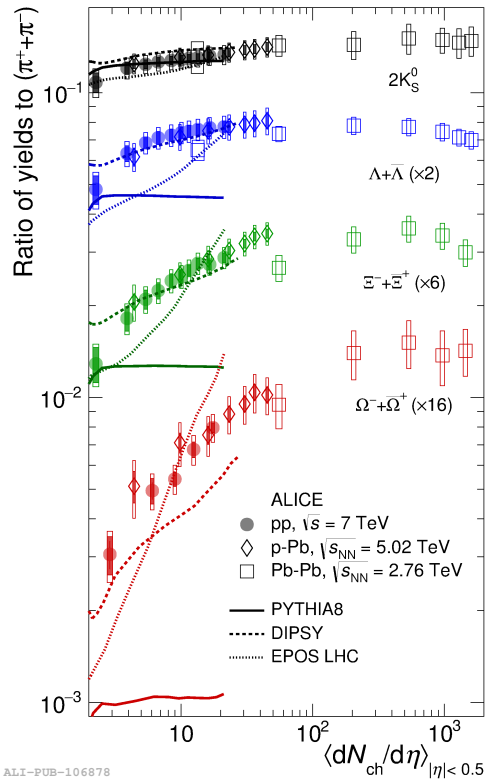

Figure 3.31 – An older version of Fig. 3.13. Yield ratios between various hadron species and pions as function of multiplicity, for pp, p–Pb, and Pb–Pb collisions, measured by the ALICE detector. This version includes predictions from phenomenological models. Figure taken from Ref. [51].

the narrowing is due to radial flow, the balance function should be narrower for heavier particles such as protons. This is indeed the case both in data and in PYTHIA, but there is almost no change with multiplicity in Δy, which is quite far off from the PYTHIA prediction [80], as is shown in Fig. 3.32b, suggesting that some other mechanism might be involved.

Finally, with the string shoving mechanism enabled in PYTHIA, a significant v_2 is generated by the model in pp collisions [81]. While a non-zero v_2 is obtained also from CR, this exhibits a very different multiplicity dependence than what is observed in data. Comparisons with data for both models are shown in Fig. 3.33a. Another promising feature of string shoving is that if the string density is allowed to vary with multiplicity, which is not unreasonable given the different densities of these environments, one can get quite close to the observed v_2 in Pb–Pb collisions (using the Angantyr package), as is shown in Fig. 3.33b. Moreover, while the flow in pp collisions from the CR mechanism always produces a positive $c_2\{4\}$, this problem disappears at high multiplicity with string shoving enabled, as is shown in Fig. 3.33c.

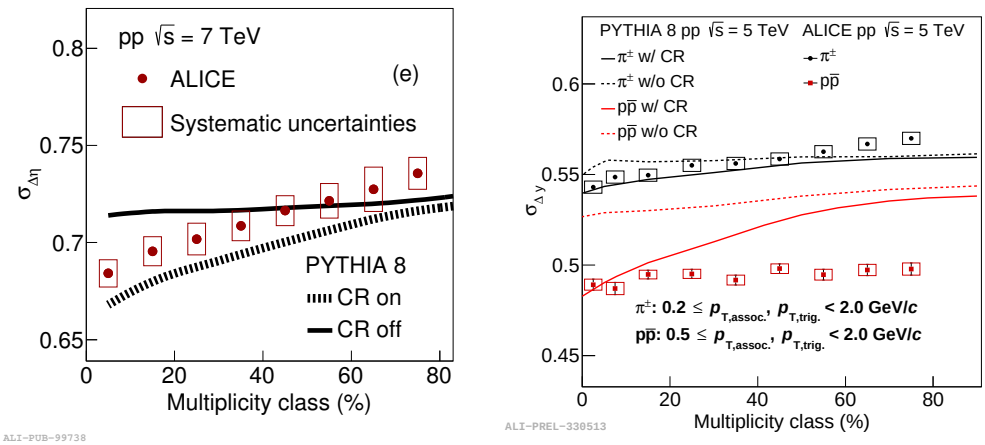

(a) **(b)**

Figure 3.32 – Balance function width in $\Delta\eta$ or Δy as a function of collision centrality, measured in pp collisions by ALICE and compared with predictions by PYTHIA8, both with CR turned on and off, for **(a)** inclusive charged particles at $\sqrt{s} = 7$ TeV, and **(b)** identified pions and protons at $\sqrt{s} = 5.02$ TeV. Figures taken from Refs. [54] and [80].

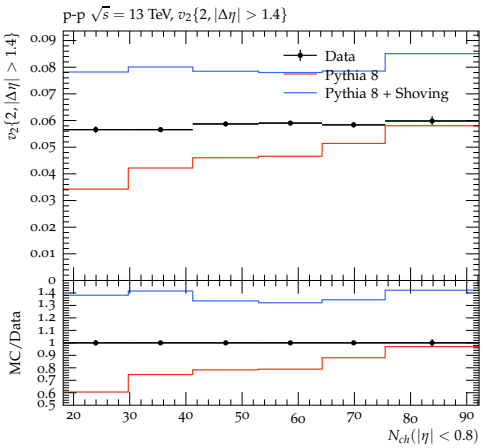

(a)

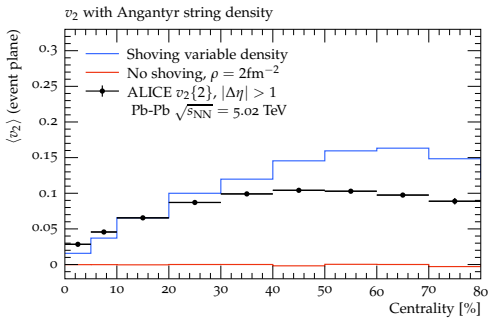

(b)

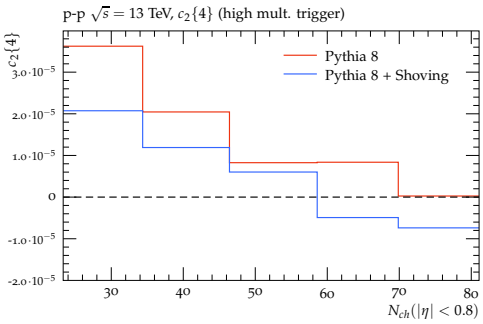

(c)

Figure 3.33 – **(a)** $v_2\{2\}$ measured with the scalar-product method as a function of multiplicity in pp collisions at $\sqrt{s} = 13$ TeV, for ALICE data and PYTHIA8 predictions both with and without string shoving enabled. **(b)** The same measurement in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV, compared with Angantyr predictions of v_2 extracted from the event plane with non-flow subtracted. Note that higher centrality = lower multiplicity. **(c)** PYTHIA predictions of $c_2\{4\}$ in pp collisions at $\sqrt{s} = 13$ TeV. Figures taken from Ref. [81].

Part II

Experimental Setup and Development

Chapter 4

The ALICE Detector

4.1 Detector Overview

ALICE – *A Large Ion Collider Experiment* – is the only LHC experiment which is dedicated to study heavy-ion physics. As such, the aim is not to find rare particles as is the primary focus of the other three major LHC experiments, but instead to collect as much data as possible from the full heavy-ion event. This puts strong demands on the detector. A Pb–Pb collision typically produces $\sim 500 - 1000$ charged particles per pseudorapidity unit ($\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}\eta$), and it can reach as high as $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}\eta \sim 2000$ in the most central collisions [91]. Ideally, one wants to individually identify every single particle in the detector acceptance and measure its position and momentum, and in this way characterise the QGP. This is not possible yet, but nevertheless ALICE has the best particle identification (PID) capabilities for charged hadrons of the LHC experiments, and particularly at low $p_{\mathrm{T}}$where the bulk of the particles are produced, which is particularly relevant when studying the QGP. Current limitations in triggering rate, p_{T} range, and resolution will be considerably improved during the ongoing upgrade (scheduled to be finished by the start of Run 3), which will be briefly described later in this chapter and in Chapter 5.

The heavy-ion data is collected during a few weeks at the end of each operational year. During the rest of the year, data from pp collisions is collected. While – in the scope of the heavy-ion programme – originally thought as a non-QGP reference to the heavy-ion data, proton collisions are now also an important part of the ALICE physics programme. Here, the collision multiplicities are much lower than for a Pb–Pb collision, so despite a higher interaction rate, this puts a less restrictive demand on the detector[1]. During periods of maintenance, some subsystems also collect data from cosmic muons, which is mainly done to monitor the detector conditions and performance.

An overview of the ALICE detector is given in Fig. 4.1. Its total dimensions are $16 \times 16 \times 26\,\mathrm{m}^3$ [94]. The detector is divided into the central barrel and the muon

[1]While the pp luminosity at ALICE is more than 1000 times greater than the Pb–Pb luminosity (derived from Ref. [92]), the large busy times of some detector subsystems cause the recorded event rates to be similar for the two run types. On the other hand, the radiation is greater during the pp data taking due to the higher luminosity, which can have a negative impact on the electronics.

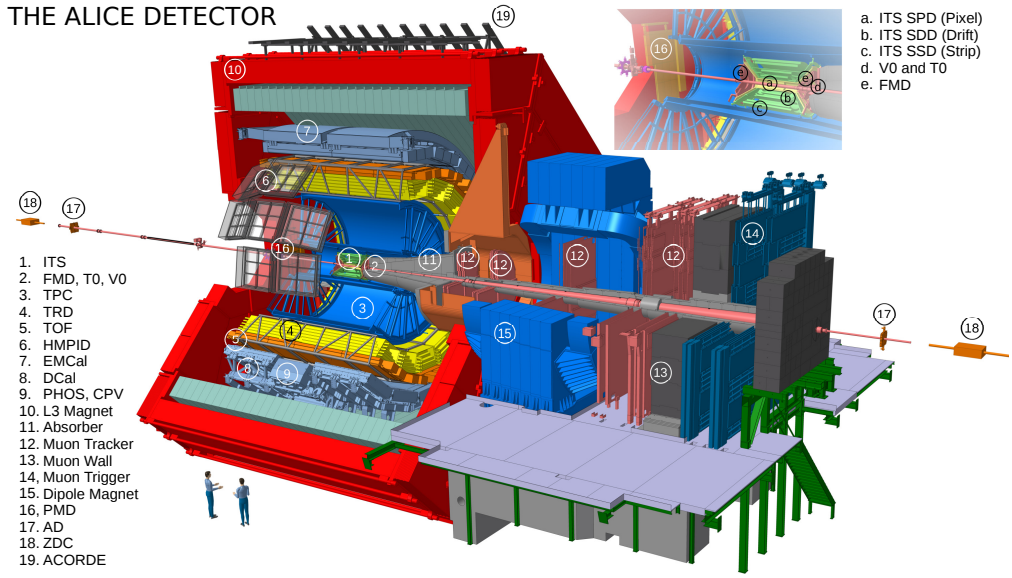

Figure 4.1 – Schematics of the ALICE detector. The subsystems relevant for this thesis are described in the text. Figure taken from Ref. [93].

spectrometer. The central barrel is contained within the L3 magnet, which is a warm solenoid magnet with a magnetic field of $B = 0.5\,\mathrm{T}$. The muon spectrometer extends only in the forward region (C side) and contains a dipole magnet ($B = 0.67\,\mathrm{T}$) which bends the muons away from the interaction vertex in the horizontal plane.

Closest to the interaction point, in the middle of the central barrel, is the Inner Tracking System (ITS), which is used for triggering and high-resolution tracking. This is further described in Section 4.2. Next to the ITS, in the forward and backward regions, are the T0 and VZERO (usually abbreviated V0) detectors. The T0 detector is used for measuring the time of the event, whereas the V0 detector is used for measuring the event multiplicity and additional triggering. The latter is described in Section 4.3. Surrounding the ITS, is the Time Projection Chamber (TPC), which is the main tracking device and is further described in Section 4.4. The next two layers of the central barrel are the Transition Radiation Detector (TRD), which is used for electron identification above $1\,\mathrm{GeV}/c$, and the Time-Of-Flight (TOF) detector, which measures the velocity of the particles and is used to complement the PID information, in particular at intermediate p_T. The TOF detector is described in Section 4.5. Outside the TOF detector, there is no detector which covers the full azimuthal angle. The most important detectors here are the electromagnetic calorimeters, EMCal and DCal, which are placed on the top and bottom of the central barrel, respectively. These measure the energy of photons and electrons.

The muon spectrometer has a conical shape which extends in the forward direction from the interaction point. First is an absorber which is used to stop less penetrable particles, so that mostly muons remain. Further away from the collision vertex are five double-planed Muon Tracking Chambers (MCH) perpendicular to the cone, which are

used to determine the momentum of the muons. These are placed both before and after the magnet. Finally, a second absorber is used to block all remaining particles which are not muons. Behind this wall are the two Muon Trigger (MTR) stations, using a similar design. The tracking of a muon is only initialised if it hits both these detectors. The technology used for the MCH is briefly described in Section 5.1.

Finally, there are a few detector components located outside the detector cavern. These include the Zero-Degree Calorimeter, ZDC, which is a set of two calorimeters located on either side of the interaction point, 116 m away along the beam pipe, and are collecting the spectator remnants of the collision. During heavy-ion collisions, these are used for triggering.

4.2 Inner Tracking System

The ITS consists of six concentric layers of silicon detectors surrounding the beam pipe – two layers each of Silicon Pixel Detectors (SPD), Silicon Drift Detectors (SDD), and Silicon Strip Detectors (SSD) – and covers the pseudorapidity window $|\eta| < 0.9$ [94]. A schematic overview is given in Fig. 4.2. The main purposes of this detector are to accurately find the position of the collision vertex and secondary vertices from weakly decaying particles (in particular strange hadrons and B and D mesons, which have typical decay lengths of a couple of cm), to complement the TPC tracking information, and in the case of the SPD to provide triggering input for the entire detector. This detector system has a very fine granularity, providing a high tracking and vertexing resolution also at extremely high multiplicity. The resolution for each subsystem is summarised in Table 4.1. Moreover, due to its proximity to the beam pipe, a high momentum resolution and large specific energy loss range, particle identification is

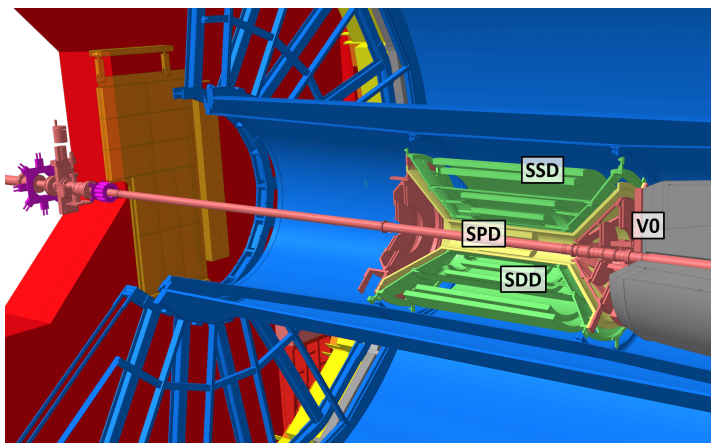

Figure 4.2 – Schematics of the Inner Tracking System and nearby detector components. The V0 detector is only shown on the C side (a similar detector disk is also mounted on the A side, but much further from the interaction point). Modified version of figure taken from Ref. [93].

Table 4.1 – Coordinate resolution in azimuthal ($r\varphi$) and longtudinal (z) directions for each subsystem in the ITS [94].

Detector	$r\varphi$ precision (µm)	z precision (µm)
SPD	12	100
SDD	35	25
SSD	20	830

possible below the lower momentum limit of the TPC (which is at $\sim 150\,\mathrm{MeV}/c$)[2].

Throughout the ITS, silicon based detector modules are used. Silicon is a semiconductor, and as such it has a valence band and a conduction band, which are separated by a band gap. When used in a detector, a pn junction is used where a reverse bias is applied, which depletes the active area of charge carriers [95, pp. 226-232]. When a charged particle (or photon) hits the detector, this will create electron-hole pairs along the trajectory[3], as shown in Fig. 4.3, which generates a current pulse. This current will traverse the electric field and eventually be collected at the electrodes, where a current amplifier is connected. Since the ITS is only used for tracking, the relevant information is whether a particle has hit the detector. This is ensured by triggering if the pulse is above a certain threshold. The advantages of using a semiconductor are that this creates a fast signal, which makes it possible to determine the interaction time at high

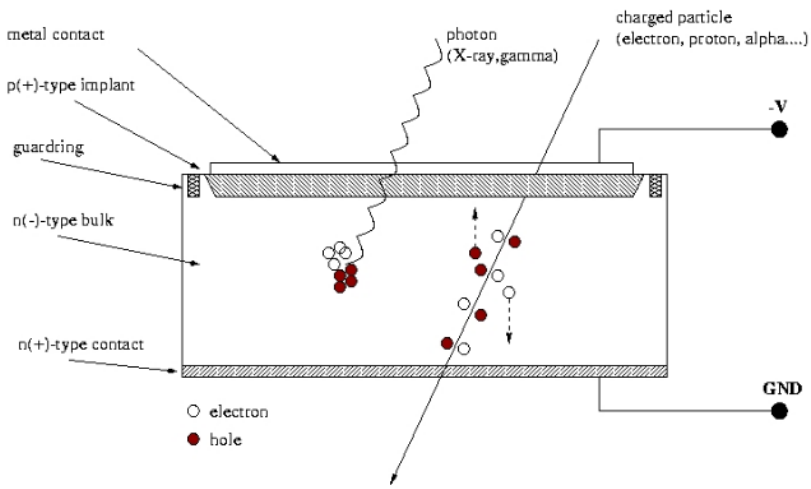

Figure 4.3 – Basic principle of a semiconductor detector. When a charged particle traverses the detector, it excites electrons from the valence band to the conduction band, which creates electron-hole pairs. These will drift in the applied electric field to the metal contacts, which generates a current pulse that can be detected. Figure taken from Ref. [96].

[2]If the momentum is too low, a charged particle will be trapped in the magnetic field and escape in the longitudinal direction, but may still be detected in the ITS. For heavier particles – in particular protons – there is also a large energy loss in the ITS, further reducing the efficiency at low momentum.

[3]For photons this will require an intermediate step of either Compton scattering or pair production.

accuracy, and that is makes it possible to achieve a very high tracking resolution, since the electron diffusion length is only a few µm.

Each subdetector of the ITS uses a different technology. The SPD is built of pixels of silicon diodes, each measuring 50×425 µm^2. These are read out by 1200 chips per layer, each covering 8192 cells, totalling $9.8 \cdot 10^6$ cells per layer. The fine granularity makes it possible to detect the position of a charged track traversing the detector to a very high precision. The layers are located 3.9 and 7.6 cm from the beam axis, respectively, and extend 282 mm in the longitudinal direction.

The SDDs are built of 260 drift cells distributed over the two layers, where the drift time of charge carriers is measured in order to precisely determine where the track has interacted with the cell. To each cell, 256 collection anodes are mounted along the z axis with a spacing of 294 µm. The drift regions are 35 mm long and extend in the φ direction. To determine the $r\varphi$ coordinate, the drift velocity is monitored by MOS injectors in the substrate (these are triggered during gaps in the LHC bunch crossing schedule as to not interfere with collisions). The position is determined by integrating the velocity over the time measured at the anode. The effective cell length is ≤ 202 µm, depending on the temperature. Moreover, these detectors are able to measure the number of collected charge carriers, which is proportional to the energy deposit, or more generally $\mathrm{d}E/\mathrm{d}x$, which is used for PID (cf. Section 8.5). The SDD layers are located 15.0 and 23.9 cm from the beam axis, respectively, and the lengths of the layers are 443 and 593 mm in the z direction.

Finally, the SSDs are built of double layers of silicon strips, which are put at an angle of 35 mrad ($\approx 2°$) relative to each other. Therefore, a particle crossing one of the SSDs will give rise to a signal in both strip layers, which gives a detection at the crossing point. Moreover, the number of collected charge carriers is measured, providing $\mathrm{d}E/\mathrm{d}x$ information in the SSD. The strips have a width of 95 µm, but due to the strip arrangement, the effective resolution is better than this in $r\varphi$ but worse in z. In total, $1.15 \cdot 10^6$ strips are used in the inner layer and $1.46 \cdot 10^6$ in the outer one, distributed on chips with 768×2 strips each. The inner and outer SSDs are located at 38 and 43 cm from the beam axis, respectively, and extend 86 and 98 cm in the z direction.

ITS Upgrade

The ITS is currently being upgraded, with the installation taking place during the LHC Long Shutdown 2, which started in December 2018 and is scheduled to end during the beginning of 2022. The main purpose of the upgrade is to be able to handle the increased collision rate during the heavy-ion data taking, or more specifically to match the readout rate enabled by the TPC upgrade. After the upgrade, the ITS readout rate will be increased from 1 kHz to 100 kHz during Pb–Pb data taking and 1 MHz during pp data taking [97, 98]. Moreover, the new ITS will have a much higher tracking and $\mathrm{d}E/\mathrm{d}x$ resolution than the old one, and since the innermost layer is closer to the beam pipe (only 22 mm from the centre), the primary vertex, as well as secondary vertices, can be determined to a higher precision.

Opposed to the current ITS, the upgraded detector will use a single technology throughout, namely light-weight high-resolution pixel detectors. These are distributed over seven layers, and a total number of $\sim 12.5 \cdot 10^9$ pixels are being used.

4.3 V0 Detector

The V0 detector consists of two circular arrays of scintillator counters, one in each of the forward regions [94]. The array on the A side (V0A; opposite to the muon arm) is placed 300 cm from the collision vertex, covering the pseudorapidity region $2.8 < \eta < 5.1$, but in order to accommodate the muon absorber, the V0C array is placed only 90 cm from the vertex, covering $-3.7 < \eta < -1.7$ (cf. Fig. 4.2). Each array consists of 4 layers with 8 scintillators each, each covering a circular segment spanning 45° in azimuthal direction. This results in quite poor granularity, but is useful for multiplicity measurements (and for some observables, the resolution in azimuthal angle is good enough that also this information can be used in measurements, as will be described in Chapter 7).

A scintillator detector makes use of a scintillating material, meaning that it will get excited by charged particles traversing it, which generates a light pulse [95, pp. 157-159, 177-180]. In the V0 detector, a plastic scintillator is used. The generated light is guided to photomultiplier tubes at the detector edges, where the photons are converted to electrons and multiplied in a strong electric field. The resulting voltage pulse can now be detected by the detector electronics. Plastic scintillator detectors are fast detectors with a good timing resolution (in the V0 better than 1 ns). When used in counting mode, such as in the V0, the number of pulses surpassing a certain threshold is measured, which here gives a measure of the event multiplicity in the forward regions. Since this is strongly correlated with the total event multiplicity[4] and hence the collision centrality, the V0 multiplicity is used for dividing the events into centrality classes, cf. Section 3.2. Moreover, the V0 detector is used for triggering, as decribed in Section 4.7.

4.4 Time Projection Chamber

The TPC is the main tracking device of the ALICE detector. This is a gaseous ionisation detector, or gas detector for brevity, meaning that it uses ionisation of a noble gas for detecting charged particles.

When charged particles traverse a gas, they will cause ionisations within it, creating electron-ion pairs. In a gas detector, an electric field is applied, which causes the electrons to drift towards the anode and ions towards the cathode [95, pp. 127-143]. In most detector setups, including a TPC, thin wires are used as anodes. Therefore, in the vicinity of the anode, the electric field is inversely proportional to the distance to the anode, resulting in a very strong electric field near the wire. Consequently, electrons accelerated in the field will gain enough energy to cause secondary ionisations, giving rise to a chain reaction. This is known as avalanche multiplication, cf. Fig. 4.4. Therefore, a few initial electrons can give rise to a substantial voltage pulse induced at the cathode readout plane, which is what is detected by the system. The voltage pulse will however

[4]Due to event anisotropies, in particular resulting from jets, there is not a complete correlation between the multiplicity in the forward region and the total multiplicity, but in heavy-ion collisions the correlation is very strong. In pp collisions on the other hand, with lower multiplicities and hence stronger fluctuations within the event, this correlation is not as pronounced. Nevertheless, it is a good idea to measure the event multiplicity in a different region than the one used for the analysis (in this thesis mid-rapidity), since that may bias the results.

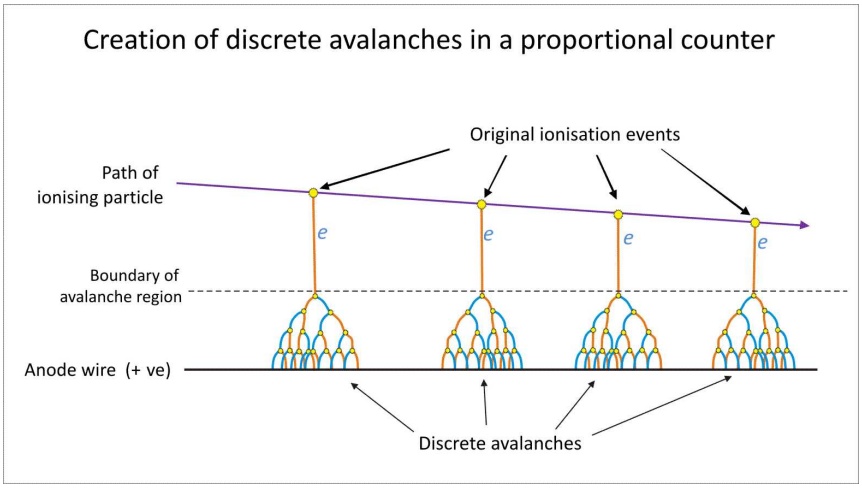

Figure 4.4 – Principle of avalanche multiplication in a gas detector. Electrons released during ionisation by an energetic particle will start a chain reaction of ionisations – avalanches – near the anode wire. In reality, avalanches may overlap however, which is not indicated by the figure. Figure taken from Ref. [99].

not stop until the ions are collected at the cathode. Since these have a much lower drift velocity than the electrons, gas detectors are quite slow.

If the voltage is not too high, the deposited energy at the electrodes is proportional to the number of initial ionisations, which in turn is proportional to the energy deposit in the gas, giving $\mathrm{d}E/\mathrm{d}x$ information of the ionising particle. This is the operational mode of the TPC, and is known as proportional mode. The voltage should not be too high, though, since then the gas atoms may get excited during avalanches. The photons released in the deexcitation process may then trigger additional avalanches. Another issue is that during recombination at the electrodes, the gas atoms may again enter an excited state, which in a similar process may cause secondary ionisations or at least prolong the avalanche. To prevent this, a molecular component – usually CO_2 or an organic molecule – is added as a quencher to the gas. Due to vibration and rotation modes, this component has more excitation levels than the noble gas, and thus these will dissipate absorbed energy without liberating further electron-ion pairs.

The ALICE TPC has a cylindrical geometry surrounding the ITS, as shown in Figs. 4.1, 4.5, and 4.6 [100]. Full tracking information is possible in the window $|\eta| < 0.9$. A longitudinal electric field is applied over the cylinder, and it is divided in the centre by a cathode plate, where the ions released during the initial ionisation are collected. At the endcaps are Multi-Wire Proportional Chambers (MWPCs), which are arrays of anode wires operating in proportional mode. When a charged particle from the collision traverses the gas, electron-ion pairs will be created along its trajectory. Since the field is longitudinal with minimal space charge distortions, the electrons will be projected onto the endcaps, giving precise information about the (r, φ) coordinates of the trajectory. Since the drift velocity is being monitored (see below), the z coordinate is obtained by measuring the arrival time of the electrons.

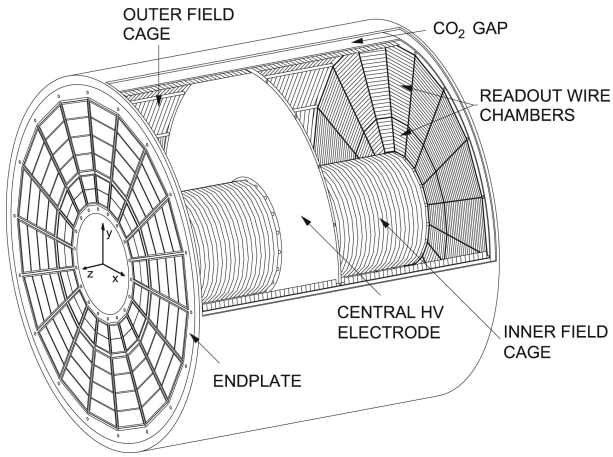

Figure 4.5 – Sketch of the TPC field cage. The electrons are collected at the readout chambers at the endcaps, whereas the ions released in the initial ionisation events are collected at the central high-voltage electrode. Figure taken from Ref. [100].

The ALICE TPC has an inner radius of 85 cm, an outer radius of 250 cm, and an overall length of 500 cm. With a voltage of 100 kV between the centre and endcaps, the resulting field strength within the detector volume is 400 V/cm. For the most frequently used TPC gas mixture – 85.7% Ne, 9.5% CO_2, and 4.8% N_2[5] – this results in an average drift velocity of 2.65 cm/µs for the electrons, or an overall drift time of 94 µs for the longest drift distances.

The endcaps are divided into 18 sectors, with four readout chambers each – two inner readout chambers (IROCs) and two outer readout chambers (OROCs). Being closer to the beam pipe, the wire density is higher for the IROCs, and the voltage slightly lower. There are no anode wires between the sectors, resulting in gaps in the detector acceptance. The readout chambers have three layers of wires. The innermost wires are thin anode wires spaced 2.5 mm from each other. The middle layer consists of thicker cathode wires, where the ions released during avalanches are collected. The signal is read out from cathode pad planes, with a size in $(r, r\varphi)$ space ranging from $4 \times 7.5\,\mathrm{mm}^2$ in the IROCs to $6 \times 15\,\mathrm{mm}^2$ in the outermost OROC sector. This results in a total number of 159 radial clusters.

The outermost layer is a gating grid consisting of thicker wires spaced by 1.25 mm, which can either be in "open" or "closed" mode. This is required since the field gets distorted by the avalanche ions, preventing further avalanches. Consequently, these ions must not enter the detector volume, which is ensured by blocking the MWPCs until the ions are collected, adding dead time to the system. This is done by applying an alternating voltage between the wires, which deflects the ions from the positively

[5]In recent years, argon has sometimes been used as the primary component, but neon is preferred since its higher drift velocity reduces space charge distortions. CO_2 was chosen as a quenching gas, since organic quenchers may dissolve during detector use, producing by-products which may impair the detector performance. They are also flammable, which would have been hazardous in such a large detector volume. The reason for including nitrogen is that this results in more stable conditions, since some nitrogen from the surrounding air will inevitably diffuse into the detector volume.

Figure 4.6 – Photograph of the TPC, when it had been taken out from the ALICE cavern for the upgrade. At this stage, the old readout chambers had been removed and the TPC was about to be moved to a cleanroom for installation of the new ones, so only the field cage (and the wire chambers, which were not yet removed) remained.

charged wires and collects them at the negatively charged ones (and vice versa for electrons originating from subsequent collisions). As a consequence, the maximum detection rate is limited to 3.3 kHz, although the data acquisition (DAQ) system further limits the recording rate to ~ 1 kHz for Pb–Pb collisions. Therefore, the gating grid is synchronised with the trigger, which is described in Section 4.7.

In order to determine the z coordinate accurately, the drift velocity is carefully monitored using a laser system, which also detects inhomogeneities in the electric field. Variations can be due to e.g. temperature fluctuations and space-charge effects. The laser system, which is run every half an hour between bunches in the LHC schedule, produces planes of tracks in the detector perpendicular to the electric field, which makes it easy to measure drift time as well as deviations caused by distortions. The lasers produce UV photons with an energy slightly lower than the ionisation potentials of the gas components. Instead, organic impurities in the gas are ionised in the process.

TPC Upgrade

The TPC is currently being upgraded, with the installation taking place during Long Shutdown 2, just as for the ITS upgrade. The main purpose of the upgrade is to avoid the greatest bottleneck in the detection system, namely the gating grid. Therefore, the MWPCs are being replaced by Gas Electron Multipliers (GEMs), which allow for continuous operation [101]. Including a major upgrade in the DAQ and readout system, this makes it possible to record tracks piled up from several events simultaneously, significantly increasing the detection rate. During Run 3, the design detection frequency is 50 kHz for Pb–Pb and 500 kHz for pp collisions, with the latter possibly being able to increase to 1 MHz[6].

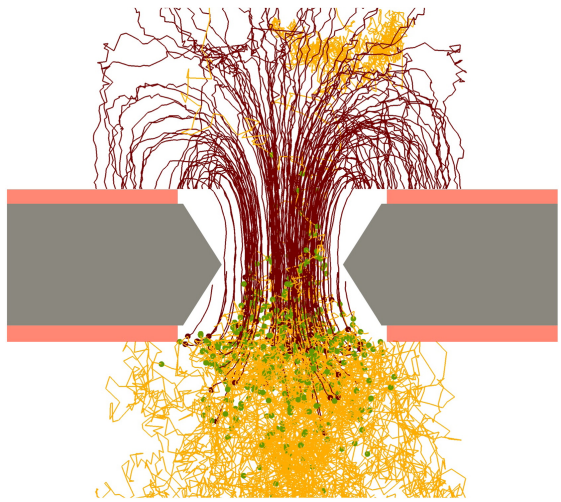

Figure 4.7 – Simulation of an avalanche in a hole in a GEM foil, originating from two electrons, projected onto the cross section plane. The green dots mark places where ionisation takes place. The dark lines mark the drift of ions and the fair lines the drift of electrons. Figure taken from Ref. [102].

A GEM consists of a thin (~ 50 µm) foil made by an insulator material, in ALICE polyimide, coated by even thinner conducting layers (copper) at the top and bottom. This is perforated by small holes, which replace the wires in the MWPCs. The layer facing the detector volume has a lower potential than the surrounding volume, whereas the other layer has a higher potential. With a potential difference of $200 - 400$ V, the electric field strength within the hole is large enough for avalanche multiplication. Following this, the released electrons will then drift towards the outer surface, whereas the ions will be collected at the inner surface, see Fig. 4.7. One important aspect which is ensured by this design is that less than 1% of the ions flow back into the detector, which is a rough limit for tolerable field distortions. In order to stabilise the operations, e.g. by reducing the risk of discharges, and to increase the gain, the GEM foils are

[6]For pp collisions, the limitation will rather be the primary vertex resolution from the ITS than the number of tracks in the TPC. If the pileup gets too large, it will not be possible to accurately assign all tracks to the right event.

stacked into four layers, with varying spacing between the holes in the foil, in order to ensure that they are not aligned.

4.5 Time-Of-Flight Detector

The working principle of the TOF detector is to measure the time-of-flight from the interaction point of a particle[7], and thus determine its velocity. Combined with the momentum information, one can extract the mass of the particle, cf. Section 8.5. The arrival time measurement is achieved by an array of Multi-gap Resistive-Plate Chambers (MRPCs), which are thin gaseous detector cells over which a high and uniform voltage is applied [94, and refs. therein]. These are divided into two half cells, each divided into five smaller modules blocked by resistive glass plates, cf. Fig. 4.8. The field is strong enough for avalanche multiplication, thus a particle traversing the cell gives rise to a detectable signal. The glass plates effectively block the avalanches, so this setup reduces the time jitter, which scales with the propagation distance. The signals from each gap sum up to the total signal, so using multiple gaps increases the signal strength. The achieved time resolution is about 40 ps.

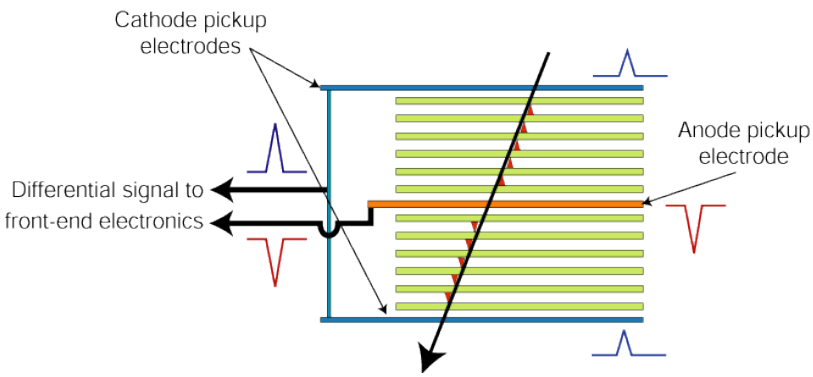

Figure 4.8 – Working principle of a TOF detector cell. A charged particle entering the cell will give rise to avalanches in the gas, which are blocked by glass plates mounted inside the cell. The resulting electrons will be picked up at the anode, which divides the cell into two half cells (in practice, there are two anode plates separated by an insulating layer). Figure taken from Ref. [103].

The TOF detector covers the full azimuthal angle and the pseudorapidity region $|\eta| < 0.9$. The inner radius is 370 cm and the active length is 741 cm. A total of 157 000 cells are used, which are arranged into strips of dimensions $122 \times 13\,\mathrm{cm}^2$, put in 18×5 modules (in (φ, z) space). The azimuthal sector boundaries are aligned with the dead areas of the TPC. The strips overlap each other to achieve full coverage within each module. Moreover, in order to minimise the path length of particles coming from the collision, the modules are tilted so that each strip is facing the interaction point perpendicularly.

[7] The collision time is measured by the T0 detector.

4.6 Tracking in the Central Barrel

To reconstruct the particle tracks from the hits in the detector, a tracking algorithm is used. A pedagogic overview is shown in Fig. 4.9. The technique used is a Kalman filter algorithm, which uses a parametrisation of the track and successively optimises the trajectories through addition of more and more space points [94]. The algorithm starts out from a seed consisting of a few clusters in the TPC. From these, an initial guess of the track is made based on the constraint that the track originates from the collision vertex, which in turn is extrapolated from the tracks in the SPD. The process is later repeated by removing this constraint, which would imply that the track originates from a secondary vertex following a decay. The tracking information is then improved by propagating the track inwards and adding space points within 4σ (standard deviations of the trajectory) from the best guess, taking into account effects from scattering and energy loss. The parameters as well as the covariance matrix of the parametrisation are updated in each iteration, improving the accuracy. If more than one detector hit fulfils the criterion, several different propagations are tested. Not until the end, the tracks with the lowest χ^2 (highest significance) are selected.

When reaching the edge of the TPC, the tracking is improved by adding space points from the ITS. For tracks where the constraint of a primary vertex is not lifted, this is first used also as a constraint for the ITS tracking. Due to the higher precision of the ITS, making it possible to find secondary vertices closer to the primary vertex, the procedure is repeated without this constraint. For tracks where the constraint has been lifted in the TPC, the ITS procedure is only done without imposing it.

When the track propagation is done in both the TPC and the ITS, the Kalman filtering is reversed by using the track parameters found in the first iteration as an initial guess, and repeating the procedure starting from the innermost layer of the ITS and propagating outwards. This is done to remove outliers from the track propagation. When reaching the edge of the TPC, there is an option to include space points from the TRD as well (not done in the analyses included in this thesis due to the low efficiency of this detector) and eventually the TOF (and in some analyses HMPID and/or CPV which however do not cover the full azimuthal angle). In the main analysis of this thesis both tracks with and without an associated TOF hit are used. Finally, a final refit is being done starting from the outer detectors and propagating inwards in order to further optimise the track parameters. While a successful TPC refit is usually required, tracks without a successful ITS refit can be used if the track has a TOF hit, or precise timing information is not required (more about this in Chapter 8).

4.7 Triggering and Data Aquisition

As already has been mentioned, the ALICE subsystems can handle an event rate of up to a few kHz. However, the LHC bunch crossing rate is 40 MHz, so even though not all bunch crossings result in a collision[8], extensive triggering is required to reduce the

[8]During pp collisions, the accelerator is tuned such that the luminosity at ALICE is much lower than at ATLAS or CMS ($L \simeq 10^{30} - 10^{31}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ at ALICE compared to $L \simeq 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ at ATLAS and CMS [92]). This reduces pile-up, but also the effective collision rate.

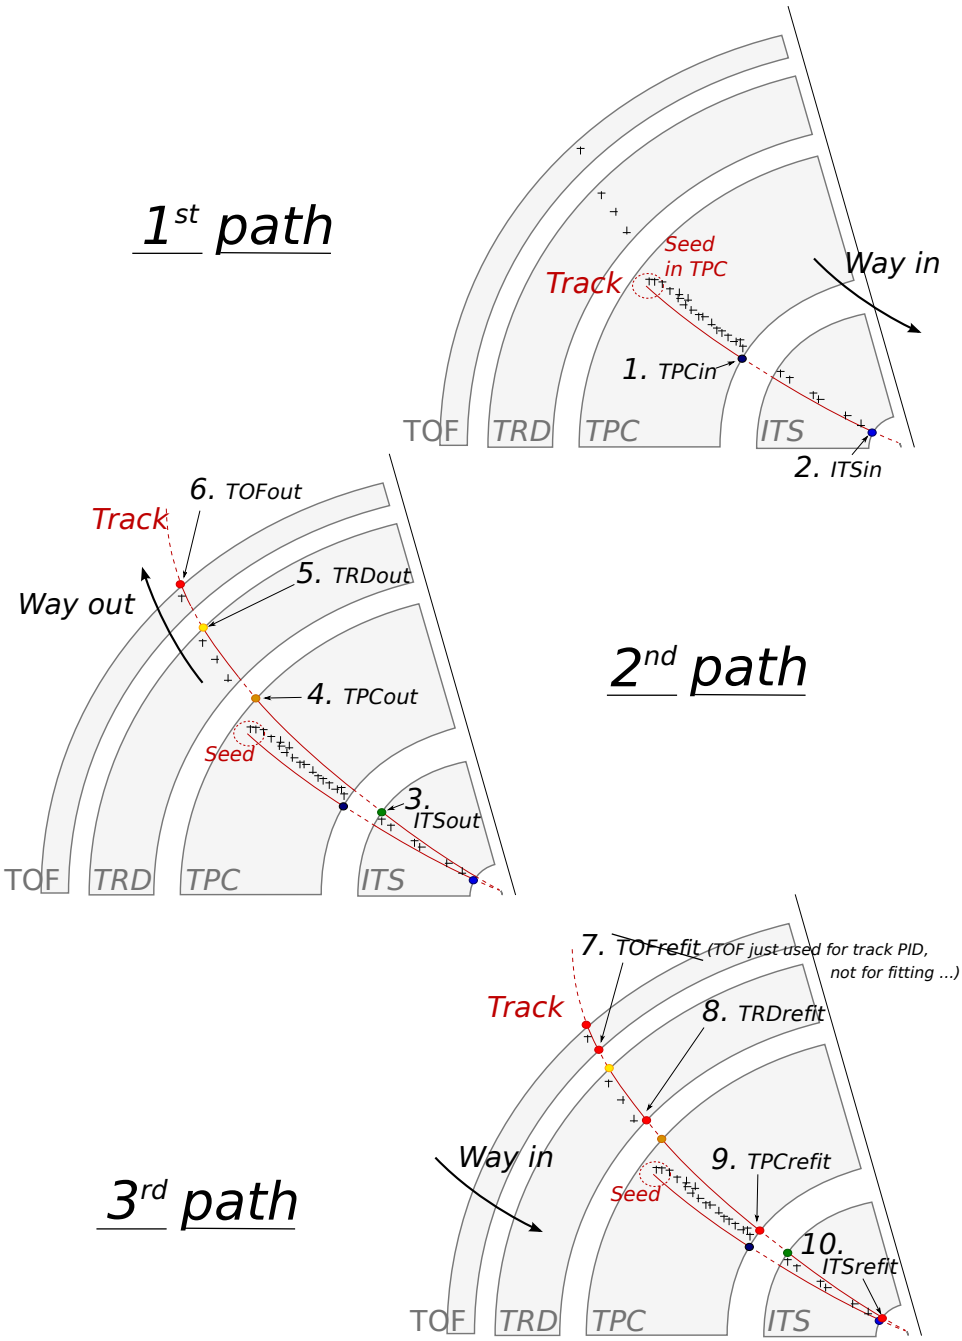

Figure 4.9 – Overview of the track reconstruction process in ALICE. See the text for details. Figure taken from Ref. [104].

event rate and ensure that all particles detected during a single event indeed are from the same collision.

The most important detectors for triggering are the V0 and SPD detectors, although the ZDC is also used in Pb–Pb collisions and several others are used for specific signals such as dimuons, photons, or jets [94]. The trigger input is handled by the Central Trigger Processor (CTP), which sends a command to all active detectors upon a successful trigger. There are three levels of triggers, L0, L1, and L2. Only events accepted by the highest level, L2, are read out. The L0 trigger is synchronised with the LHC bunch crossing clock at ALICE and it only includes the fastest inputs. All input at this level are processed within 1.2 µs of the collision. The remaining fast inputs are handled by the L1 trigger, which takes 6.5 µs to complete. The L2 trigger handles all slower inputs, and must be completed before the TPC is read out, i.e. within 88 µs from the collision.

The detector signals used by the trigger depend on the trigger configuration. Two important configurations are minimum bias (MB), which basically accepts events following the full multiplicity distribution, and high multiplicity (HM), which is only active in pp collisions and only triggers on the $\sim 0.1\%$ of the events with highest multiplicity. The analyses described in this thesis all use the MB trigger, since part of their purpose is to study how certain observables depend on multiplicity. In order to maximise the number of rare particles, the HM trigger would have been a much better choice however, since this records a much larger number of tracks (in particular in pp collisions). The MB trigger requires a hit in both V0 detectors (V0-AND), which is sufficient for the low multiplicities in pp and p-Pb collisions. Additional constraints are later applied on the V0 timing information and the correlation between number of tracklets and clusters in the SPD, in order to reduce beam-gas events [105].

For Pb–Pb and Xe–Xe collisions, the V0-AND trigger is polluted by electromagnetic events, i.e. interactions between the nuclei originating from interactions in the large electric fields associated with such ions. To get rid of these, extra constraints were applied, namely a threshold on the V0 amplitude in the 2015 Pb–Pb run and required hits in both ZDC detectors in the 2017 Xe–Xe run and 2018 Pb–Pb run.

For most of the data taking, the HM trigger was based on a threshold on the V0 multiplicity, but in particular in 2018, it was instead based on the SPD. In the latter case, at least 85 hits were required. Moreover, there was a cut on the number of V0 channels where the hits were flagged as beam-gas interactions (based on the timing) and it was required that neither the previous nor the next bunch crossing would have any activity, in order to avoid pile-up, which could otherwise erroneously be tagged as HM events.

Upon a successful trigger, the call from the CTP will start a series of read-out operations in each active detector (how this is done depends on the detector; for the TPC this is described in Chapter 5), until eventually the trigger input is recorded locally at a Local Data Concentrator (LDC). During processing, the read-out system will be busy. If so, a signal will be sent back to the CTP, which blocks the triggering until all detectors used in the partition are operational. During recording, selected portions of the data sent to the LDCs is duplicated and sent to the High Level Trigger (HLT), where it is processed and decided whether to store or reject the event. This is the final triggering stage, where also some compression is done by cutting out irrelevant information. If accepted, the selected part of the data from the full event will be sent

via the DAQ system to Global Data Concentrators (GDCs), where the whole event is built, from where it will further recorded at a Transient Data Storage. Part of the data from the event will also be processed online through data quality monitoring and detector algorithms to ensure good quality. If not passing the quality evaluation (done manually in Run 2, but this will be replaced with automatic procedures by the start of Run 3), the ALICE shift crew will be notified and the shift leader may decide to stop the run until the problem is solved. Finally, the data will be sent to GRID for final storage and post-processing and be copied to tape.

The offline processing is carried out over a distributed computing network called GRID. Here, for instance all tracks are reconstructed following the procedure described in Section 4.6, as well as candidates for weakly decaying particles (cf. Section 8.5.2). The processed event data is stored in Event Summary Data (ESD) files, which can be analysed directly by the analysis software. Due to the amount of information included, these are however very performance heavy to analyse, so a further processing stage is carried out, where the most relevant information for most analyses is selected. This information is stored in Analysis Oriented Data (AOD) files. The processing is then repeated in several stages, called analysis passes, where the detector calibration has been tuned between each pass[9]. Also, associated Monte Carlo (MC) simulation runs are carried out based on the detector and beam conditions of each run period, which is done in a similar way as for data.

While AOD files are used in most analyses, the main analysis described in Chapter 8 is based on ESD files, partly due to a few bugs which occurred during the processing into AOD files. Therefore, the most relevant information for our analyses was extracted from the ESDs and stored into data trees, which could be analysed offline.

Upgrade: Removing the Trigger

Upon the upgrade to continuous read-out, the current trigger system will be removed[10]. There will still not be enough capacity to store all events and therefore some selection will still be required, but this will not be done online (see below). The focus will now be on selecting rare processes. Moreover, a MB sample will be collected during a limited period each year using a weaker magnetic field than is currently being used, in order to increase the capabilities at low p_{T} [106]. For the MB sample, all events will be recorded, but during most of the data taking, the number of events will be reduced by a factor of $\sim 1.7 \cdot 10^3$.

As a consequence of the new trigger operation and increased data rates, the DAQ system needs to be replaced, both on the hardware and the software side. The processing will be done through the new Offline–Online system, O^2. This will use a temporary storage space with a capacity of 60 PB, where all detector information will be stored until the detector calibration for the specific run period is completed. The ESD files will be replaced by Compressed Time Frame (CTF) data, which will be used for creating AOD files, just as now done in two passes. However, unlike now, the AOD files will

[9] The tracking information is used in the calibration of the later passes (e.g. for tuning PID parameters, cf. Section 8.5), requiring that a coarse calibration has already been carried out for the previous pass.

[10] Some detectors, such as the TRD and the HMPID will still operate in triggered mode. These will use a dedicated hardware trigger, but it will not be nearly as extensive as the one currently used.

contain the full event information, but in a compressed format, so they can be used in all analysis. Consequently, the CTF files will be removed after the processing stage, in order to save storage space. The processing will be done at GRID, and in the case of pp data taking – which requires less resources than Pb–Pb data taking – part of the local O^2 computing resources will be used for this as well. Events which do not contain interesting physics will be discarded during the first pass of the offline processing.

The O^2 system will also replace the current analysis framework, AliPhysics. The reasons – apart from the new requirements with data files in a new format – are that it will use the computing resources more efficiently, it will enable multicore and GPU computing, which have seen much development recently, and it will make use of recent developments in the C++ language. These improvements are required with the immensely increased data volumes expected in Run 3 and beyond.

Chapter 5

Chip Testing for the TPC

A complete experimental scientific project should contain all steps of the scientific process, from deciding how to address the scientific question, develop and build the experiment, acquire and analyse the data, and finally interpret and publish the results. In a large project like ALICE this chain takes almost a scientific career and is far beyond the scope of a PhD thesis. As a fruitful compromise, I have had the opportunity to use the knowledge gained from analysing data from the present ALICE detector to contribute to the development of the detector upgrades, which is known as a service task. In Section 10.5 I will elaborate on how my present analysis will benefit from the upgraded detector.

5.1 SAMPA: a New Readout Chip Design

To be able to read out the analogue information produced by the TPC readout chambers and convert it to a useful format, which can later be stored to disk or tape, a fast analogue and digital electronic system is required. In Run 2, this was handled by a series of chips on a frontend card, as schematically shown in Fig. 5.1. This setup, which is based on the ALTRO (ALICE TPC Read-Out) chip, is however limited to a sampling window of 100 µs [107], matching the gated operation of the TPC. Therefore, for continuous readout, a new system for the readout electronics has been developed.

The central unit of the new readout system is the SAMPA chip [108]. This contains both the preamplifier/shaper and the Analogue-to-Digital Converter (ADC) and data processing units, i.e. it basically merges the functionality of the PASA and ALTRO chips. An overview of the SAMPA chip design is shown in Fig. 5.2. The signal from the TPC comes as a current pulse. This induces a charge on a capacitor in the preamplifier feedback loop, which creates a voltage pulse. The preamplifier is followed by a shaper, which gives the signal a smooth, semi-Gaussian shape, integrating the signal on a time scale relevant for the duration of the signal in the detector. In the ADC, the analogue voltage pulse is converted to a code with 10 digital bits (1024 channels) using successive approximation, at a sampling rate which can be set as high as 20 MHz. There is an option to compress null samples on the chip using a series of data processing units, but due to the low fraction of null samples in the TPC, this is not used here. Instead,

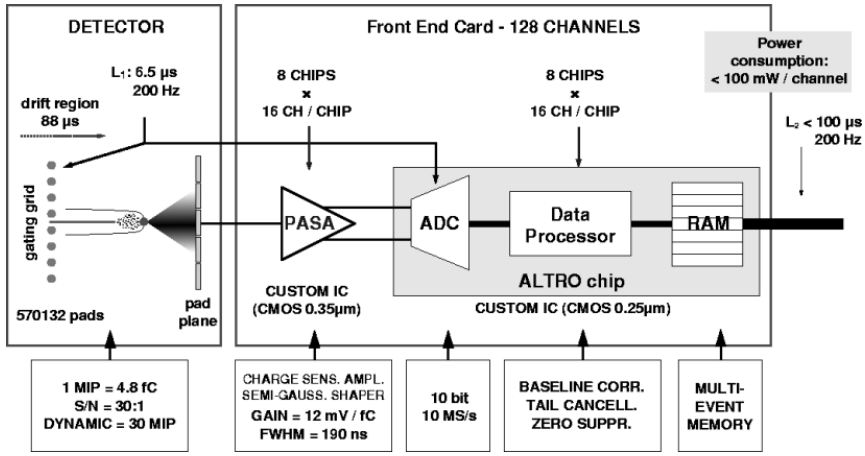

Figure 5.1 – Block diagram of the current TPC frontend electronics. The detector signal from each channel is sent to a preamplifier/shaping amplifier (PASA). The output from the PASA is converted to a digital signal and processed at an ALTRO chip. A local RAM is used to store data taken during a time window of 100 µs, before it is sent via link to the data acquisition and detector control systems. Image taken from Ref. [107].

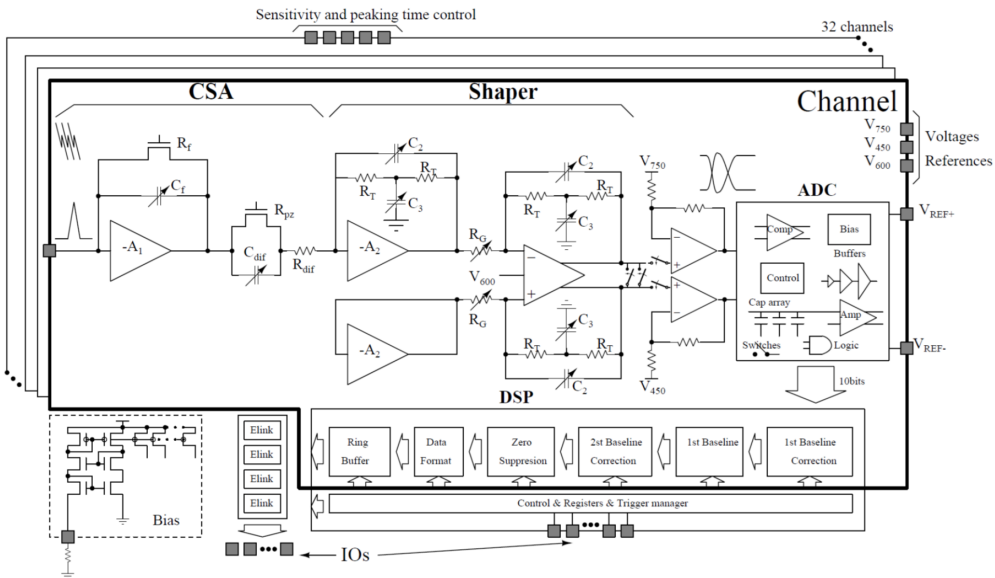

Figure 5.2 – Block diagram of the SAMPA chip design. This contains a preamplifier (Charge Sensitive Amplifier, CSA), shaper, ADC, and successive data processing units. Figure taken from Ref. [108].

all samples are transferred via 10 parallel fast-connection links and the processing is done afterwards, both at hardware and software level. An advantage of the SAMPA chip is that it has a flexible design, making it possible to tune several parameters such as shaping time and gain. Thus it can be used for both the TPC and MCH readout systems, although the detector characteristics are quite different.

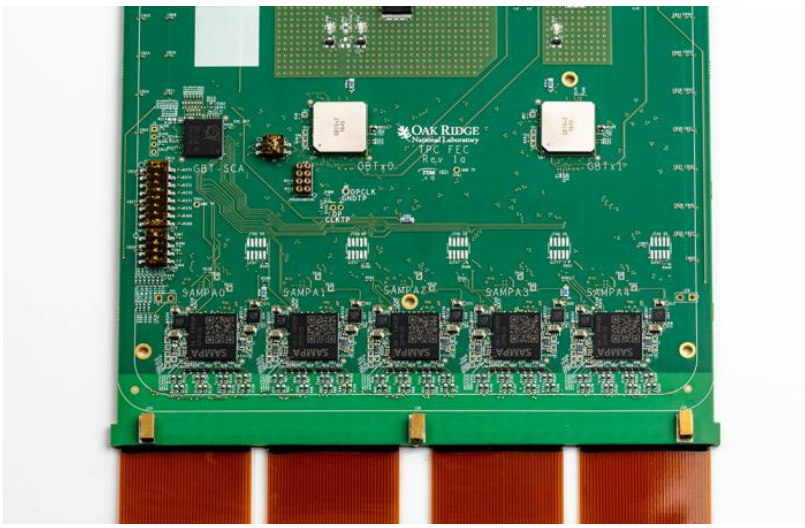

Figure 5.3 – Photograph of a frontend card. The five SAMPA chips are mounted along the side facing the Kapton cables connected to the readout chambers (bottom of the image). Figure taken from Ref. [109].

The SAMPA chip is packaged in a ball grid array (BGA) package, with dimensions of 15×15 mm^2. Each chip contains 32 channels corresponding to different cathode readout pads. The SAMPA chips are mounted on frontend cards, each containing five chips, cf. Fig. 5.3 (see Fig. 5.5b for an individual chip). A total of ∼17 000 chips are used for the TPC and ∼37 000 for the MCH.

Usage of SAMPA Chips in the Muon Tracking Chambers

The muon tracking chambers use cathode pad chambers based on MWPCs for tracking the muons [94]. However, the current readout system is unable to handle the increased event rates during Run 3, and therefore the frontend electronics and readout chain are upgraded to allow for continuous readout [110]. For the new frontend electronics, the SAMPA chips have been selected. Due to a quite different environment than for the TPC, the specifications for the chips used for the MCH are not the same as for the TPC. In particular, the gain will be set to 4 mV/fC for the MCH, compared to 20 or 30 mV/fC for the TPC, and most of the data processing options not used in the TPC are used here. See the next section for specification details.

5.2 Testing Overview

The Lund ALICE group is responsible for verifying that the chip works as intended and ensuring that the chips follow the specifications – both in the prototyping phase and for every chip which will be mounted in the detector – which requires extensive testing. Consequently, the test setup saw some development between the testing phases.

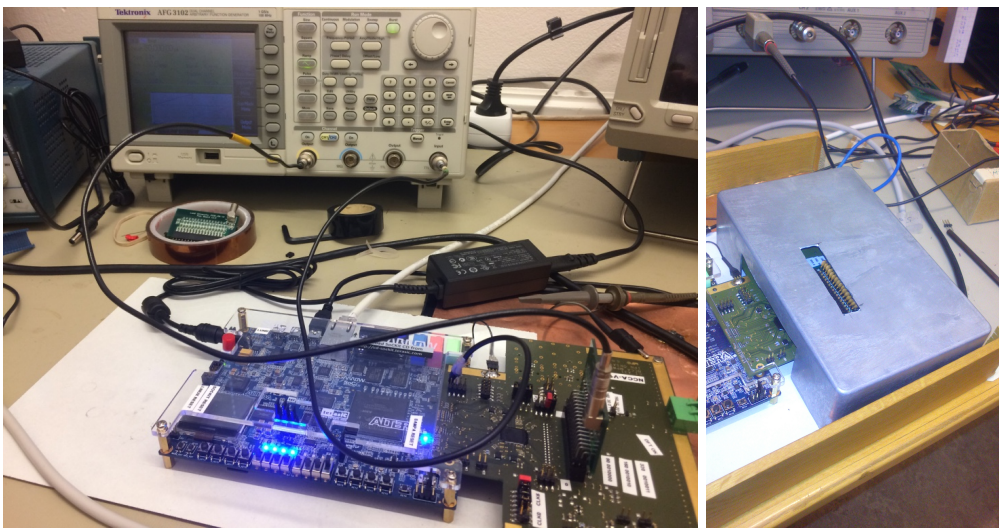

(a) **(b)**

Figure 5.4 – Test setup for characterising the SAMPA prototype. The chip is soldered onto an NCCA board in the lower right corner of panel **(a)**. This is mounted onto a SoCKit board to its left. A pulse generator has its input channel connected to the sampling clock of the NCCA and its output connected to the injector card on top of the NCCA board, which converts the voltage pulse to a charge pulse that is sent to the chip. The data is transferred to a computer (not shown in the figure) via an Ethernet cable. In panel (**b**) an aluminium box is placed over the NCCA board for noise suppression. Here, the connector to the pulse generator is replaced by a series of capacitors, but the box was normally used also during pulsed measurements.

The initial setup, shown in Fig. 5.4a, was mainly used for developing the testing code and verifying basic chip functionality. Here, a SAMPA chip prototype was soldered onto an NCCA[1] board, which was mounted to a SoCKit board – a programmable hardware platform accessible through the FPGA language, which enables fast data transmission [111]. The NCCA board had a 32-channelled port connected to the SAMPA chip, where an injector card could be mounted to send a test pulse to the chip. Depending on the test parameter, a few different injector cards were used to be able to switch between pulsing all channels or just a subset of them (cf. Section 5.3). Alternatively, a set of capacitors could be connected to the port, which was done for noise measurements. A capacitor acts as an antenna, so when increasing the capacitance, the noise level will

[1] Negatively polarised Circuit Card Assembly.

increase. For optimal performance (and to mimic the conditions in the detector), one should however reduce the noise as much as much as possible, which requires shielding. This was done by placing an aluminium box over the equipment, cf. Fig. 5.4b, which acts as a Faraday cage. To synchronise the pulsing with the readout, another port at the NCCA board was connected to the sampling clock. The pulsing was done by using a pulse generator, which produced a voltage step function which could be adjustable by using a specific input file. This used the signal from the sampling clock as trigger. The pulse generator was connected to the injector card, which converted the voltage pulse to a charge pulse, which was injected to the chip.

The readout control was handled by a computer interface made by our colleague Arild Velure in the Bergen group. During measurements, the output from the chip was sampled by an FPGA programme and sent to the computer through an Ethernet connection. The output was stored in a data tree which could be analysed through the ROOT package.

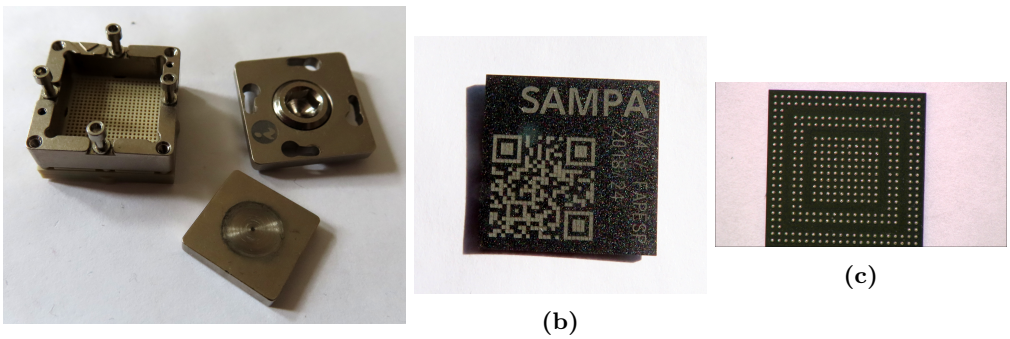

(a) **(b)** **(c)**

Figure 5.5 – **(a)** Socket used for placing chips for testing, along with a pressure plate and a lid. The pressure plate was used to keep an even pressure on the chip, which was ensured by tightening the screw in the lid. **(b)** Backside of a SAMPA chip. The QR code is used for identification. **(c)** The other side of the chip. The pins are connectors to the electronics.

In the second phase, the NCCA board with the soldered chip was replaced by one where a socket box (cf. Fig. 5.5a) was used in its place. Here, the SAMPA chip could be placed in the box. Connection was ensured by pins in the bottom of the box being connected to the connector pins of the SAMPA chip (Fig. 5.5c). A slight tension was applied to ensure good connection. This setup enabled testing of multiple chips, which was necessary to characterise chip-to-chip variations.

Finally, for testing individual chips, an automatic procedure involving a robot was required. This required to rewrite the testing code to executable files (instead of relying on ROOT macros as was the case before) which could be used in the automatic routine. The procedure was as follows: First, a set of ~ 500 chips was placed manually on trays next to the robot. The robot (Fig. 5.6) then used a suction device to lift the chip and move it to a small tray close to one of the test stations[2], in order to align it, and

[2]Two test stations were used for the testing.

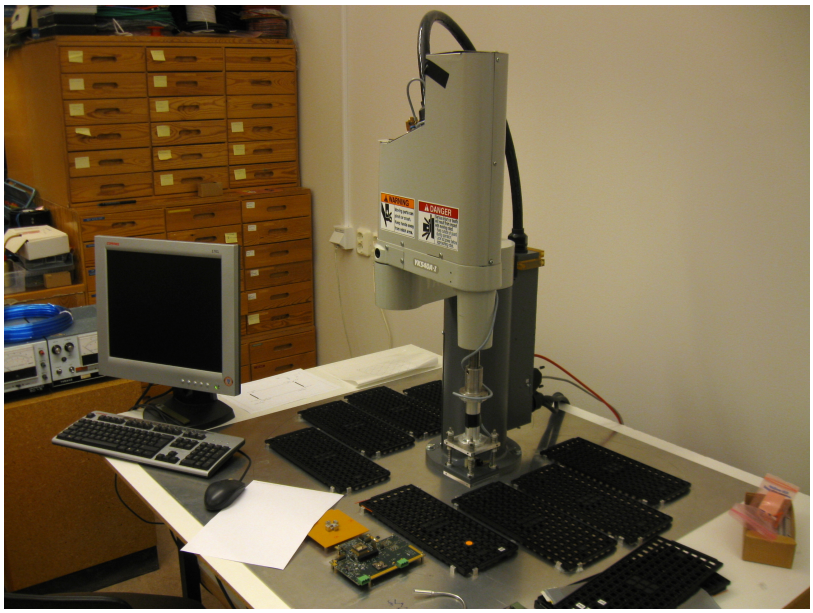

Figure 5.6 – Robot used for the automatic testing, before setting up the full testing apparatus.

further moved it to a PCCA board similar to the one with the socket box used for the prototype testing. A pressure plate was added to the socket to apply tension to the chip, as before, and pressure was ensured by a special pressurising wagon (Fig. 5.7), which used rails to move it to the test socket. Following this, a test routine was run using a few different pulse sequences for the individual tests (here odd and even channels were connected to different output channels in the pulse generator using a similar connector as before). In the meantime, the robot was moving another chip to the second test station, so the two chips could be tested in parallel. The output from the chips was written to text files, which were analysed by the executables, extracting the relevant information for each channel in the chip. If passing the specifications, the chip was moved to one tray, and if rejected it was moved to another one. The rejected chips were typically retested to make sure the failure was not due to e.g. bad connection with the socket, external noise, or instabilities in one of the test routines[3]. An overview of the test station is shown in Fig. 5.8 and a demonstration of the robot testing is available at `https://www.youtube.com/watch?v=r8HM-wFboOM&feature=youtu.be`.

The analogue parameters tested were gain, rise time, cross-talk and noise. Moreover, a couple of digital tests built into the SAMPA design were carried out. These include verification of all memory cells and registers, toggling flip flops, and verifying digital inputs/outputs by JTAG. Altogether fabrication faults in the BGA packaging should be found for almost all signals and DC connections (as long as the same signal is not routed several times in parallel through the package, like supply voltages and grounds).

[3]This was the case a few times, and consequently some of the test programmes had to be rewritten during the testing.

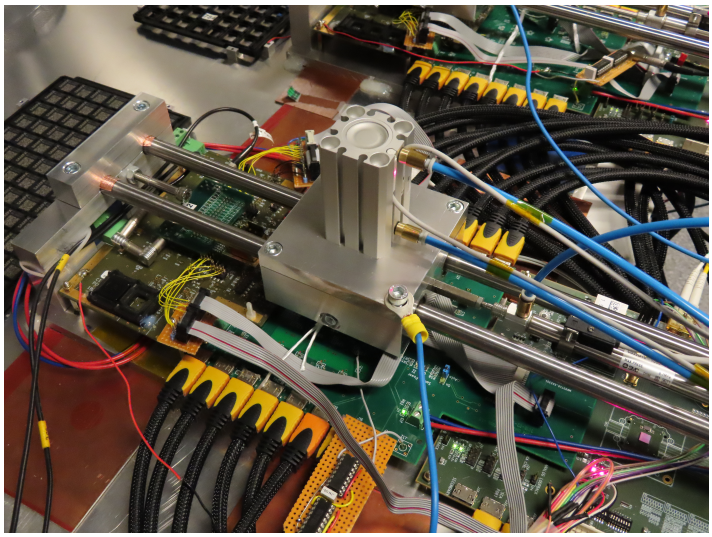

Figure 5.7 – View of one of the test stations. The tower-like structure in the middle is a pneumatic wagon used for applying an even pressure to the chip in order to ensure good connection. This can move freely on the metal rails, and its position is ensured by a stopping screw between the rails (to the left in the image).

Table 5.1 – Summary of a selection of the SAMPA chip specifications at TPC and MCH settings [112].

Specification	TPC setting	MCH setting
Reference voltage	1.25 V	1.25 V
Detector capacitance	18.5 pF	40-80 pF
Peaking time (shaper)	160 ns	320 ns[a]
Gain	20 or 30 mV/fC	4 mV/fC
Sampling frequency	5 MHz[a]	10 MHz
Equivalent noise charge	$< 600\,\mathrm{e}$	$< 950\,\mathrm{e}$ (40 pF) $< 1600\,\mathrm{e}$ (80 pF)
Cross-talk	$< 0.3\%$	$< 0.2\%$

[a]Changed from the original specification.

A summary of the specifications for most parameters, for both TPC and MCH setting, is given in Table 5.1. The following sections describe how the (external) tests were carried out, both in the prototype characterisation and the robot testing. During the robot testing, it was ensured that the capacitance over the chip was lower than during operation in the detector. The sampling frequency was set to 5 MHz due to technical limitations of the equipment. For most measurements, the peaking time and gain were set according to the specifications (i.e. for the TPC, measurements were carried out both at 20 and 30 mV/fC).

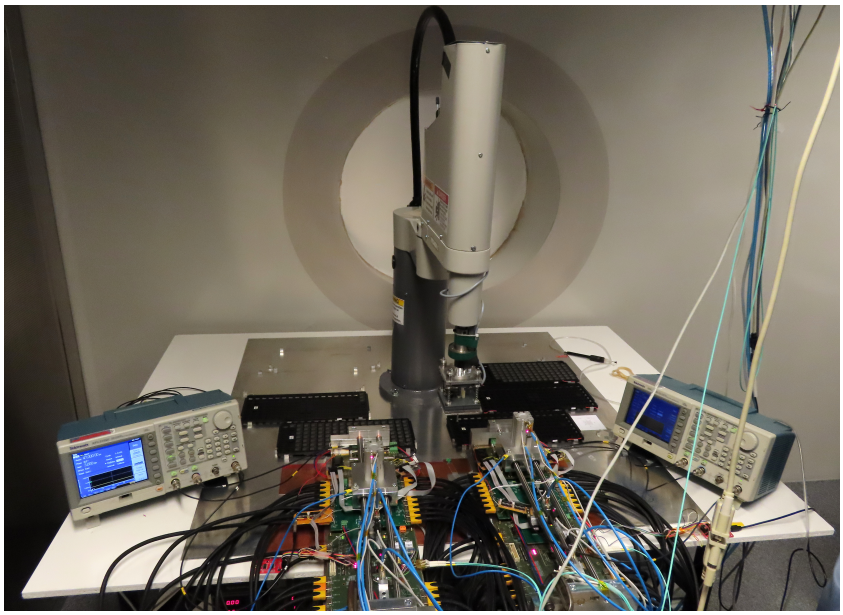

Figure 5.8 – Overview of the test setup for the automatic testing, with the robot in the middle and the two test stations in the front.

5.3 Cross-Talk Measurements

Capacitive coupling between charge input channels in an analogue frontend card is expected on some level, and therefore they may pick up signals from neighbouring channels on the chip. This effect is known as *cross-talk*. It is important that the cross-talk is not too large, since otherwise a signal in one channel may trigger a signal above the threshold in another one, effectively triggering a false measurement. Acceptable cross-talk limits are given in Table 5.1.

To measure cross-talk, a large pulse was sent to a subset of the channels and the output in the other channels was measured as a fraction of the pulse amplitude in the pulsed channels. This was done for several different pulsing configurations; during the prototyping phase the following were used: pulsing all odd channels, all even channels, every third channel, just a single channel, and all channels but one. In the prototyping stage, the test pulse card, which is a source of cross-talk in itself, was disconnected from the non-pulsed channels. In the production testing, this could not be done and cross-talk measurements were reduced to be sensitive to short circuits (e.g. touching bond wires) between neighbouring channels. Here a simultaneous test pattern of all even or all odd channels was always used.

Two important aspects are to synchronise the measurements so they are done at the peak amplitude and to use a large statistical sample to suppress asynchronous noise. The former turned out to be a challenge during the characterisation phase, since the data transferred through the link was packaged such that the output from three consecutive channels was bunched together, but the transfer rate was different for different packets,

resulting in a quite disordered output file for the analysis. During the robotic testing, the output from all channels was written to a binary file (which was later decoded into a text file) before sending the next pulse, so this problem was not present here.

The measured parameter is ADC output at the peak time minus the average baseline value (measured outside the peak time window). Some results from the prototype testing are shown in Fig. 5.9. First, one should notice that these measurements have been performed on a single chip. Therefore, any large differences in a single channel between measurements where a similar number of channels have been pulsed are likely due to differences in measurement conditions. This is the likely explanation for the significantly lower cross-talk levels observed when pulsing odd channels compared to even pulsing. Moreover, a large deviation is observed in channel 30 when pulsing every third channel, with a crosstalk nearly a factor five greater than any other similar measurement. This is likely due to cross-talk within the connector rather than within the chip, and therefore this measurement is not shown in the figure. Apart from these caveats, the cross-talk levels are similar when pulsing every third channel as when pulsing every other channel, and therefore the former measurements are redundant. When pulsing all channels except for one, the levels are significantly higher, but it is unlikely that this situation would ever occur within the detector, so due to the challenge of setting up this measurement for automatic testing, this was not investigated further. The most common condition would be a pulse in a single or just a few channels, which generates much

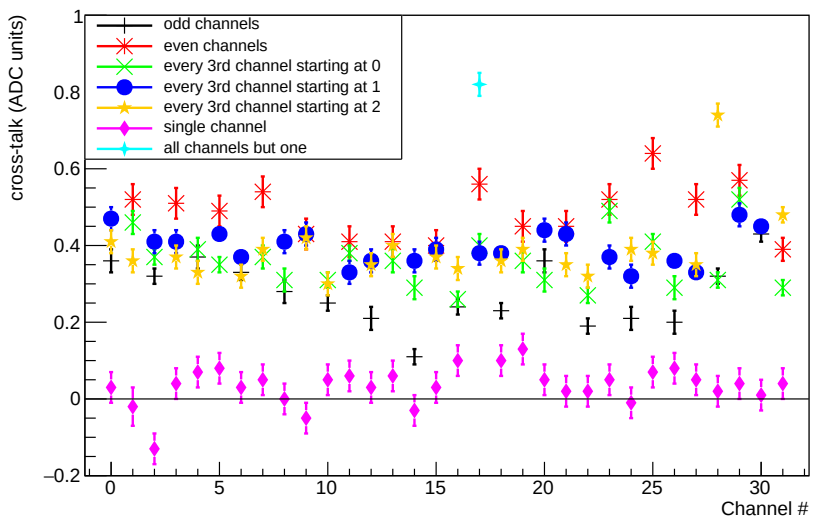

Figure 5.9 – Measured cross-talk (in ADC units; the amplitude of the test pulse was about 600 ADC units) for the first version of the SAMPA chip for a few different configurations. The legend indicates which channels were pulsed. For the single channel measurement, channel 17 was pulsed, and this was the only channel tested with all the others pulsed. In the case where every third channel starting at channel 2 were pulsed, channel 30 had a cross-talk of ~ 3 ADC units, which is outside the range of the plot.

lower cross-talk levels (and only significantly different from zero in nearby channels). Therefore, if the odd and even pulsing yields acceptable results, one can conclude that the cross-talk is acceptable for any situation within the detector.

During the robotic testing, the cross-talk was normalised to the peak amplitude. The chip was rejected if this exceeded 5% in at least one channel, and also if it was $< -1\%$, which could indicate other errors. These limits were set quite loose since other cross-talk sources were present in the test set-up, in particular from the connector itself. The cross-talk measured at the front-end cards satisfies the specifications.

5.4 Noise Measurements

The noise measurements are quite straightforward. Here, simply the standard deviation of the output from the ADC, σ_{ADC}, in each channel was measured, with no signal at the input. This can then be converted to Equivalent Noise Charge (ENC) in electrons by using

$$\mathrm{ENC} = \frac{\sigma_{\mathrm{ADC}} \cdot 2.2\mathrm{mV/ADC}}{g \cdot e}, \tag{5.1}$$

where g is the gain factor measured in mV/fC and $e = 1.602 \cdot 10^{-4}$ fC is the elementary charge. There is however one problem with this approach, which will be discussed carefully in Chapter 6, and that is that the measured output from the ADC is discretised, so if the analogue noise is low compared to the ADC step size, the measured σ_{ADC} will be subject to a substantial discretisation effect. Therefore, the ENC cannot be measured accurately in this way using a too low gain setting (in particular 4 mV/fC which will be used in the MCH), although, as will be shown, a gain of 20 mV/fC is enough. Consequently, this higher gain setting was used for the noise evaluation of the chips used in the MCH during the automatic testing.

In the performance tests, noise was measured as a function of input capacitance, which was done by connecting a capacitor from each input channel to ground, as described in Section 5.2. One should note that there is a stray capacitance in the chip and in particular on the board, so this approach is not perfect. Some results for 30 mV/fC are shown in Fig. 5.10. This chip satisfies the requirement of < 600 electrons at low capacitance (expected conditions in the detector). The channel-to-channel spread increases when increasing the capacitance, which could be due to many different reasons, such as the surrounding environment, or how things are connected. Evidently the point at 0 pF does not match an extrapolation from the other data points. There could be several reasons for that. One contribution is that the stray capacitance in the system is non-negligible and also that the digital noise on the board contributes. It is difficult to say how large this is not knowing exactly how the circuit is connected, but it should be roughly as large as the external capacitance at the end of the linear regime, i.e. $\sim 15 - 20\,\mathrm{pF}$.

During the automatic testing, all noise measurements were carried out using a gain of 20 mV/fC. The automatic setup was not ideal for this measurement since the test pulse card (connected to the pulse generator) had to be connected during the measurement. Also the environment was noisier due to the servo motors for the robotic movement. Here the limit was set to 1.3 ADC units (≈ 900 electrons) for chips used in the TPC.

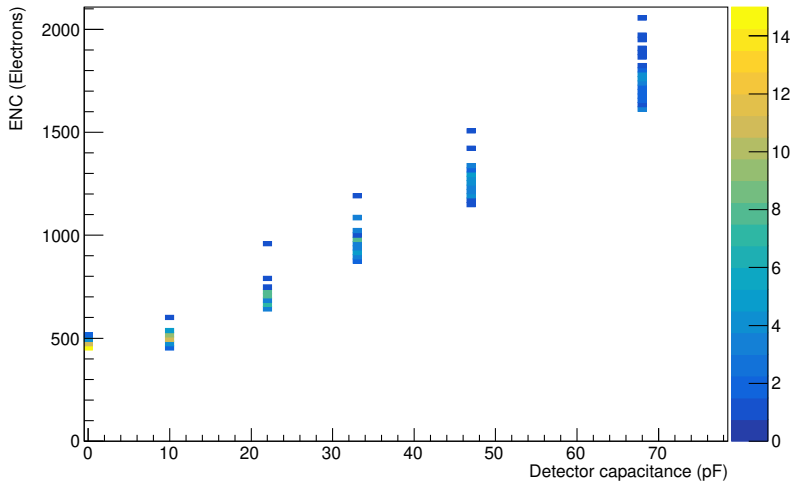

Figure 5.10 – ENC as a function of external capacitance for a V2 SAMPA prototype, using a gain of 30 mV/fC. The colour indicates the density of channels in a narrow ENC interval.

For the MCH, the limit was set to 1.6 ADC units for one of the test stations and 2.0 ADC units for the other one, due to a noisier environment. Channels 30 and 31 picked up more noise, so here looser limits were used. There was also a lower limit of 0.6 ADC units ($\approx$ 400 electrons). A lower noise level than this could be due to stuck bits or a lost bond wire, but then of course, the chip would fail other tests as well.

In connection with the noise tests, the pedestal voltages were tested as well. These were required to lie between 40 and 110 ADC units (90-240 mV), as to not interfere too much with the detection range.

5.5 Rise Time Measurements

Following Table 5.1, the rise time of the analogue pulse, which is related to the shaping time, should be $\sim 160\,\text{ns}$ at TPC setting and $\sim 320\,\text{ns}$ at MCH setting. The hardware and firmware used for the automatic testing limited this to a sampling rate of 5 MHz, i.e. a sampling interval of 200 ns, which is too wide for measuring the rise time to adequate precision. To circumvent this limitation, a few different approaches were tested.

First, a $\Gamma(4)$ fit function was fitted to the sampled peak. This is on the form

$$f(t) = \begin{cases} b, & t < t_0, \\ b + A \cdot \exp\left(-4\left(\dfrac{t-t_0}{\tau} - 1\right)\right) \cdot \left(\dfrac{t-t_0}{\tau}\right)^4, & t \geq t_0, \end{cases} \tag{5.2}$$

where t_0 is the start time of the pulse, b is the baseline (can usually be fixed to the average measured value at $t < t_0$), A is the amplitude, and τ is the rise time. While

this provides a good description of the pulse shape, it turned out that the measured rise time following this approach is biased by the position of the sampling points in the peak region, cf. Fig. 5.11, which led to the conclusion that a more clever approach was required. The solution was to use several pulses, where each pulse is shifted relative to each other by a known delay – 20 ns turned out to be enough – and use this information to reconstruct the pulse shape.

At first, this was done by taking several measurements synchronised with the

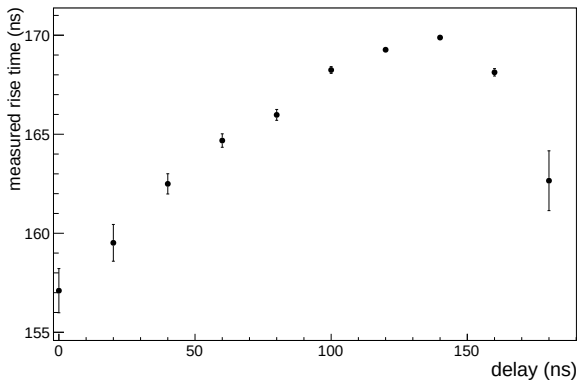

Figure 5.11 – Rise time extracted from a fit, using a sampling window of 200 ns, shown as a function of the delay of the input pulse. The true rise time for this measurement is ≈ 175 ns, so this approach yields neither precise nor accurate information.

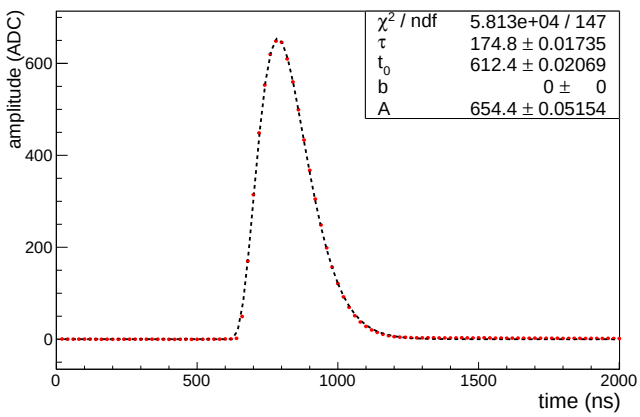

Figure 5.12 – Extracted pulse shape from a measurement of a chip using TPC settings, i.e. a nominal shaping time of 160 ns, and a gain of 30 mV/fC. The data points are marked by dots and the fit by the dashed line. The top right corner shows the extracted fit parameters (τ and t_0 are in ns and A and b are in ADC units) as well as the quality of the fit. This is measured after baseline subtraction, so this parameter is fixed to zero.

sampling clock, but applying the shift to each measurement. This has the advantage that the pulses can be sufficiently spaced in time to not affect each other, but it turned out to be a tedious measurement and difficult to obtain stable conditions, so this approach was abandoned. Instead, a pulse train of charges produced from 10 identical step voltages was sent to the chip, each separated by 3 µs and shifted by 20 ns, reaching an effective sampling rate of 50 MHz. To get sufficiently large statistics, a few hundred pulse trains were sent to the chip during each measurement. A signal with negative amplitude was used to mark the start of each pulse train. With this refined measurement, the fit function, Eq. (5.2), could be applied without any measurable bias. An example of a measurement during the automatic testing is shown in Fig. 5.12. While the χ^2/d.o.f. in this example indicates a poor fit to the data, one should keep in mind that the statistical sample is large, so any systematic disturbance will have a negative impact on this. Upon close inspection the dominant source is likely a slowly decreasing

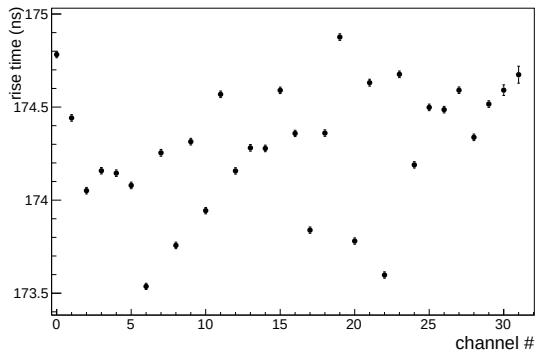

(a)

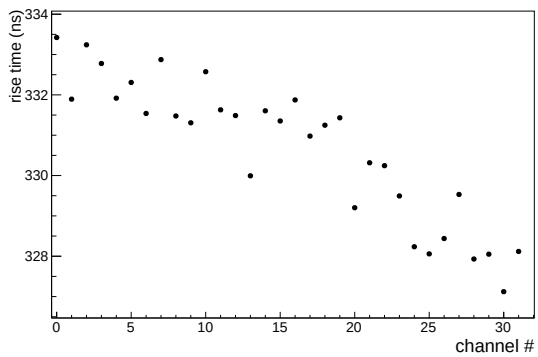

(b)

Figure 5.13 – Rise time as a function of channel number for two different chips, using **(a)** a gain of 30 mV/fC and a nominal rise time of 160 ns, and **(b)** a gain of 4 mV/fC and a nominal rise time of 320 ns. In both cases, the channel-to-channel spread is less than 2%.

tail far away from the peak (at $t \gtrsim 1250\,\text{ns}$), which is not accounted for by the fit function. Moreover, the onset of the pulse is slightly more abrupt than indicated by the fit. These deviations should not affect the rise time significantly (and if they do, it happens in a systematic way), so one can conclude that this fit function describes the pulse shape adequately.

Typically, the measured rise time was slightly higher than the design value, but one should keep in mind that its definition in a design perspective may be slightly different than what is obtained from the fit. Interestingly, the value was dependent on the gain setting to some degree. There was also a significant spread between the channels on a chip, although still acceptable. Examples for two different chips, one with TPC setting and one with MCH setting, are shown in Fig. 5.13. Acceptable values were set to $155 < \tau < 180\,\text{ns}$ for the TPC (measured at 20 mV/fC) and $300 < \tau < 350\,\text{ns}$ for the MCH.

5.6 Gain Measurements

The gain could in principle be derived from the fit parameter A in Eq. (5.2), which was also done as a complementary measurement. This measurement was however too tedious for checking both gain and linearity, and therefore another approach was chosen. The approach was to apply pulses of different amplitudes to the chip. The pulses in the train were synchronized with the sampling clock such that each sample was situated at the peak value. The relation between input charge and digitised peak amplitude was then analysed. During the robotic testing[4], this was done using a pulse

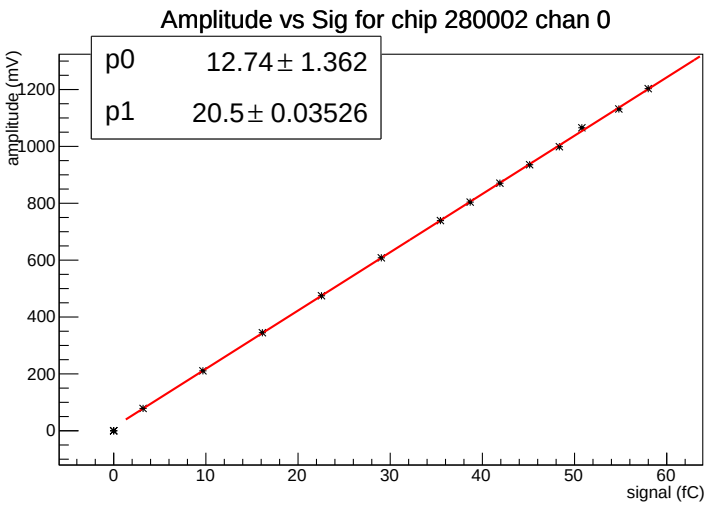

Figure 5.14 – Example of a calibration curve for a gain measurement of a chip with a nominal setting of 20 mV/fC. The parameter "p1" indicates the slope of the curve.

[4]Since I was not involved in the gain measurements during the characterisation phase, this part is not included in this thesis.

train similar to the one used during the rise time measurement, except that instead of different trigger delays, the different pulses had different amplitude. A linear fit was applied to the measurements, which determined the gain factor of each channel in the chip, see Fig. 5.14 for a calibration curve. To begin with, a measurement at a setting of 20 mV/fC was used for both the TPC and the MCH. Acceptable values were set to $18.5 < g < 20.7\,\mathrm{mV/fC}$. To further test if the gain was correct at other settings, a measurement was done at 30 mV/fC for both detectors, with coarse limits of $25 < g < 35\,\mathrm{mV/fC}$. For the MCH, an additional measurement was carried out at 4 mV/fC, with limits of $3.9 < g < 4.4\,\mathrm{mV/fC}$.

5.7 Bit Issues: the Odd-Even Effect

Ideally a 10-bit ADC like the one used in the SAMPA chips divides the analogue range into 1024 equally wide intervals. In reality the intervals are defined by analogue voltages to comparators, resulting in some variation of the interval widths. This is most sensitive for the comparator defining the least significant bit. For the ADC in the SAMPA one can see an effect that the odd and even ADC codes are not equally populated, as shown in Fig. 5.15. Therefore this effect has been dubbed the "odd-even effect". To some extent it will be present in all ADC designs which aim at reaching maximum sensitivity. To limit the voltage division to a level where it responds 100% correctly would sacrifice sensitivity which is not desired by the SAMPA users in ALICE.

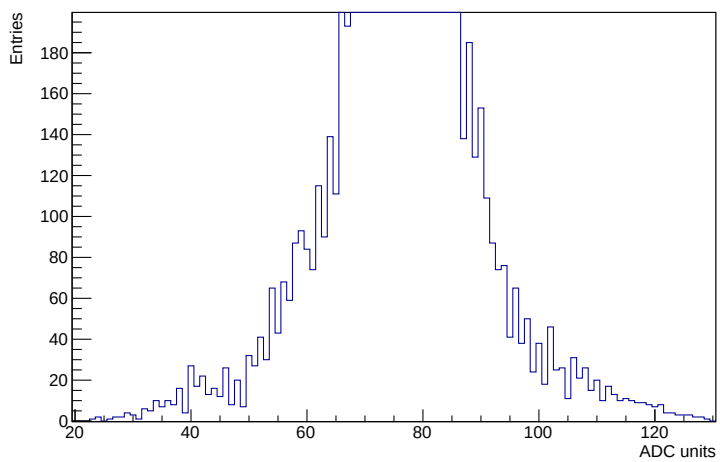

Figure 5.15 – Histogram over number of entries per ADC code intended for a cross-talk measurement, but due to inadequate noise suppression, there is a large spread around the pedestal level at ≈ 75 ADC units. While this caused issues with the measurement (it had to be redone with proper shielding), the increased spread clearly demonstrates the odd-even effect, i.e. a larger number of entries with an even ADC count than for neighbours with an odd count. Some periodicity with every four and eight bins can also be seen, indicating higher-order bit issues.

A similar effect (perhaps with different origin) is that ADC codes spaced by powers of 2 (and in particular 16) are less populated than their neighbouring channels. The net effect is small such that there is a shift of one ADC channel (at most $\sim 1‰$) for the affected ADC values.

Both these effects are of minor importance for the use of the SAMPA in real data taking since the information is vastly oversampled either by many samples per pulse as in the MCH case or many pad values per track as for the TPC. Supported by physics performance simulations of the detectors it has been found that the benefit of improving the ADC performance on these points was not necessary, and the required effort would have risked missing the target of having all SAMPA chips ready for the TPC upgrade during Long Shutdown 2. However, the odd-even effect has an impact on the noise performance results, as will be shown in Chapter 6, where it also will be described how to deal with it.

The cause of the odd-even effect is a bit unclear, but it clearly originates from the hardware (but not necessarily the chip design). The ADC used in the SAMPA chips is based on successive approximation, meaning that it compares the amplitude of the analogue signal from the shaper (at the sampling time) with a reference voltage in several successive steps (here 10) [113, pp. 143-146]. If it is higher than the reference, a one is set, otherwise a zero. To be able to do this, a sample of the signal is stored in a capacitor and this charge is compared to an analogue reference signal obtained through a digital-to-analogue conversion. After each comparison, the capacitor is discharged with a charge equal to the reference signal, so only the remainder remains. In each step, the reference voltage gets successively smaller. Ideally, each voltage should be exactly a factor two smaller than the previous one, so the signal gets divided in two equal halves, but this cannot be the case here. Therefore, I personally suspect that the reference voltages in the ADC deviate from their design values, such that the voltage division is not perfect. This should result in an effect that is similar to what is observed.

Chapter 6

Noise Measurements of Chips for the Muon Chambers

6.1 Introduction

When converting the analogue signal to digital data, one will inevitably lose information due to the discretisation process. This will affect the resolution in time as well as in amplitude. In Section 5.5, it was described how one can improve the temporal resolution in order to accurately measure the shaping time of the chip. This chapter will instead focus on the amplitude. If the ADC resolution is too low, it will affect the noise measurements. In particular, the discrete distribution will deviate significantly from the underlying analogue one, meaning that the measured (discrete) standard deviation will no longer correspond to its analogue counterpart, invalidating the approach taken in Section 5.4. In the following sections, it is described how one can circumvent this problem and obtain useful results also if the analogue noise is very low compared to the level spacing in the ADC.

If the gain is sufficiently high, the spread over the discrete ADC levels due to noise will become large enough for it to describe the analogue distribution, which in the case of white noise is Gaussian. This is clearly the case for the TPC settings, where a gain of 20 or 30 mV/fC is used. For the muon chambers, however, using a gain of 4 mV/fC, this is not the case. Following Eq. (5.1), the desired maximum noise level of 950 electrons corresponds to a standard deviation of $\sigma_{\mathrm{ADC}} = 0.28$ ADC units, meaning that most of the noise data will be distributed over just one or two ADC channels. Therefore, the Gaussian approximation is not valid here. The impact of this is illustrated in Fig. 6.1, which is obtained using data from the same phase of the prototype testing as Fig. 5.10. Here, the measured noise level seems to stabilise at a fairly large value at low capacitance, of roughly 0.4 ADC units, corresponding to 1400 electrons. Even more problematic, the spread between channels is very large. Both of these effects are largely due to the discretisation process, which shows the importance of correcting for this effect.

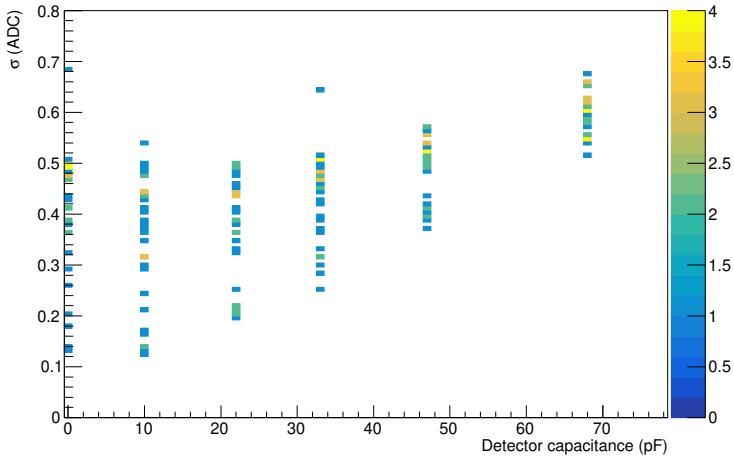

Figure 6.1 – Measured noise (standard deviation, σ) in ADC units for a SAMPA prototype, as a function of detector capacitance, using a gain of 4 mV/fC. The colour indicates the density of channels in a narrow σ interval.

6.2 Correlation Between Measured Baseline and Noise Levels

A noise level of 0.4 ADC units should be close to the limit of what can be reliably measured. At this level, assuming a Gaussian distribution, the discrete output will only occupy a few (2-4) different ADC bins for a single channel. Therefore, the measured standard deviation σ will depend on where the float remainder of the mean value μ is located, or mathematically μ modulo 1. In the extreme case of only having occupation in two bins, the following estimates can be made:

$$\mu_0 \equiv \mu \bmod 1 = \frac{N_1}{N_0 + N_1}$$

where N_0 is the occupation in the lower bin, and N_1 is the occupation in the upper bin, and

$$\sigma = \sqrt{\frac{N_0\mu_0^2 + N_1(1-\mu_0)^2}{N_0 + N_1 - 1}} \simeq \sqrt{\mu_0(1-\mu_0)} \tag{6.1}$$

for large samples. This gives a strong correlation between μ and σ.

To test whether a correlation between μ and σ exists for the data set used to obtain Fig. 6.1, the measured noise in each channel was plotted against the baseline modulo 1, for each capacitance. The results are shown in Fig. 6.2, where the result for the limiting case of two bins (Eq. 6.1) also is included. Indeed, there is a strong correlation between μ and σ for low capacitances (low noise), bending around the limiting case, proving that the discretisation effect plays a large role in the obtained results. At higher capacitance, where σ is further away from the limiting case, the effect is less obvious.

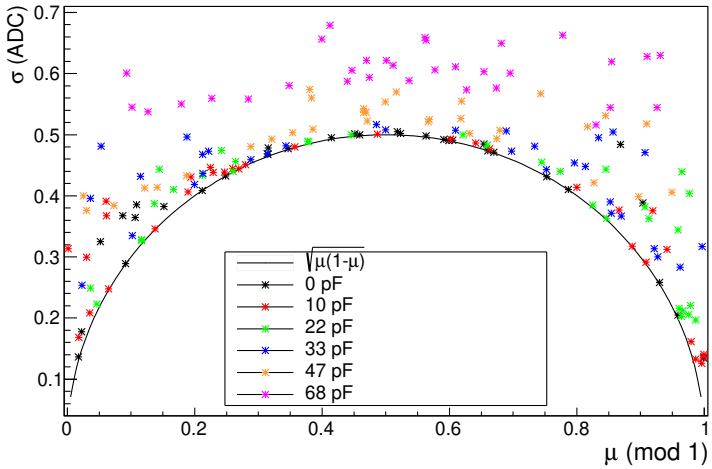

Figure 6.2 – Measured noise (σ) as a function of mean modulo 1 (μ_0), for different capacitances. The solid line indicates the limiting case of only having data in two adjacent ADC bins.

Now remains to find a method to get around this effect and extract the analogue part of the noise.

6.3 Models for Discretisation

6.3.1 Ideal Model

The expected behaviour of an ideal ADC, is that it simply truncates the analogue value (incoming charge) to the closest smaller discrete ADC level. If the equipment is well-shielded, as is supposed in this case[1], the analogue noise is expected to behave like white noise, following a Gaussian distribution. This has the density function

$$f_A(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right),$$

where μ and σ are the analogue mean and standard deviation, respectively. After truncation, the corresponding discrete density function f_D is obtained by integrating the analogue density function between each ADC level and the next one, i.e.

$$\begin{aligned} f_D(x) &= \sum_{n=-\infty}^{\infty} \delta(x-n) \int_n^{n+1} f_A(x')\mathrm{d}x' \\ &= \sum_{n=-\infty}^{\infty} \frac{1}{2}\delta(x-n)\left(\mathrm{erf}\left(\frac{n+1-\mu}{\sigma\sqrt{2}}\right) - \mathrm{erf}\left(\frac{n-\mu}{\sigma\sqrt{2}}\right)\right), \end{aligned} \tag{6.2}$$

[1]This seems to be true for most channels, but as will be seen, channel 30 is a notable exception.

where x is in ADC units and δ denotes the Dirac delta.

From this, the corresponding discrete mean μ_D and standard deviation σ_D may be calculated. Starting out from the definitions of the expection value and standard deviation, yield

$$\mu_D = \mathbb{E}_{f_D}[x] = \int_{-\infty}^{\infty} x f_D(x) \mathrm{d}x = \sum_{n=-\infty}^{\infty} \frac{1}{2} n \left(\mathrm{erf}\left(\frac{n+1-\mu}{\sigma\sqrt{2}} \right) - \mathrm{erf}\left(\frac{n-\mu}{\sigma\sqrt{2}} \right) \right),$$

and

$$\begin{aligned} \sigma_D^2 &= \mathbb{E}_{f_D}[(x-\mu_D)^2] = \int_{-\infty}^{\infty} (x-\mu_D)^2 f_D(x) \mathrm{d}x \\ &= \sum_{n=-\infty}^{\infty} \frac{1}{2} (n-\mu_D)^2 \left(\mathrm{erf}\left(\frac{n+1-\mu}{\sigma\sqrt{2}} \right) - \mathrm{erf}\left(\frac{n-\mu}{\sigma\sqrt{2}} \right) \right). \end{aligned}$$

In practical calculations, one cannot take the sums up to infinity. Instead, one should choose some appropriate limits a and b, such that the error function becomes constant at the limits (this happens when the density function approaches zero). In this case, after some simplifications,

$$\mu_D \simeq \frac{1}{2} \left(b + a - 1 - \sum_{n=a}^{b} \mathrm{erf}\left(\frac{n-\mu}{\sigma\sqrt{2}} \right) \right), \tag{6.3}$$

and

$$\sigma_D^2 \simeq \frac{1}{2} \left((b-\mu_D)^2 + (a-1-\mu_D)^2 \right) + \sum_{n=a}^{b} \left(\frac{1}{2} + \mu_D - n \right) \mathrm{erf}\left(\frac{n-\mu}{\sigma\sqrt{2}} \right). \tag{6.4}$$

Now, Eqs. 6.3 & 6.4 define a mapping between (μ, σ) and (μ_D, σ_D). If one finds an inverse to this function, one should be able to use this to extract analogue values μ and σ from a discrete (i.e. digital) distribution, such as the one shown in Fig. 6.2. First, however, one should do some fact checks, to convince oneself that the obtained distributions make sense.

Asymptotically, as $\sigma \to \infty$, $\mu_D \to \mu - 1/2$ and $\sigma_D \to \sigma$, which is the expected behaviour. For the more interesting regime of small σ, Fig. 6.3 shows the contours of μ_D and σ_D in the (μ, σ) plane. For very small values of σ, $\mu_D \approx 0$ (modulo 1) and $\sigma_D \approx 0$ for a large part of the phase space. This corresponds to the situation where all analogue values occupy a small region far from the discrete levels, and thus almost all values are truncated to the same value. From the figure, it appears that for $\sigma \gtrsim 0.6$, μ_D and σ_D approach their asymptotic values, which means that the discrete output should give a reliable measurement of the noise in this case.

6.3.2 Taking Into Account the Odd-Even Effect

As has been shown in Section 5.7, the ideal model for the ADC used in Section 6.3.1 does not seem to describe reality. For a more accurate description, the odd-even effect,

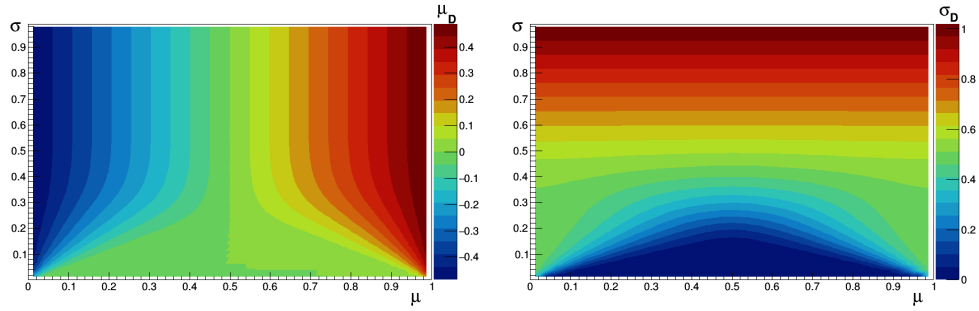

Figure 6.3 – Mean μ_D **(a)** and standard deviation σ_D **(b)** of the discrete distribution, as functions of analogue mean μ (defined on the interval $[0, 1)$) and standard deviation σ.

i.e. the imbalance between odd and even ADC levels, must be taken into account. This can be modelled by introducing a shift in each level relative to the ideal case. Using the parameter $\Delta_{(n)} \in (-1, 1)$ ADC units for the shift at ADC level n, where a positive value of $\Delta_{(n)}$ corresponds to an increase in the level relative to the ideal case (which means that some levels may have a negative value of $\Delta_{(n)}$), yields the density function

$$\begin{aligned} f_D(x) &= \sum_{n=-\infty}^{\infty} \delta(x-n) \int_{n+\Delta_{(n)}}^{n+1+\Delta_{(n+1)}} f_A(x')\mathrm{d}x' \\ &= \sum_{n=-\infty}^{\infty} \frac{1}{2}\delta(x-n) \left(\mathrm{erf}\left(\frac{n+1-\mu+\Delta_{(n+1)}}{\sigma\sqrt{2}} \right) - \mathrm{erf}\left(\frac{n-\mu+\Delta_{(n)}}{\sigma\sqrt{2}} \right) \right). \end{aligned} \quad (6.5)$$

Following similar steps as in Section 6.3.1, one obtains

$$\mu_D \simeq \frac{1}{2} \left(b + a - 1 - \sum_{n=a}^{b} \mathrm{erf}\left(\frac{n-\mu+\Delta_{(n)}}{\sigma\sqrt{2}} \right) \right), \quad (6.6)$$

and

$$\sigma_D^2 \simeq \frac{1}{2} \left((b-\mu_D)^2 + (a-1-\mu_D)^2 \right) + \sum_{n=a}^{b} \left(\frac{1}{2} + \mu_D - n \right) \mathrm{erf}\left(\frac{n-\mu+\Delta_{(n)}}{\sigma\sqrt{2}} \right). \quad (6.7)$$

While this should work for any ADC with irregular output levels, Eqs. (6.5)-(6.7) are not very useful for these measurements unless one finds an expression for $\Delta_{(n)}$. The output pattern indicates that the problem occurs in the comparators, and more specifically when each bit is set. Therefore, $\Delta_{(n)}$ should be a superposition of the shifts at each comparison in the ADC. Introducing the parameters Δ_{2^m} for the shift at bit m starting at the smallest comparison, the term Δ_{2^m} should be included if this bit is set and otherwise excluded, yielding the expression

$$\Delta_{(n)} = \sum_{m=0}^{N} \Delta_{2^m} \left(\lfloor n/2^m \rfloor \mod 2 \right),$$

where N is the total number of ADC bits (here 10). Since the output used in this chapter only occupies a few different ADC bins, and a fitting procedure will be used to obtain the best Δ values, one cannot include all bits. In particular, if one bit is set for all different bins, this will not affect the standard deviation, and there is no way to disentangle this particular Δ_{2^m} from μ, and thus it should be excluded. There is also a risk of overfitting if including too many levels. Therefore, only two levels, Δ_1 and Δ_2, are used in this analysis.

Figure 6.4 illustrates the case of a non-zero Δ_1, but with all higher Δs set to zero (an example with both Δ_1 and Δ_2 being non-zero is shown in Appendix D.1). Here, $\Delta_1 = -0.15$ is chosen, since this, as will be seen later, is a relevant choice for studying data from the SAMPA chip. As seen, there is a clear asymmetry between odd and even levels. Moreover, the asymptotic behaviour is reached at a slightly greater σ than in the $\Delta_1 = 0$ case, namely at $\sigma \approx 0.8$.

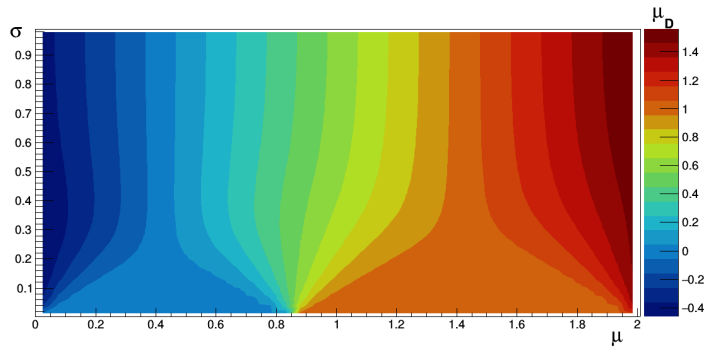

(a)

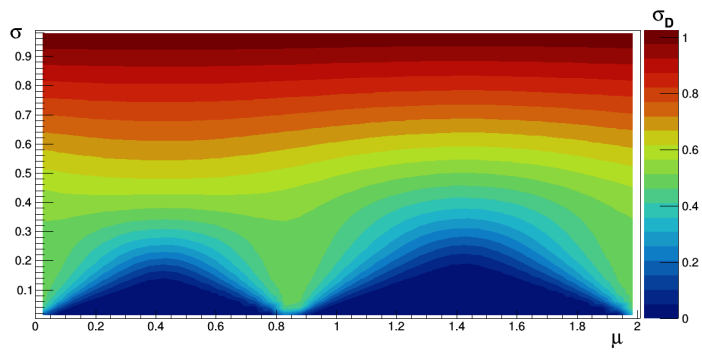

(b)

Figure 6.4 – Mean μ_D **(a)** and standard deviation σ_D **(b)** of the discrete distribution when using $\Delta = -0.15$, as functions of analogue mean μ (defined on the interval $[0, 2)$) and standard deviation σ.

6.4 Inverse

To be able to obtain analogue values (μ, σ) from the discrete distribution obtained from the digital output, one needs to invert the function defined by Eqs. 6.6 & 6.7. In Appendix A.1, it is shown that a local inverse exists whenever $\sigma > 0$, which when taking into account the continuity and asymptotic behaviour of the function on $(\mathbb{R} \times \mathbb{R}_+) \mapsto (\mathbb{R} \times \mathbb{R}_+)$, strongly implies that a global inverse exists in its domain.

Since no analytical expression exists for the inverse, this was found by minimising

$$f(\mu, \sigma) = (\mu_D(\mu, \sigma) - \tilde{\mu}_D)^2 + (\sigma_D(\mu, \sigma) - \tilde{\sigma}_D)^2 \tag{6.8}$$

for given discrete values $(\tilde{\mu}_D, \tilde{\sigma}_D)$, where μ_D and σ_D are given by Eqs. 6.6 and 6.7, respectively, using Newton's method. This function has the nice feature of a minimum of $f = 0$ exactly at the inverse. Newton's method is an iterative method where the iteration step is given as (modification of the work shown in Ref. [114, pp. 65-66])

$$\begin{pmatrix} \mu_{n+1} \\ \sigma_{n+1} \end{pmatrix} = \begin{pmatrix} \mu_n \\ \sigma_n \end{pmatrix} - \gamma_n \left(\epsilon_n \mathbf{I} + \mathbf{H} f(\mu_n, \sigma_n)\right)^{-1} \nabla f(\mu_n, \sigma_n), \tag{6.9}$$

where

$$\mathbf{H} f = \begin{pmatrix} \dfrac{\partial^2 f}{\partial \mu^2} & \dfrac{\partial^2 f}{\partial \mu \partial \sigma} \\ \dfrac{\partial^2 f}{\partial \mu \partial \sigma} & \dfrac{\partial^2 f}{\partial \sigma^2} \end{pmatrix} \tag{6.10}$$

is the Hessian matrix of f, ϵ_n is a constant to make

$$\epsilon_n \mathbf{I} + \mathbf{H} f(\mu_n, \sigma_n)$$

positively definite, and γ_n is a constant to reduce the step size, should it become

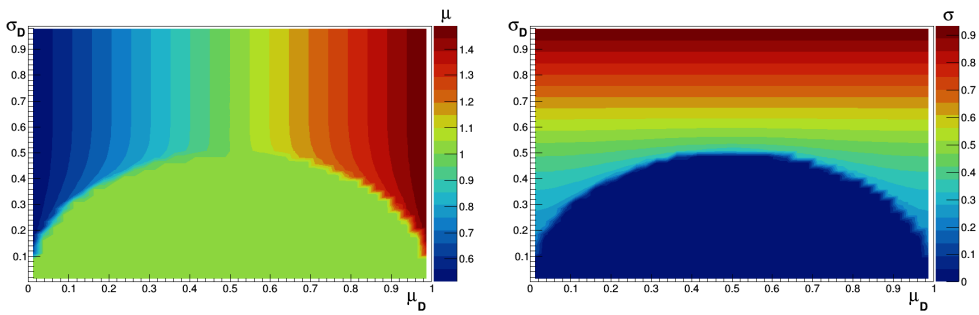

(a) **(b)**

Figure 6.5 – Analogue mean μ **(a)** and standard deviation σ **(b)**, as functions of mean μ_D (defined on $[0, 1)$) and standard deviation σ_D of the discrete distribution when using $\Delta_{(n)} = 0$, extracted from the inverse of Eqs. 6.6 & 6.7. Note that the inverse is not defined for $\sigma_D \leq \sqrt{\mu_D(1 - \mu_D)}$. Due to plotting limitations, $(\mu, \sigma) = (1, 0)$ are used for this region in the plots, since these values are obtained when approaching the limit.

too large. Here, it was concluded that $\mathbf{H}f(\mu_n, \sigma_n)$ is positively definite if its smallest eigenvalue $\lambda_{\min} \geq 0.5$. If not, ϵ_n was chosen as $0.5 - \lambda_{\min}$. Initially, $\gamma_n = 1$ was used. If $\sigma_{n+1} < 0$, $\sigma_D(\mu_{n+1}, \sigma_{n+1}) < \sigma_D(\mu_n, \sigma_n)/10$, or $f(\mu_{n+1}, \sigma_{n+1}) > f(\mu_n, \sigma_n)$, it was concluded that the step was too large and γ_n was decreased until neither of these criteria were fulfilled. The derivatives were computed numerically.

As a starting point, $(\mu_0, \sigma_0) = (\tilde{\mu}_D - 0.5, \tilde{\sigma}_D)$ was used. The iteration proceeded until $f(\mu_n, \sigma_n) < 10^{-13}$, where convergence was assumed. If the function did not converge, a different starting point was chosen by slightly increasing the value of σ_0 and moving μ_0 closer to the centre of the interval.

Contour plots over the resulting inverse functions for $\Delta_{(n)} = 0$ are shown in Fig. 6.5 and analogous ones for $\Delta_1 = -0.15$ in Fig. 6.6. Corresponding plots for $\Delta_1 = -0.15$ and $\Delta_2 = 0.05$ are shown in Appendix D.1. One should note that no inverse exists in the forbidden region of

$$\sigma_D \leq \sqrt{\mu_D(1 - \mu_D)},$$

where μ_D is defined modulo 1, since this is outside the domain of the inverse (cf. Eq. 6.1). This causes some interpolation issues in the plots. In the case of a non-zero Δ_1, it is

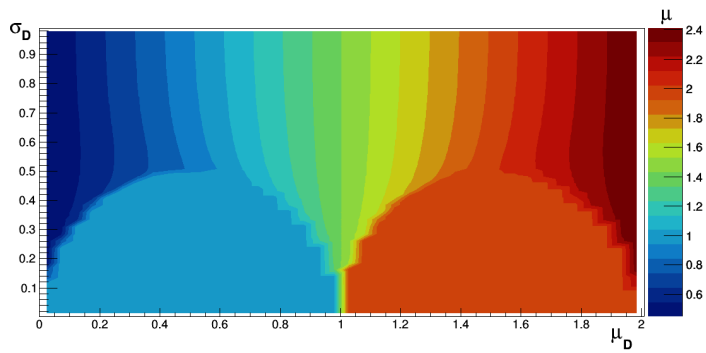

(a)

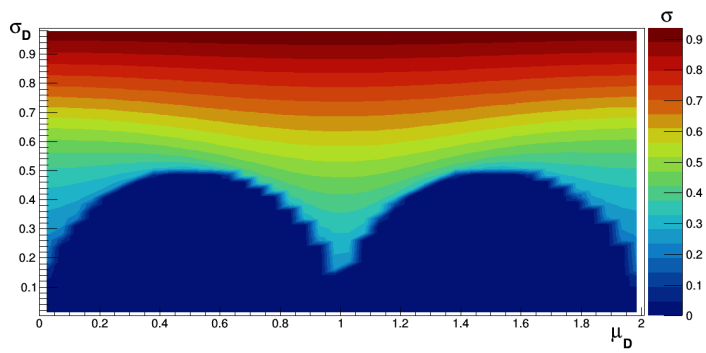

(b)

Figure 6.6 – Analogous to Fig. 6.5, these plots show μ and σ as functions of μ_D (defined on the interval $[0, 2)$) and σ_D, for $\Delta_1 = -0.15$.

clear that for the extracted value of σ, it is quite significant whether the discrete mean is close to an odd bin or an even one.

6.5 Application to the Data Set

6.5.1 Hypotheses

With an inverse available, it is possible to extract estimates of analogue μ and σ to the SAMPA data set used to generate Figs. 6.1 and 6.2. For this, five different hypotheses were tested:

1. The discrete ADC scale is the only effect, i.e. there is no odd-even effect and $\Delta_{(n)} = 0$.

2. The magnitude of the odd-even effect is the same for all channels, but no higher-order bit errors are present, i.e. Δ_1 is a constant non-zero value and $\Delta_2 = 0$.

3. The magnitude of the odd-even effect is different for different channels, but there are still no higher-order bit errors, i.e. Δ_1 is channel-dependent, but $\Delta_2 = 0$.

4. Higher-order bit errors are present, but the magnitude of all of them is the same for all channels, i.e. Δ_1 and Δ_2 have constant non-zero values.

5. Higher-order bit errors are present, and different for different channels, i.e. Δ_1 and Δ_2 are channel-dependent.

For Hypotheses 2-5, optimal values of Δ_1 and Δ_2 were found by minimising the χ^2 function w.r.t. (Δ_1, Δ_2), here defined as

$$\chi^2 = \sum_j N_j \sum_i \frac{\left(N_i^{obs}/N_j - p_i^{model}\left(\Delta_1^{(j)}, \Delta_2^{(j)}\right)\right)^2}{p_i^{model}\left(\Delta_1^{(j)}, \Delta_2^{(j)}\right)}, \tag{6.11}$$

where j denotes the channel number, N_j is the total number of entries in this channel, i denotes the ADC bin, N_i^{obs} is the observed number of counts in bin i, and p_i^{model} is the probability to get an entry in this bin according to the density function defined by Eq. 6.5. For Hypothesis 2 and 4, the same Δs were used throughout all channels, whereas for Hypothesis 3 and 5, a new minimisation was carried out for each channel.

The χ^2 function was minimised using a combination of conjugate gradients and Brent's method on $(|\Delta_1|, |\Delta_2|) \leq (0.5, 0.5)$. These methods were chosen since the χ^2 function did not appear to be quadratic, which is required for Newton's method. The conjugate gradient method is a method for choosing directions in multidimensional minimisation, which is reliable and generally efficient, since it uses that most well-behaved functions are nearly quadratic close to the minimum [114, pp. 78-79]. To find the actual minimum, this must be combined with a line search algorithm. Here, Brent's method was chosen since this fits a parabola to the function whenever possible and otherwise uses the slower but more reliable Golden Section line search, making it both reliable and efficient [115, pp. 73-75].

It should be noted that the data from channel 30, located at the corner of the socket, appeared to not follow any distribution close to a Gaussian one (or its discrete analogy), with a significant amount of data points far from the mean value. For this reason, this channel was excluded from the global minimisation of χ^2 used for Hypothesis 2. A similar issue seemed to appear also for the 33 pF measurement in channel 2. Therefore, this data point was excluded from the minimisation as well.

6.5.2 Error estimate

To test the statistical uncertainty of the obtained μ and σ for each hypothesis, a simulation was carried out. This was done by drawing N Gaussian samples with the obtained parameters and discretising each sample through truncation (including the odd-even effect; this is one correct way to draw from the discrete distribution), where N is the number of entries in the original measurement. The mean and standard deviation of the resulting distribution were then inverted. The entire process was repeated 100 times and the spreads in the resulting μ and σ were used to estimate the errors.

It should be noted that this only gives an estimate of the statistical error, so e.g. uncertainties in the obtained Δ values (see below) are not included in this error estimate.

6.5.3 Validation

The computed χ^2 value for a given fit gives a measure of the quality of fit. Therefore, the χ^2/d.o.f. values obtained from each hypothesis (excluding channel 30, where it is known that the fit is really poor) were compared to test the significance of the odd-even effect, and to what extent the models used here represent reality. Moreover, it was tested whether there is any significant difference in extracted analogue noise levels between the hypotheses. If not, one should be able to use a simpler model and skip a tedious minimisation procedure to obtain the underlying σ values.

For Hypothesis 3 and 5, the 98% confidence intervals for Δ were estimated by finding the limits where the p-values cross $P = 0.02$, for each channel where, it is possible (i.e. $P > 0.02$ for the best fit). These are not confidence intervals in the traditional sense, but rather limits for how far away one can depart from the optimal value and still have a reasonable chance of obtaining the measured data.

6.6 Results

Figs. 6.7 and 6.8 show the extracted analogue values of σ plotted against μ_0, as well as the spread in σ plotted against detector capacitance (these plots are analogous to Figs. 6.1 and 6.2), for Hypothesis 1 and 5, respectively. The results for the other hypotheses are very similar to those from Hypothesis 5, and are thus only shown in Appendix D.1. In Fig. 6.7 (Hypothesis 1), there appears to remain a weak correlation between μ and σ, which is not present for any of the other hypotheses. The spread in σ between channels is also greater for this hypothesis. Therefore, it is probably best to include the odd-even effect in one way or another. There are a few significant outliers present in all of the plots. These are almost all from channel 30, which even though

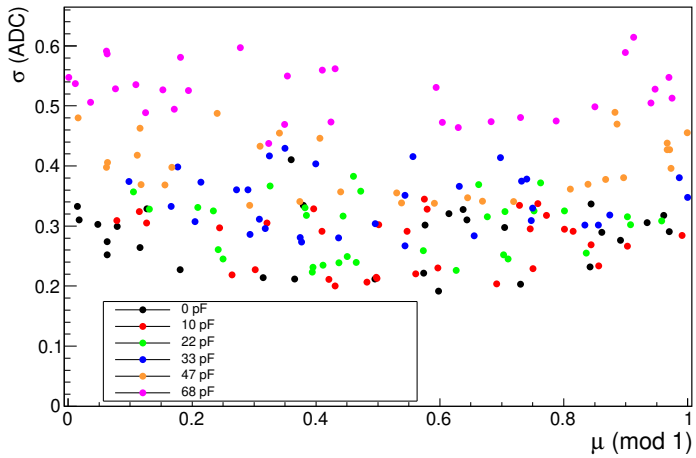

(a)

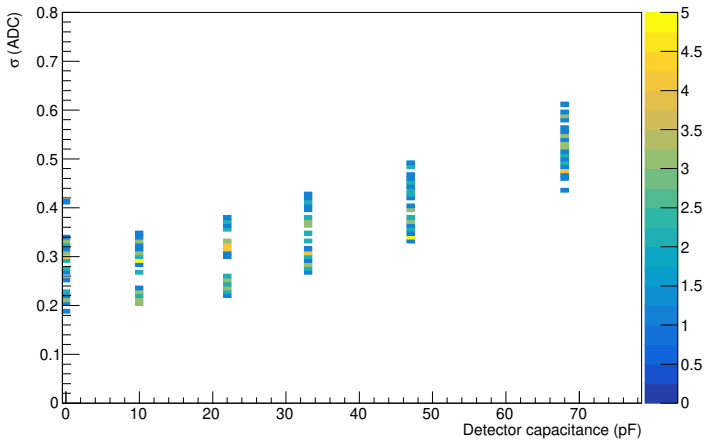

(b)

Figure 6.7 – **(a)** Extracted analogue noise σ versus truncated mean μ_0 for different detector capacitances, computed for $\Delta_{(n)} = 0$ (Hypothesis 1), i.e. the assumption of no odd-even effect. For a given capacitance, each point corresponds to a single channel. Statistical error bars are included, but not visible since they are smaller than the points themselves. **(b)** Spread in σ between channels, as a function of detector capacitance. The colour indicates the number of channels in a narrow σ interval.

the most outlying data points have been removed, still has a greater standard deviation than most other channels at low capacitances.

The obtained σ values are summarised in Fig. 6.9 for all hypotheses, including the raw data itself (marked as "Hypothesis 0"), as functions of detector capacitance. These

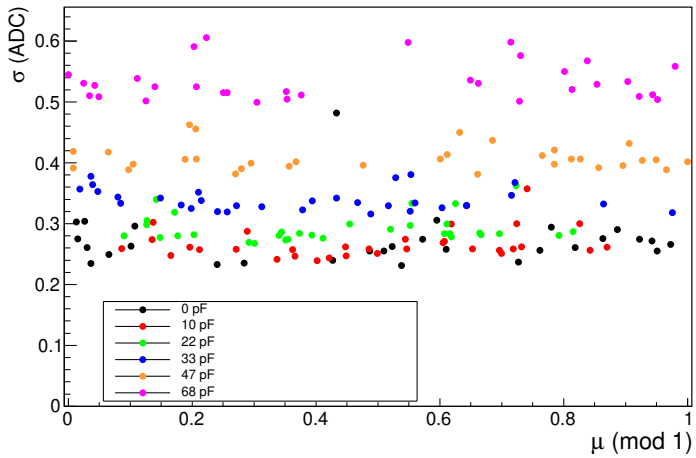

(a)

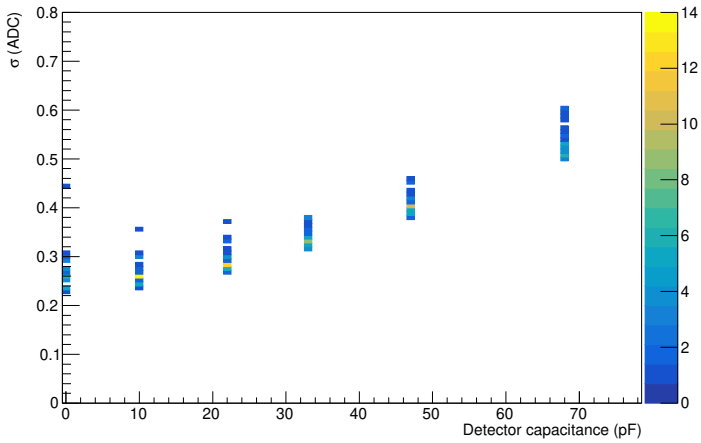

(b)

Figure 6.8 – Analogous plots to Fig. 6.7, computed using optimised values of Δ_1 and Δ_2 in each channel (Hypothesis 5).

values are in all cases computed while excluding channel 30. As seen, both the values of σ and their spread between channels drop compared to the discrete data for all methods described here, and in particular at lower capacitances. There is not any significant difference in average value of sigma between any of the hypotheses described in 6.5.1, but the spread between channels is significantly smaller for Hypothesises 2-5. One should also note that σ is rather constant up to 22 pF also in the inverted case, varying around 0.27 ADC units, or ~ 900 electrons, which is close to the specifications (although these are defined for a slightly higher input capacitance).

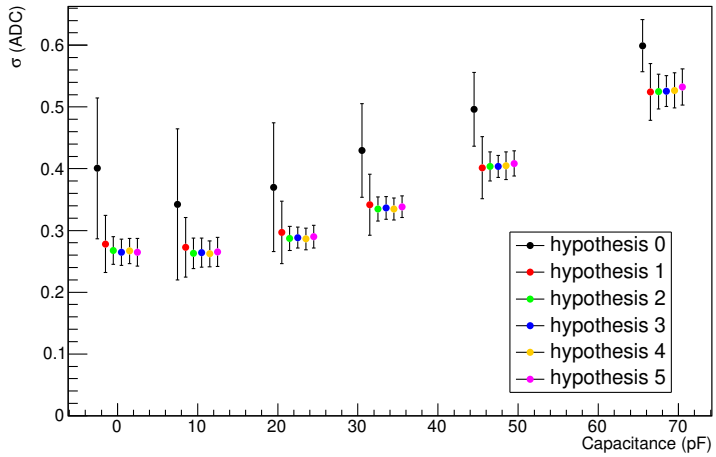

Figure 6.9 – Average values of σ obtained for each of the hypotheses described in Section 6.5.1, compared with analogous results directly from the discrete distribution (denoted "Hypothesis 0"). The error bars indicate spread between channels. The points are separated horizontally to resolve the points.

The magnitude of the odd-even effect obtained for the best fit in Hypothesis 2 is $\Delta_1 = -0.1378$, which is close to the value used to generate Figs. 6.4 and 6.6. For Hypothesis 4, the best values are $\Delta_1 = -0.1484$ and $\Delta_2 = 0.0639$, so Δ_1 is clearly dominant.

For Hypothesis 3 and 5, the Δ values for the best fits, along with 98% confidence limits, are summarised in Fig. 6.10a and 6.11a, respectively. As seen, there are quite large variations in Δ. For Hypothesis 3, the minimum is at $\Delta_1 \approx -0.21$ and the maximum is at $\Delta_1 \approx -0.08$. For Hypothesis 5, Δ_1 varies between -0.25 and -0.08, and Δ_2 between -0.03 and 0.26. For Δ_2, there are two very distinct regimes, with a majority of channels having a value close to zero, or at most 0.1, and a few channels being > 0.2. For many of the former, the fit is quite good using a single parameter, whereas for the latter, both parameters are required to explain the data. Upon closer inspection, the channels with a large Δ_2 all have a μ just below 80 ADC units, i.e. they are close to a transition between 0 and 1 for the four lowest bits. Therefore, the results are likely due to influences from Δ_3 and Δ_4, indicating that the true value of Δ_2 is quite weak. Nevertheless, when the spread is so low as in these measurements, Δ_2 may be used to mimic a large fraction of the combined effect of these bit errors. It should also be noted that the uncertainty of Δ_2 is quite large for some channels, but generally smaller for Δ_1.

Resulting χ^2 values for each hypothesis are summarised in Table 6.1, along with corresponding numbers of degrees of freedom. Individual χ^2 values for each channel for Hypothesis 3 and 5 are summarised in Fig. 6.10b and 6.11b, respectively. Of all these models, only Hypothesis 5 has a χ^2/d.o.f. close to 1, indicating that one needs to include several Δ parameters to accurately describe the data. For 14 of the channels,

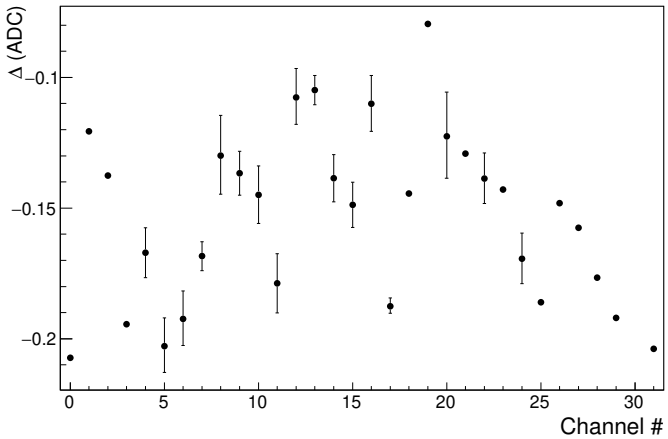

(a)

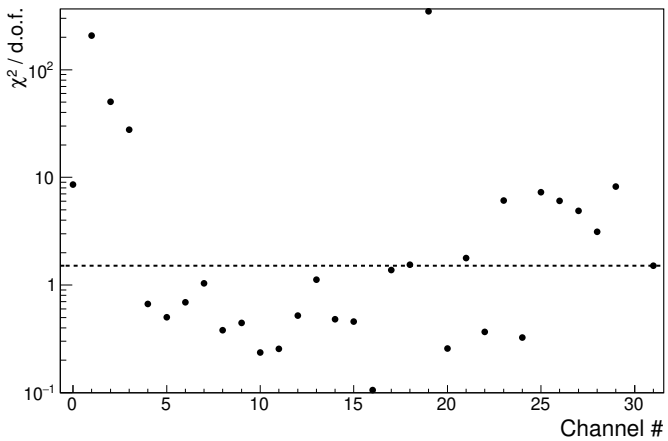

(b)

Figure 6.10 – **(a)** Best fits of Δ_1 for each channel (except channel 30) obtained for Hypothesis 3. For channels where $P > 0.02$ for the best fit, 98% confidence limits are given. For channel 2, the 33 pF measurement has been excluded. **(b)** Corresponding χ^2/d.o.f. values for the best fits. The dashed line indicates where $P = 0.02$ for 40 degrees of freedom, which is a typical number of degrees of freedom for a single channel.

however, χ^2/d.o.f. < 1.0 by only including Δ_1, so this may not be necessary in all channels. Including Δ_2 in the fit, this increases to 18, with a p-value $P > 0.02$ for 21 channels.

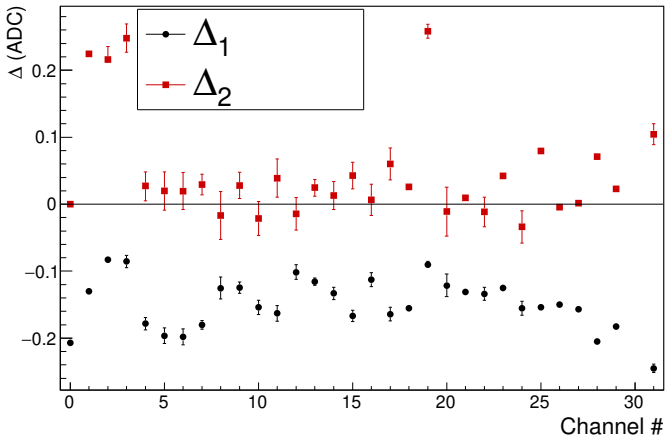

(a)

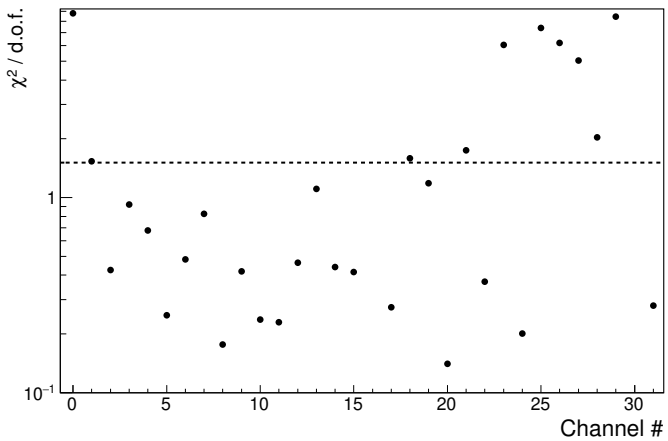

(b)

Figure 6.11 – **(a)** Best fits of Δ_1 and Δ_2 for each channel (except channel 30) obtained for Hypothesis 5. For channels where $P > 0.02$ for the best fit, 98% confidence limits are given. For channel 2, the 33 pF measurement has been excluded. **(b)** Corresponding χ^2/d.o.f. values for the best fits. The dashed line indicates where $P = 0.02$ for 40 degrees of freedom, which is slightly above the typical number of degrees of freedom for a single channel.

6.7 Application to the Automatic Testing

6.7.1 Motivation

In the previous sections, this method has only been applied to a single chip in the development phase. To put this in a wider context, and increasing the data sample,

Table 6.1 – χ^2 values and degrees of freedom (d.o.f.) computed for the results from each hypothesis described in Section 6.5.1. Data from channel 30, as well as the 33 pF measurement in channel 2, have been excluded from the computations. Resulting values of Δ are also indicated.

Hypothesis	χ^2	d.o.f.	χ^2/d.o.f.	Δ_1	Δ_2
1	$1.04 \cdot 10^6$	1231	844	0	0
2	$6.06 \cdot 10^4$	1207	50	-0.1378	0
3	$2.89 \cdot 10^4$	1178	25	-0.08 – -0.21	0
4	$4.03 \cdot 10^4$	1195	34	-0.1484	0.0639
5	$2.07 \cdot 10^3$	1151	1.8	-0.25 – -0.08	-0.03 – 0.26

this should also be tested on data from the robot testing. While the noise measurements at 4 mV/fC were never used for the validation of the MCH chips, they were still carried out, and can therefore be used for this study. This also makes it possible to test whether there is any strong correlation between the corrected results at 4 mV/fC and the raw results at 20 mV/fC, which were used during the testing campaign.

6.7.2 Measurement Details

For this study, a total of 10 chips tested for the MCH are included. Of these, five were analysed at each test station to test whether the choice of station has any significant effect on these results. For each chip, each of the hypotheses defined in Section 6.5.1 were tested. Finally, the same analysis was repeated for the 20 mV/fC results of the same chips, which was done both to ensure that the discretisation effect is not significant here, and to compare the results between different gain settings. This was done separately from the previous measurements, since the increased gain caused the pedestal levels to drift, making it difficult to combine measurements from different gain settings (if desiring to do so, one would need to include more than two Δ levels).

6.7.3 Results

Figure 6.12 shows the measured (uncorrected) σ as a function of measured baseline μ for the 10 chips included in this measurement, along with the spread between channels for each individual chip. As seen, the discretisation effect is present also here, with data similar to the lower-capacitance ($\lesssim 30\,\text{pF}$) measurements of the prototype chip, so similar corrections as done previously are clearly necessary.

Results from the corrections defined by Hypothesis 1, 4, and 5 in Section 6.5.1 are shown in Figs. 6.13, 6.14, and 6.15, respectively. The results from Hypothesis 2 and 3, which are similar to Hypothesis 4 and 5, respectively (but have larger χ^2/d.o.f.), are shown in Appendix D.1. For Hypothesis 1 and 4, similar observations can be made as for the prototype testing, namely that some correlation between μ and σ remains if setting $\Delta_{(n)} = 0$, but this is removed by using non-zero Δ_1 and Δ_2. For Hypothesis 5, however, the measured noise level seems to be largely reduced from Hypothesis 4, with many channels having an apparent σ around 0.2. This is almost exclusively observed in channels with $\mu \gtrsim 0.2$, raising questions about measurement

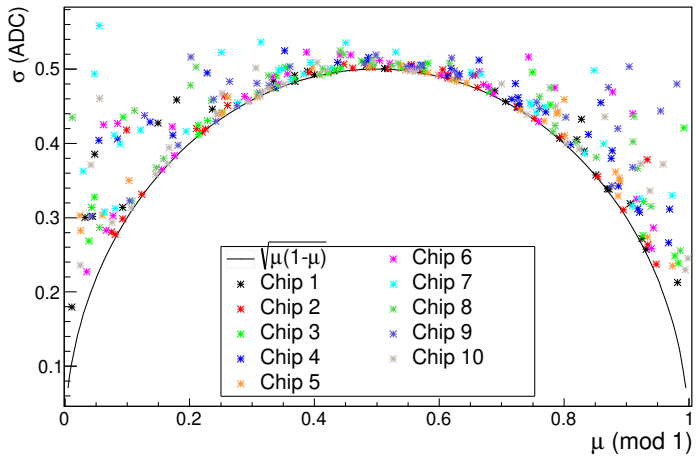

(a)

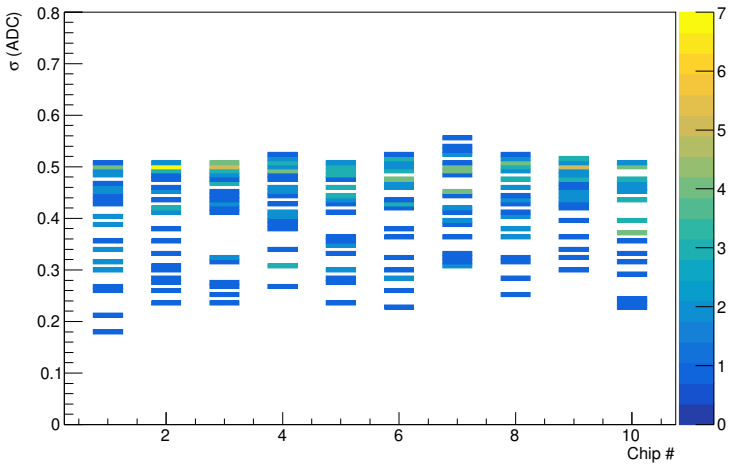

(b)

Figure 6.12 – **(a)** Uncorrected noise σ as a function of truncated mean μ_0, measured for 10 different chips from the automatic testing campaign. Each data point corresponds to a single channel. The solid line marks the limiting case defined by Eq. (6.1). **(b)** Spread in σ between channels, divided between the different chips. The colour indicates the number of channels in a narrow σ interval.

bias. Upon closer inspection, most of the low-σ channels have either a negative or large positive value of Δ_1, which also generally has a much larger uncertainty than what is observed in Figs. 6.10a and 6.11a. This is due to a much smaller data set ($\sim$ 37 000 samples/channel compared with $\sim$ 2 000 000 samples/channel distributed

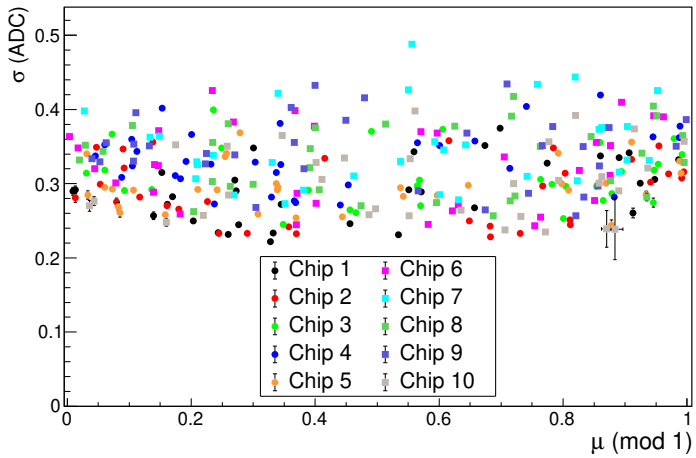

(a)

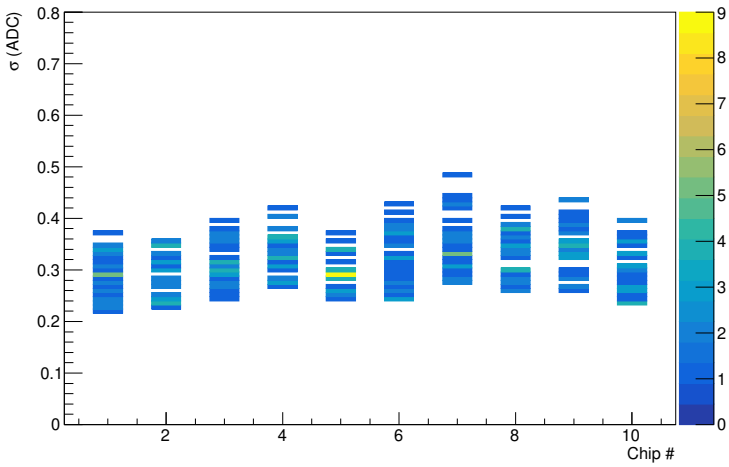

(b)

Figure 6.13 – **(a)** Extracted analogue noise σ versus truncated mean μ_0 for the same chips as were used to obtain Fig. 6.12, computed for $\Delta_{(n)} = 0$ (Hypothesis 1). Each data point corresponds to a single channel in a chip. The marker type indicates the station where the chip was tested; circular and square markers are used for test station 1 and 2, respectively. The error bars are statistical only (large error bars may result from failed convergence during the error simulation, since this reverted to a generic estimate of $(\mu, \sigma) = (0.5, 0)$). **(b)** Spread in σ between channels, divided between chips. The colour indicates the number of channels in a narrow σ interval.

over the six capacitance measurements of the prototype) along with very few degrees of freedom per channel (resulting from only carrying out a single measurement per chip).

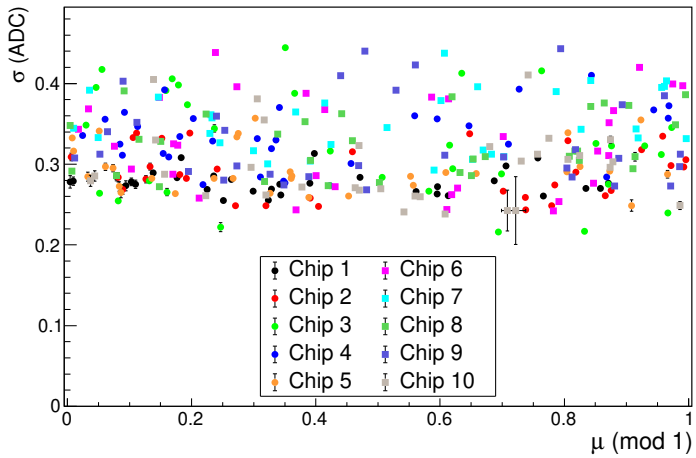

(a)

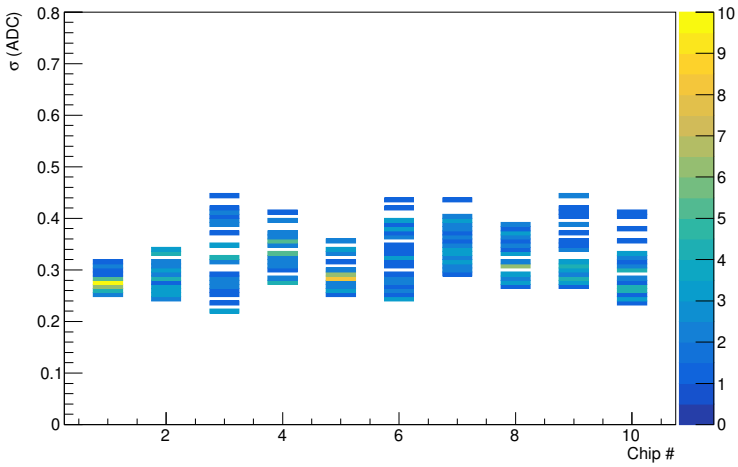

(b)

Figure 6.14 – Plots analogous to Fig. 6.13, computed for non-zero Δ_1 and Δ_2, optimised for each chip, but assumed to be constant within the chip (Hypothesis 4).

Therefore, trying to fit a unique Δ for each channel may result in overfitting, causing the minimisation of χ^2 to end up at a solution with a lower σ but higher Δs than the true parameters of the chip. Therefore, Hypothesis 3 and 5 cannot be applied reliably to this data set and one should rather use Hypothesis 4 instead. The results from the prototype test, and Fig. 6.9 in particular, indicate that Hypothesis 4 gives very similar results to Hypothesis 5, meaning that the results in Fig. 6.14 should accurately describe the noise levels in the chips from the automatic testing.

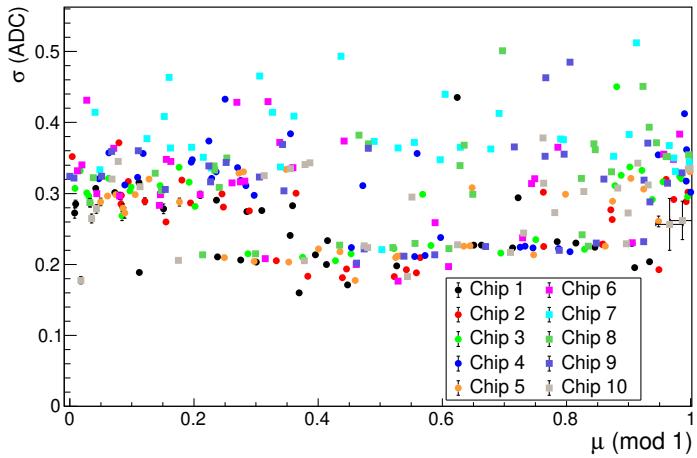

Figure 6.15 – Plot analogous to Fig. 6.13a, computed for non-zero Δ_1 and Δ_2, obtained by minimising the χ^2/d.o.f. for each individual channel in every chip (Hypothesis 5).

Uncorrected noise as a function of baseline for a gain of 20 mV/fC, i.e. the data used for assessing the noise level of these chips, is shown in Fig. 6.16. As seen, the measured noise is far above the measurement limit indicated by Eq. (6.1). There appears to be two outliers, with a single channel being exposed to much more noise than the others. These are both in chips tested at test station 1, and both in channel 4, but it was not further analysed whether this is a coincidence or a systematic effect.

Figure 6.17 shows the converted values based on Hypothesis 5. Due to the larger number of degrees of freedom in these measurements, the problem with overfitting is not present here. The results from the other hypothesis are very similar and have thus been moved to Appendix D.1. As expected, the conversion does not significantly change the results, showing that one does not need to account for the discretisation effects when the noise level is this high. This is even clearer when studying Fig. 6.18, which shows a comparison between the obtained σ from the uncorrected data and corrected results using two of the hypotheses, for each chip and gain setting. The selected hypotheses are Hypothesis 1, since it has no odd-even effect, and Hypothesis 4 for 4 mV/fC and Hypothesis 5 for 20 mV/fC, since these are considered the most accurate ones given the overfitting problem. For 20 mV/fC, the correction does not change the results very much, but taking the odd-even effect into account seems to bring the results slightly closer to the raw data. For 4 mV/fC on the other hand, corrections are necessary just as for the prototype testing. For both hypotheses, the extracted analogue noise is considerably lower than the measured digital noise and the channel spread is reduced, but including the odd-even effect reduces the spread even more.

Finally, Fig. 6.19 shows the channel-by-channel correlation between measured noise at 20 mV/fC and the best correction at 4 mV/fC. Quite remarkably, there is only a weak correlation between the two measurements. In particular, in most chips, there seems to be two distinct noise levels at the 20 mV/fC measurement, which have no

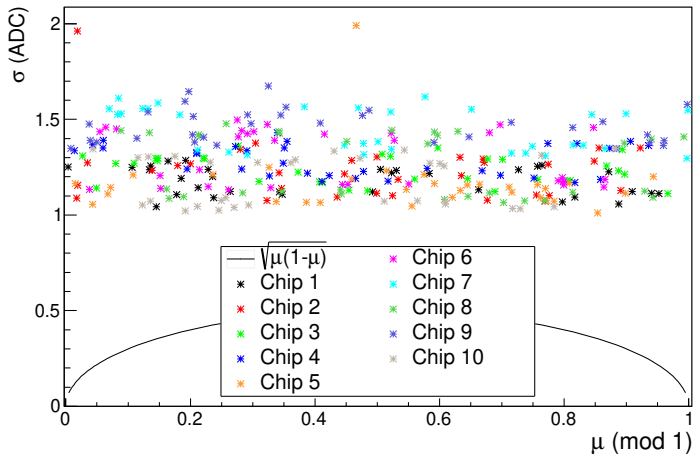

(a)

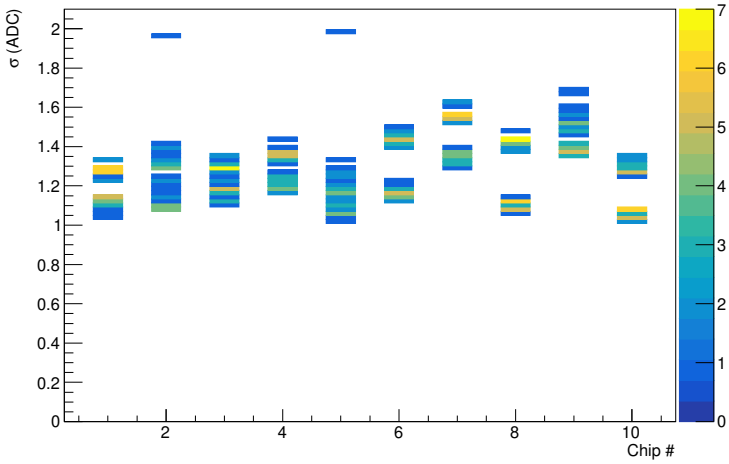

(b)

Figure 6.16 – Analogous plots to Fig. 6.12, measured using a gain of 20 mV/fC.

correspondence in the measurement at 4 mV/fC. Moreover, the two outliers observed at 20 mV/fC do not show up at the lower gain setting. Apart from these feeatures, there seems to be a weak correlation, but with a much larger spread at 4 mV/fC than at 20 mV/fC. Measured in electrons, the noise is almost exclusively higher at 4 mV/fC, with some channels having an ENC as high as 1500 electrons, compared to only a few channels above 1150 electrons at 20 mV/fC.

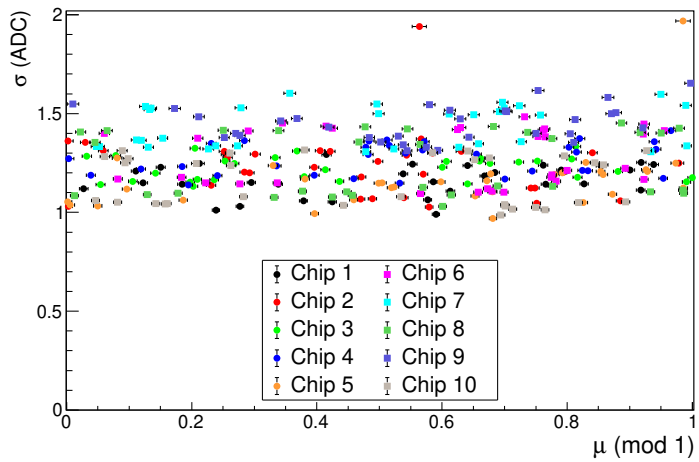

(a)

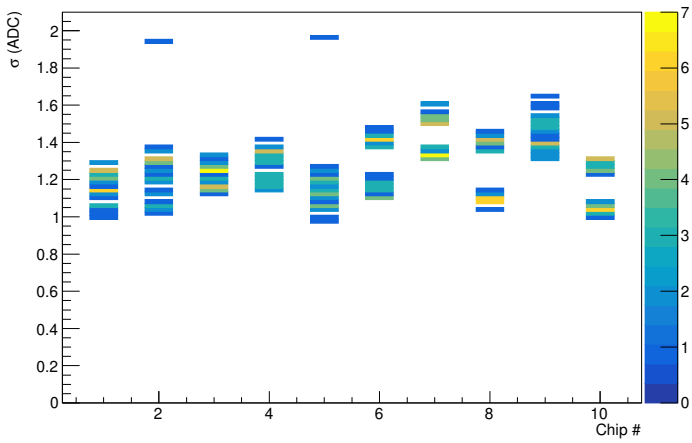

(b)

Figure 6.17 – Plots analogous to Fig. 6.13 for the 20 mV/fC data set, computed for non-zero Δ_1 and Δ_2, obtained by minimising the χ^2/d.o.f. for each individual channel in every chip (Hypothesis 5).

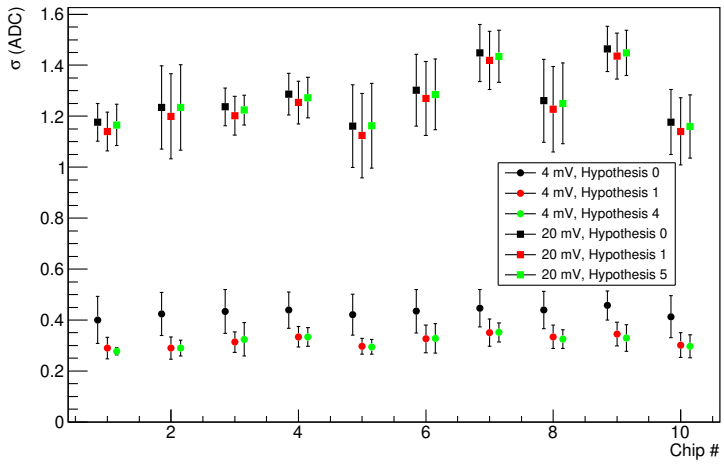

Figure 6.18 – Average values of σ for the 10 chips from the automatic testing, obtained for a selection of the hypotheses described in Section 6.5.1, compared with analogous results directly from the discrete distribution (denoted "Hypothesis 0"). The error bars indicate spread between channels. The points are slightly separated horizontally to resolve them.

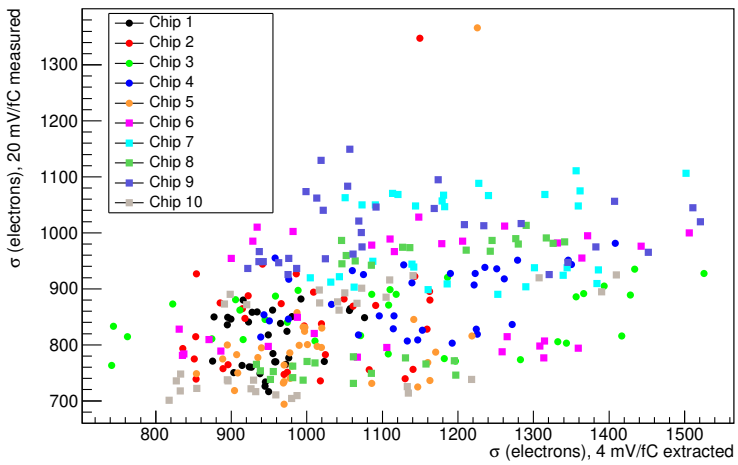

Figure 6.19 – Correlation between uncorrected σ at 20 mV/fC and extracted σ at 4 mV/fC measured in electrons, using Hypothesis 4. Each marker corresponds to a single channel. Circular markers are used for chips tested at station 1 and square markers for chips tested at station 2.

6.8 Discussion

This work shows that when trying to measure noise in a sampled data set, this will suffer from a discretisation effect if the noise is too low, but also that we can do something about it. The odd-even effect complicates matters considerably, especially since this does not just affect the last bit, but every comparison in the ADC seems to be affected to some extent (cf. Fig. 5.15). Therefore, one should ideally include the seven lowest Δ_{2^k} to cover the full range of the pedestal levels in the chip, but due to the low spread in each channel, this is not feasible. This may be possible if assuming that the bit errors are the same for all channels within a chip, but since other measurements have indicated that this is not really the case, and there is a risk of overfitting, there was no point in pursuing this. If the channel-to-channel spread is low enough (which may be the case), it could perhaps be worthwhile to test on the 20 mV/fC data set for a more robust test of the accuracy of the model description in Section 6.3.2 (using only two levels gives a quite high χ^2/d.o.f. for this data set), but this was assessed to be beyond the scope of this analysis.

When not being able to include the full parametrisation, fitting Δ_1 and Δ_2 independently for each channel (Hypothesis 5) seems to give a close-to-accurate description in most channels. Fitting each channel independently seems to be enough to account for most higher-order bit effects, especially given the large apparent Δ_2 when the baseline is close to 80 ADC units. For a couple of channels the fit is quite poor, however. These are in particular those where μ is between 2 and 3 modulo 4, indicating that the spread in ADC should extend into the next interval of length 4 in both directions. Therefore, including Δ_4 would probably result in better results here, but given the low occupancy in many ADC channels, this may result in large uncertainties on some of the parameters, which may have a similar impact on the extracted μ and σ as for the robotic testing.

The results from Fig. 6.9 advice against trying to include more parameters. By just assuming the ideal model, the correction puts the average of σ close to what is obtained by the other hypotheses (and most likely the true value), but the spread is still too high. By just using the simplest model for the odd-even effect (Hypothesis 2), the resulting σ gets very close to the more advanced hypotheses and the correlation between μ and σ is virtually gone, indicating that this model is probably good enough. Correcting for higher-order effects by including more Δs reduces the χ^2/d.o.f., so this clearly gives a more accurate description, but it does not change the end results.

For the robotic testing, one can largely draw similar conclusions as for the prototype test with different external capacitances. The noise level is similar to an external capacitance of $\sim 20\,\mathrm{pF}$, and there does not seem to be any significant difference between the two test stations, at least for this data set. The biggest difference is the overfitting issue occurring when testing Hypothesis 3 and 5. What happens is that due to the relatively low statistics, statistical fluctuations cause the distribution to appear to be similar to one with larger Δ_1 and/or Δ_2, than what the actual parameters of the chip are. This shows that it is important to either use a large data set (2000 000 events seems to be enough, but not 30 000), or use more measurements, preferably covering more ADC levels. The latter is ensured when measuring over several different input capacitances. A lucky consequence of this is that this also seems to shift the baseline level somewhat, increasing the ADC range. The best way to remedy this in the robotic testing might

be to try to combine the measurements between 4 mV/fC and 20 mV/fC, but since this has a quite large impact on the baseline level (and uncertain impact on $\Delta_{(n)}$), it may require a bit of work and likely more than two Δ_{2^k} parameters.

If doing so, in order to ensure that there is no overfitting, one should calculate the systematic errors of extracted μ and σ from the uncertainties in the different Δ parameters (instead of just the statistical ones), which also could be done in the current analysis, but was deemed to be too tedious. Increasing the number of fit parameters would further increase the effort, which is one of the reasons why this was not attempted.

A simpler solution was to stick to Hypothesis 4, which according to the results in Section 6.6 should be good enough. A possible issue with this, which was not studied further, is that if non-Gaussian noise is present in any of the channels (such as channel 30 in the SAMPA prototype), this may strongly affect the obtained χ^2 and possibly have a large impact on the minimisation. This was not studied further, and given the large spread in σ for some chips in Fig. 6.14, such noise could definitely be present, so studying this could mean a possible improvement. In any case, the difference between Hypothesis 1 and 4 (cf. Fig. 6.18) is generally not larger than for the prototype (cf. Fig. 6.9), so this is likely only a minor effect.

For the measurements at 20 mV/fC, the results show that no correction is necessary, meaning that this can safely be used for the testing campaign. A bigger question is if the results obtained from such a measurement really can be applied to the gain setting where the chips will be used, namely 4 mV/fC. Figure 6.19 indicates that this may not really be the case, given the weak correlation between the measurements. In particular, outliers showing up in the 20 mV/fC data set may not have a correspondence in the low-gain setting, which may lead to functional chips being discarded unnecessarily. This indicates however that the outliers likely are due to conditions during the testing – e.g. non-Gaussian noise from the robot motor which was identified as a significant noise sorce – and since all chips which failed the test were retested, such issues should go away during the second iteration. The other issue is that the ENC range is larger at 4 mV/fC than at 20 mV/fC, making it possible that a chip well passing the criteria at the higher gain setting, may have a high ENC in some channels at the lower one. The test criteria are rather set such that most functional chips pass the criteria than to strictly follow the specifications, so this may result in some somewhat noisy chips passing the tests. This should however not be a real issue, unless a chip with a very high noise level passes the criteria. To test this, one should check if a chip which is really noisy at 4 mV/fC (indicating some bad channel) also will be so at 20 mV/fC. Unfortunately, no such chip was included in the data set, so this remains an open question.

Now to the most important result of this work: the measured noise levels at 4 mV/fC – both low capacitances of the prototype, and in general for the automatic testing – are largely due to a discretisation effect. More realistic noise level for small capacitances is around 0.27 ADC units for the measurements with no external capacitance, or 900 electrons, and ~ 0.3 ADC units (1000 electrons) for the robot testing, instead of the previously measured values of $0.35 - 0.4$ ADC units, or $1200 - 1400$ electrons.

Finally, in this study, the correction for the noise measurement has been applied to a quite limited use case. In principle, this should be applicable to any measurement of white noise where a digitisation step is involved and the noise level is below the ADC resolution. I can think of tests of shielding etc., where this approach may be

used instead of investing in expensive equipment to improve the resolution. The major drawback is that it is assumed that the noise is Gaussian, which is not always the case, and depending on the frequency interval may, installations of low- or high-pass filters may be required to achieve this. The odd-even effect is likely a general feature of ADCs, since electronic equipment can never be perfect, although the effect probably is less pronounced than in these chips, usually. The work here shows that even this effect is something that can be corrected for.

6.9 Conclusions

Although it is not possible to apply the full description of the model developed in this chapter to the present data sets, it is a promising way to extract the analogue noise from data which suffer from a discretisation effect. Moreover, while the odd-even effect does not need to be included for an accurate measurement of the mean noise level, given the presence of such an effect in these SAMPA chips, it is important for increasing the precision. A simple model turns out to be enough, however.

An estimate of the noise for a chip with an applied external capacitance of less than 22 pF is about 900 electrons, which satisfies the specifications reasonably well. This is a significant improvement from the measured values when not accounting for the discretisation effect. Similar conclusions can be drawn for chips from the automatic testing campaign, although the noise from the setup is slightly higher here.

The corrected results from the measurements at 4 mV/fC (which will be used in the detector) are only mildly correlated with the data at 20 mV/fC used for the testing. Given the data set tested here, this should not cause any problem, but a larger data set is necessary for a more conclusive statement.

The method demonstrated here should be possible to extend to more general applications of noise measurements in e.g. tests of electric shielding where the equipment resolution for the measurement is limited, provided that the noise is Gaussian.

Part III

Main Analyses

Chapter 7

Measurement of Flow in Xe–Xe Collisions

My contribution to this analysis was quite limited, so I will only describe it briefly. There is quite a large overlap with my conference proceedings from Hot Quarks 2018 [116], which in turn is a summary of Ref. [117]. I will not describe the models tested here in much detail, but the interested reader may check out the references.

7.1 Introduction

In 2017, a short Xe–Xe run at $\sqrt{s_{\mathrm{NN}}} = 5.44\,\mathrm{TeV}$ was carried out, which gave an excellent opportunity for testing models that predict a dependence with system size. The isotope used was ^{129}Xe, which is significantly lighter than ^{208}Pb, but still in the heavy-ion regime. In this study, anisotropic flow, v_{n} (cf. Section 3.7.2), was measured for $2 \leq \mathrm{n} \leq 4$ in Xe–Xe and compared with Pb–Pb measurements at $\sqrt{s_{\mathrm{NN}}} = 5.02\,\mathrm{TeV}$, in order to test initial-state models. In particular, it is assumed that the scaling with eccentricity described by Eq. (3.8) holds. Since the scaling depends on initial state (and to a lesser extent by η/s of the medium), any model preserving this scaling is a viable candidate for the initial state in heavy-ion collisions.

7.2 Analysis Methods

In this study, inclusive charged tracks reconstructed by the ITS and TPC were used (cf. Sections 4.2 and 4.4). Only tracks in the transverse-momentum region $0.2 < p_{\mathrm{T}} < 10$ GeV/c were included in the analysis (for the results presented here, p_{T} was limited to $0.2 < p_{\mathrm{T}} < 3\,\mathrm{GeV}/c$). The events were divided into multiplicity classes by using the V0M estimator, where the events are divided into multiplicity classes (cf. Section 3.2) based on the combined number of hits in the two V0 counters (cf. Section 4.3). The track selection was largely the same as the one used in Table 8.1, excluding the PID cuts, with the most notable difference that different distance-of-

closest approach (DCA) cuts were used, namely $\mathrm{DCA}_z < 3.2\,\mathrm{cm}$ and $\mathrm{DCA}_{\mathrm{xy}} < 2.4\,\mathrm{cm}$[1]. Particle efficiencies were calculated like described in Section 8.6, but using a different event generator (HIJING [118]), and since flow measurements are quite sensitive to detector non-uniformities, a full efficiency correction in $(p_{\mathrm{T}}, \eta, \varphi)$ space was carried out. Systematic uncertainties were calculated as described in Section 8.10, but using different variations of course.

The flow coefficients were measured using multi-particle cumulants – defined in Section 3.7.2 – which were calculated by summing generalised flow vectors

$$Q_{\mathrm{n},p} = \sum_{k=1}^{M} w_k^p \exp(\mathrm{i}\mathrm{n}\varphi_k),$$

recursively through the generic framework [60], where w_k are particle weights applied to account for efficiency losses in the detector. The exact details are quite complicated and will not be given here, but the most important end results are that

$$\langle\langle 2\rangle\rangle = \frac{|Q_{\mathrm{n},1}|^2 - Q_{0,2}}{Q_{0,1}^2 - Q_{0,2}}$$

and

$$\langle\langle 4\rangle\rangle = \frac{|Q_{\mathrm{n},1}|^4 - 4|Q_{\mathrm{n},1}|^2 Q_{0,2} + 2Q_{0,2}^2 + |Q_{2\mathrm{n},2}|^2 - 6Q_{0,4} + \mathrm{Re}(8\,\mathrm{Q}_{\mathrm{n},1}^*\mathrm{Q}_{\mathrm{n},3} - 2\,\mathrm{Q}_{2\mathrm{n},2}^*\mathrm{Q}_{\mathrm{n},1}^2)}{Q_{0,1}^4 - 6Q_{0,1}^2 Q_{0,2} + 3Q_{0,2}^2 + 8Q_{0,1}Q_{0,3} - 6Q_{0,4}}.$$

Inserting these into Eqs. (3.11)–(3.15) then yield expressions for $v_{\mathrm{n}}\{2\}$ and $v_{\mathrm{n}}\{4\}$.

To suppress non-flow in two-particle cumulants, correlations were taken between tracks separated by an η gap larger than 1 in the TPC, or by using the scalar product method to provide an η gap larger than 2 and increase the statistics [58]. In the latter case, $v_{\mathrm{n}}\{2, |\Delta\eta| > 2\}$ was calculated as

$$v_{\mathrm{n}}\{\mathrm{SP}\} = \langle\langle \exp(\mathrm{i}\mathrm{n}\varphi_k Q_{\mathrm{n}}^*)\rangle\rangle \,/ \sqrt{\frac{\langle Q_{\mathrm{n}} Q_{\mathrm{n}}^{\mathrm{A}*}\rangle \langle Q_{\mathrm{n}} Q_{\mathrm{n}}^{\mathrm{B}*}\rangle}{\langle Q_{\mathrm{n}}^{\mathrm{B}*} Q_{\mathrm{n}}^{\mathrm{A}*}\rangle}},$$

where Q_{n} is measured in V0A (located at $2.8 < \eta < 5.1$) by estimating φ_k from the centre of each sector and using the signal amplitude as weight, $Q_{\mathrm{n}}^{\mathrm{A}*}$ is measured in V0C ($-3.7 < \eta < -1.7$), and $Q_{\mathrm{n}}^{\mathrm{B}*}$ is measured from the tracks in the ITS and TPC.

7.3 Results

Measurements of v_2, v_3, and v_4 in Xe–Xe collisions, using two- and multiparticle cumulants, are shown in the top panel of Figure 7.1a. Following Eq. (3.16), the ratio $v_2\{4\}/v_2\{2\}$ is sensitive to flow fluctuations, and hence to the initial state (IS). Therefore, to test the IS condition and hydrodynamic description, this ratio is compared to

[1] With the much larger number of tracks in a heavy-ion collision than a pp collision, the requirements for the DCA cuts become vastly different. Moreover, it is less important to reject decays from secondary particles in an inclusive-charged particle analysis, which makes it possible to loosen the DCA_{xy} cut.

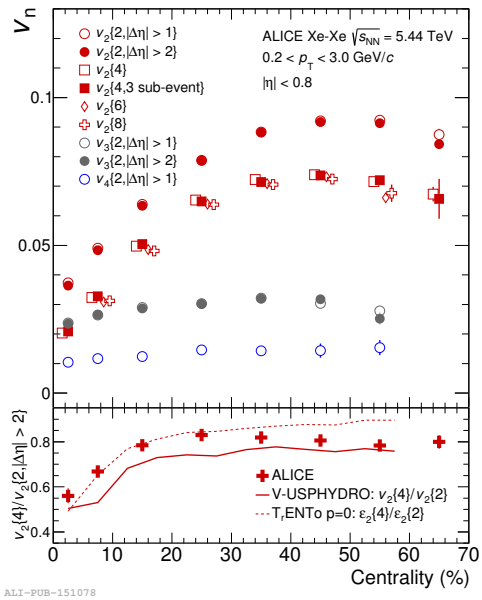

(a)

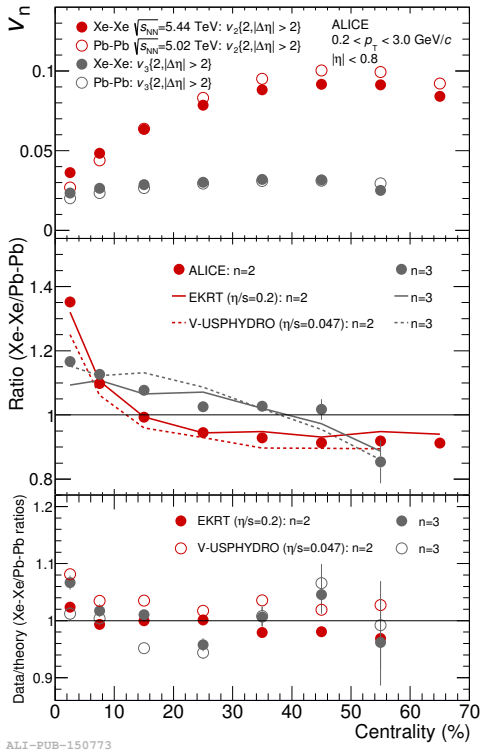

(b)

Figure 7.1 – **(a)** Upper panel: $v_n\{m\}$, $2 \leq n \leq 4$, as a function of centrality in Xe–Xe collisions at 5.44 TeV, for various orders m and pseudorapidity gaps. Lower panel: Comparison of the ratio $v_2\{4\}/v_2\{2, |\Delta\eta > 2|\}$ with a hydrodynamic calculation where a T$_R$ENTo IS model has been propagated using V-USPHYDRO [119, 120]. The results are also compared with $\varepsilon_2\{4\}/\varepsilon_2\{2\}$ from the IS model. **(b)** Top panel: Comparison of $v_n\{2, |\Delta\eta > 2|\}$, n = 2, 3, between Xe–Xe and Pb–Pb collisions as a function of centrality. Middle panel: Ratios between flow coefficients in Xe–Xe and analogous ones in Pb–Pb as a function of centrality, compared to hydrodynamic calculations from EKRT ($\eta/s = 0.2$) and V-USPHYDRO ($\eta/s = 0.047$) [119, 123], respectively. Bottom panel: Double ratio of data and theory for the two models. Figures taken from Ref. [117].

a hydrodynamic calculation, using V-USPHYDRO with $\eta/s = 0.047$ [119], as shown in the lower panel of Fig. 7.1a. The IS is modelled by T$_R$ENTo [120], which is an effective model for modelling the entropy distribution in the event, and hence the eccentricity on an event-by-event basis. In this model, the nuclear deformation in Xe is taken into account, which modifies the distribution of nucleons (Eq. (3.1)) to [121]

$$\rho(r, \theta) = \frac{\rho_0}{1 + \exp\left((r - R(1 + \beta_2 Y_{20}(\theta)))/a\right)}, \tag{7.1}$$

where θ is the polar coordinate and Y_{20} is a Bessel function of the second kind (w has been assumed to be zero in this parametrisation). The nuclear deformation is given by β_2, which for ^{129}Xe is predicted as $\beta_2 = 0.162$ [122], but has never been measured directly.

The results are also compared with the eccentricity ratio $\varepsilon_2\{4\}/\varepsilon_2\{2\}$ obtained directly from the IS model. Both of these predictions follow the same trend as the data, although both deviate somewhat from the data in parts of the centrality range.

In the upper panel of Figure 7.1b, $v_2\{2, |\Delta\eta > 2|\}$ and $v_3\{2, |\Delta\eta > 2|\}$ are compared between Xe–Xe collisions at 5.44 TeV, and Pb–Pb collisions at 5.02 TeV, as a function of centrality. In the middle panel, the ratios between the results from these two systems are compared with two different hydrodynamic calculations, EKRT using $\eta/s = 0.2$ [123] and V-USPHYDRO using $\eta/s = 0.047$. The deformation described by Eq. (7.1) is taken into account in both of these models.

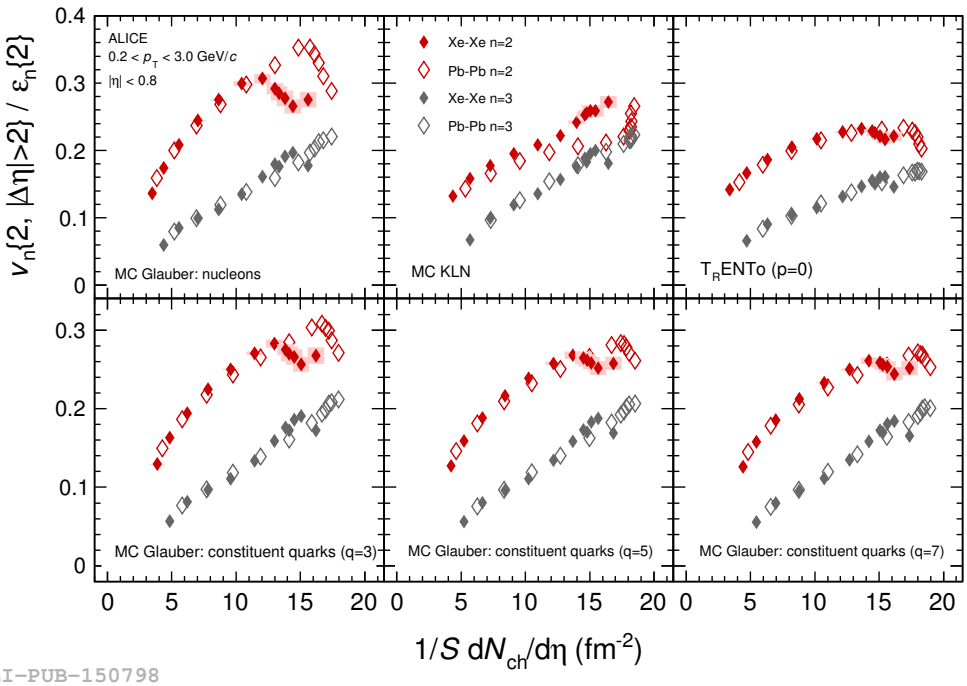

ALI-PUB-150798

Figure 7.2 – $v_n\{2, |\Delta\eta > 2|\}/\varepsilon_n\{2\}$ as a function of transverse density in Xe–Xe and Pb–Pb for n = 2 and n = 3, for various IS models. Top left: MC Glauber using nucleons as sources [35]. Top centre: MC KLN [125]. Top right: T$_R$ENTo [120]. Bottom: MC Glauber using 3, 5, and 7 constituent quarks as sources [35], respectively. Figure taken from Ref. [117].

Following Eq. (3.8), the ratio

$$\frac{v_n\{2, |\Delta\eta > 2|\}}{\varepsilon_n\{2\}}$$

provides an estimate of the scaling factor κ_n, which for n = 2 and n = 3 is expected to scale with $1/S\,\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}\eta$. The $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}\eta$ were provided from other studies [91, 124]. In Figure 7.2, this scaling is tested for a few different IS models: MC Glauber with nucleons and 3, 5, and 7 constituent quarks ($q = 3, 5, 7$) [35], respectively, as sources (cf. Section 3.2.1); MC KLN [125], which is a CGC model; and the T$_{\mathrm{R}}$ENTo model. MC Glauber with $q = 5$ or $q = 7$, as well as T$_{\mathrm{R}}$ENTo, generally yield good scaling, although there is a sharp decrease in the ratio at high transverse density. MC Glauber with nucleons as sources, as well as MC KLN, scale poorly across the systems.

7.4 Discussion

The fact that both the hydrodynamic prediction of $v_2\{4\}/v_2\{2, |\Delta\eta > 2|\}$ and the corresponding eccentricity ratio agree reasonably well with the data in Fig. 7.1a indicates that Eq. (3.8) holds approximately and that flow fluctuations are preserved in the hydrodynamic expansion. Moreover, it seems T$_{\mathrm{R}}$ENTo models the initial eccentricity fairly well.

Figure 7.1b indicates that both EKRT and V-USPHYDRO are able to model differences in IS and medium response between Pb and Xe fairly well. The lower v_2 in Xe at mid-central collisions is expected to be due to a larger viscous damping in Xe [119, 123], since ε_2 should be quite similar for the two systems. The peak in central collisions is mostly due to the deformation of the Xe nucleus, enhancing flow at central collisions. Since the Xe nucleus is smaller, it is expected to be more affected by flow fluctuations, which is the most likely reason for the larger v_3 (and v_2) in central Xe collisions.

The sharp decrease in $v_2\{2, |\Delta\eta > 2|\}/\varepsilon_2\{2\}$ at high transverse density seen in most models for Pb–Pb in Fig. 7.2 may indicate some shortcomings in the modelling of the IS in ultra-central collisions [117]. One reason may be that the model-based value used for β_2 is not entirely correct. Therefore, this measurement can give some insight into the deformation. By interpolating between the Pb–Pb data points using one of the models with good scaling, one can determine how large decrease in $\epsilon_2\{2\}$ for central Xe–Xe that is required to match the interpolated $v_2\{2\}/\varepsilon_2\{2\}$ for Pb–Pb, and then rescale β_2 to match this reduction in eccentricity. This yields $\beta_2 = 0.14 \pm 0.01$, where the uncertainty is model dependent[2].

The requirement to use $q \geq 5$ in MC Glauber for good scaling shows that one needs to use several constituent quarks as sources, i.e. take into account subnuclear structure, for this approach to work properly (also indicated in Ref. [126]). Moreover, MC KLN can be ruled out from this study, whereas also these results favour T$_{\mathrm{R}}$ENTo. It is quite intriguing that the effective models succeed in describing the initial states, whereas the only tested model based on QCD inspired effects fails, which means that still much work is required to understand the underlying mechanisms.

[2]This calculation was done by Anthony Timmins, who was one of the members in the paper committee for Ref. [117], but it never made it into the final publication.

7.5 Conclusions

Measurements of v_2, v_3, and v_4 in Xe–Xe collisions have given valuable information about the initial state and hydrodynamic propagation in heavy-ion collisions. These results indicate that flow fluctuations are preserved through the hydrodynamic propagation. Moreover, both EKRT and V-USPHYDRO can describe differences between Xe–Xe and Pb–Pb collisions. These models show that the Xe nucleus is deformed. Finally, the data favour the effective models $\mathrm{T_R}$ENTo IS model and MC Glauber with multiple quarks as sources, but rule out the CGC model MC KLN, as well as MC Glauber models not taking into account subnuclear structure.

Chapter 8

Analysis of Ξ-Hadron Correlations in 13 TeV pp Collisions

8.1 Introduction and Overview

One of the QGP signatures which has been probed experimentally is the enhanced yields of multistrange hadrons in heavy-ion collisions as compared to small systems (pp or p–Pb collisions), known as *strangeness enhancement*, cf. Section 3.6. Until recently, this was thought to be well understood, until a similar enhancement was observed also in small systems with high density. As discussed in Chapter 3, this is just one of several QGP observables where similar effects have been observed in small systems, but it has the advantage that there are two model very different models that both predict the observed results reasonably well, namely DIPSY with the rope hadronisation framework now enabled in PYTHIA, and the core-corona model EPOS. Comparisons between data and predictions from these two models is shown in Fig. 3.31. While DIPSY does a better job in quantitatively describing strangeness enhancement, this measurement alone is not enough to favour the rope hadronisation mechanism over the core-corona approach; it mostly tells that this model is better tuned to data. Therefore, the purpose of this analysis is to test the hadronisation mechanisms proposed by these models through angular correlations, and in the extension – by studying key features of the correlation function – go beyond these predictions to get a better understanding of what is happening in these collisions.

As strangeness S is conserved in strong interactions, one expects that when hadrons are formed, hadrons with $S < 0$ are exactly compensated by hadrons with $S > 0$, resulting in correlations between these hadrons. A prediction of the string/rope model is that strange hadrons are produced directly by $s\bar{s}$ pair breakings, so there will be strong local correlations (in phase space) between strange and anti-strange hadrons in the same event, even between different hadronic species. On the other hand, if strange quarks are produced early in the collision and strange hadrons later, as one would

expect in a thermal picture (this is the way strangeness is produced in EPOS), these correlations might be non-local (in phase space).

In this study, these predictions are tested by comparing angular correlations between the multistrange baryon Ξ^- (quark content dss) – or its antiparticle – and single-strange hadrons (kaons and Λ baryons, where $\mathrm{K}^+ = \mathrm{u\bar{s}}$ and $\Lambda = \mathrm{uds}$) with correlations between Ξ^- and non-strange hadrons (pions and protons). To begin with, $\Xi^- - \mathrm{K}^+$ ($\overline{\Xi}^+ - \mathrm{K}^-$) correlations are compared with the reference $\Xi^- - \pi^+$ ($\overline{\Xi}^+ - \pi^-$), in pp collisions at $\sqrt{s} = 13\,\mathrm{TeV}$. Since these correlations also may be influenced by non-local jet-like correlations and the underlying event, a subtraction to same-quantum number ($\Xi^- - \mathrm{h}^-$ and $\overline{\Xi}^+ - \mathrm{h}^+$, where h is either π or K) correlations is done as well, which really is the most relevant observable for this analysis. As an extension, $\Xi^- - \overline{\mathrm{p}}$ ($\overline{\Xi}^+ - \mathrm{p}$) correlations are compared with $\Xi^- - \overline{\Lambda}$ ($\overline{\Xi}^+ - \Lambda$) correlations. Here, baryon number conservation also affects the measurements, and therefore this sheds light on the currently quite poorly understood baryon production mechanism, but this also complicates the interpretation of the results somewhat. Finally, an attempt is made at measuring $\Xi - \Xi$ correlations, but for this the statistical sample is quite limiting, making it difficult to draw any strong conclusions for the time being.

The results are compared to theoretical predictions by PYTHIA8 and EPOS LHC. For PYTHIA8, an unofficial tune of the rope hadronisation model is also included in the comparison (as no official tune exists yet), as well as an official tune of the junction model (cf. Section 3.8.2).

8.2 Introduction to Some Kinematic Variables: Rapidity Versus Pseudorapidity

Before moving on to the measurements, it is useful to discuss two kinematic variables which are frequently used in the analysis. These are *rapidity* and *pseudorapidity*.

While the azimuthal coordinate of the detector volume is simply defined by the angle, using a similar definition for the longitudinal component would be problematic. In particular, the distribution of the particle yields close to the centre of the detector does not scale with the longitudinal angle. Instead a much better variable is rapidity,

$$y \equiv \frac{1}{2} \ln \left(\frac{E + p_\parallel}{E - p_\parallel} \right),$$

where $E = \sqrt{m^2 + p^2}$ (where m and p are the particle mass and momentum, respectively) is the particle energy and $p_\parallel$ its longitudinal momentum, which is ideal since it is additive under Lorentz boosts [27, pp. 27-30], resulting in a smooth scaling of the particle yields with rapidity, which is Lorentz invariant. Due to the mass dependence, this is however impractical to use as definition of the detector coordinates, so instead pseudorapidity is used, which is the rapidity limit for massless particles and is obtained as

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

where θ is the longitudinal angle relative to the beam direction. When $p_{\mathrm{T}} \gg m$, η approaches y, but since the correlation function also involves low-momentum particles and there is a flattening of the pseudorapidity distribution with decreasing p_{T}, it is preferred to use rapidity in analyses. The detector volume is however limited in η, so in order to maximise the statistical sample, a cut has been applied in η instead of in y in this analysis. The two quantities are related to each other as

$$y = \ln\left(\frac{\sqrt{m^2 + p_{\mathrm{T}}^2 \cosh^2\eta} + p_{\mathrm{T}} \sinh\eta}{\sqrt{m^2 + p_{\mathrm{T}}^2}}\right). \tag{8.1}$$

8.3 Measuring the Correlation Function

The correlation function $\mathbb{C}(\Delta y, \Delta\varphi)$, where Δy is the difference in rapidity and $\Delta\varphi$ is the difference in azimuthal angle between trigger and associated particles, is defined as

$$\mathbb{C}(\Delta y, \Delta\varphi) = \frac{1}{N_{\mathrm{pairs}}} \frac{\mathrm{d}^2 N_{\mathrm{pairs}}}{\mathrm{d}\Delta\varphi \mathrm{d}\Delta y}, \tag{8.2}$$

where N_{pairs} is the number of trigger-associated particle pairs. For this analysis, the number of particles associated with each trigger particle is also relevant, and thus it makes more sense to normalise to the number of triggers N_{trig}, which yields the per-trigger yields

$$\mathbb{Y}(\Delta y, \Delta\varphi) = \frac{1}{N_{\mathrm{trig}}} \frac{\mathrm{d}^2 N_{\mathrm{pairs}}}{\mathrm{d}\Delta\varphi \mathrm{d}\Delta y}. \tag{8.3}$$

Regardless of normalisation, this is simply a distribution in relative angles between trigger and associated particles, and can thus ideally be constructed by filling a two-dimensional histogram over Δy and $\Delta\varphi$ in the selected events. However, for a real detector, where the phase space is limited and the detector acceptance is not uniform, the resulting distribution would then become convoluted with the detector acceptance. To mitigate this, one needs to divide the measured per-trigger yields $S(\Delta y, \Delta\varphi)$ with a distribution $B(\Delta y, \Delta\varphi)$ resulting from the convolution alone, i.e.

$$\mathbb{Y}(\Delta y, \Delta\varphi) = \frac{S(\Delta y, \Delta\varphi)}{B(\Delta y, \Delta\varphi)}, \tag{8.4}$$

where

$$S(\Delta y, \Delta\varphi) = \frac{1}{N_{\mathrm{trig}}} \frac{\mathrm{d}^2 N_{\mathrm{pairs}}^{\mathrm{signal}}}{\mathrm{d}\Delta y \mathrm{d}\Delta\varphi},$$

where 'signal' denotes that this is the distribution of pairs from the same event.

An unbiased way to do this is through event mixing, which means that each trigger is associated with particles in a set of *different* events, which should all be as similar as possible as the original event. Thus

$$B(\Delta y, \Delta\varphi) = \frac{1}{\mathcal{C}} \frac{\mathrm{d}^2 N_{\mathrm{pairs}}^{\mathrm{mixed}}}{\mathrm{d}\Delta y \mathrm{d}\Delta\varphi},$$

where $\mathcal{C}$ is a normalisation constant and 'mixed' denotes that these are from the mixed event. Here, the normalisation is chosen such that the central value equals unity, which was ensured by projecting $B(\Delta y, \Delta\varphi)$ in Δy over the full $\Delta\varphi$ range and fitting a function to it of the form

$$b(\Delta y) = a_0 + a_1|\Delta y| + a_2(\Delta y)^2, \tag{8.5}$$

where a_i, $i = 0 - 2$, are fit parameters.

To ensure large similarity between signal and mixed events, events with a difference in number of tracks $\Delta N \leq 5$ and difference in z coordinate of the collision vertex $\Delta\mathrm{vtx}_z < 1\,\mathrm{cm}$ were used for the event mixing, but apart from when measuring $\Xi - \Xi$ correlations, where it was done to increase statistics, it was not required that the mixed event contained a Ξ trigger (the statistics would not have allowed that). This was done in batches of 1000 000 events each. To minimise statistical errors, up to 60 mixed events were used per signal event, but fewer if 60 events matching the selection criteria did not exist within the batch. To further maximise the similarity between signal and mixed events, the pairs were sorted such that at first hand $|\Delta N|$ was minimised, and secondly $|\Delta\mathrm{vtx}_z|$.

For $\Xi - \Xi$ correlations, an additional constraint of $p_\mathrm{T}^\mathrm{trig} > p_\mathrm{T}^\mathrm{assoc}$ was applied to make sure that each pair was just counted once, and since a random ordering could potentially bias the correlation function. A demonstration of how the correlation function is calculated is shown in Fig. 8.1.

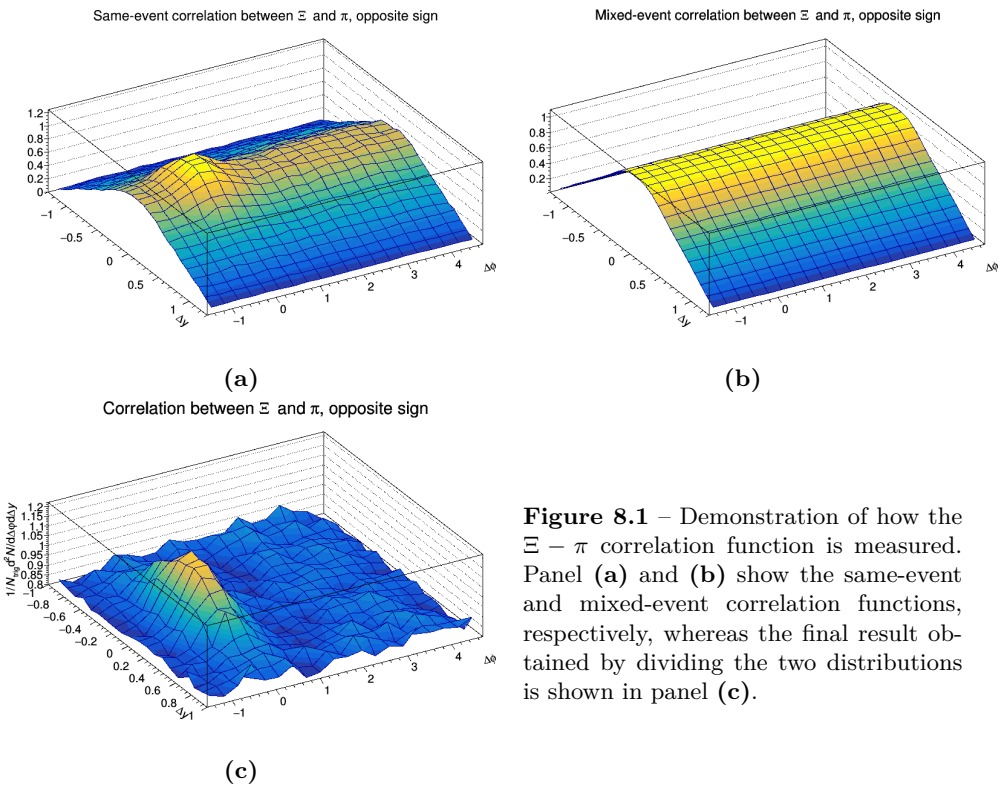

Figure 8.1 – Demonstration of how the $\Xi - \pi$ correlation function is measured. Panel **(a)** and **(b)** show the same-event and mixed-event correlation functions, respectively, whereas the final result obtained by dividing the two distributions is shown in panel **(c)**.

8.4 Event Selection

For the results presented here, a subset of the Run 2 pp data set covering one run period from 2016 (pp16k) and most of the 2018 data (9 run periods), corresponding to 820 million events, has been used. This data set was selected since it it has passed the quality assurance during previous correlation studies in ALICE, and the same gas mixture (based on argon) was used in the TPC for all these run periods[1]. This is associated with a set of Monte Carlo (MC) samples totalling 190 million events. This was obtained from simulations by the PYTHIA8 event generator [74] and all tracks were propagated through the detector using the available GEANT3/FLUKA simulation [127]. Using a MC run is important for several parts of the analysis, in particular for validating the method and calculate efficiencies, as will described later in this section. The GEANT3 software is unfortunately a bit dated, resulting in quite large uncertainties on some parameters, but unfortunately the simulation using the newer GEANT4 software is not yet finished for the full detector, so this could not be used for this thesis.

Following the original event selection, a couple of quality cuts were applied to filter away bad events. The following cuts were applied:

- The standard physics selection using the minimum-bias trigger was applied. For pp collisions this means that a hit is required in both the V0A and V0C detectors.
- Events with an incomplete DAQ were rejected.
- Events flagged as pile-up (i.e. containing overlapping data from other bunch crossings) based on the timing information from the V0 or SPD were rejected.
- It was required that the primary vertex was found by the SPD, determined to a resolution of $\leq 0.25\,\text{cm}$ and a dispersion of $< 0.04\,\text{cm}$.
- Events where the difference in z position of the primary vertex reconstructed using global tracks (ITS+TPC) and only the SPD was $\geq 0.5\,\text{cm}$ were rejected.
- The distance in z coordinate between the collision point and the centre of the detector was required to be $|\text{vtx}_z| < 10\,\text{cm}$ in order to ensure uniform efficiency within the detector.

Finally, events were divided into multiplicity classes based on the V0M centrality estimator (cf. 7.2). The correlations were measured both for minimum-bias events (no multiplicity selection), and for a low- and high-multiplicity sample, covering the 60% of the events with lowest multiplicity, and the 5% with the highest multiplicity, respectively. These intervals were selected since it was found that each contain about 20% of the total Ξ sample, which was chosen for statistical reasons.

[1] In this way, it was assured that tracking efficiencies etc. would not fluctuate too much between different run periods.

8.5 Particle Identification

For this analysis, it is important to disentangle the contributions to the correlation function from different particle species, which requires to use several of the PID techniques offered by the ALICE detector. In principle, there are two different categories of particles, namely long-lived particles that can be identified directly from the tracks they leave within the detector, and short-lived particles, which are identified from their decay products. All unstable particles relevant for this analysis decay through the weak interaction and as such, they are relatively long-lived compared to particles decaying through the strong or electromagnetic interactions. In fact, all particles used in this analysis live long enough to on the average have travelled a significant distance from the primary vertex before they decay, which has implications on the decay topology and make them easier to identify, as will be described below. A summary of some properties of the particles included in this analysis, along with a few other related particles, is given in Appendix C.

8.5.1 Identification of Long-Lived Particles (Direct Detection)

The only charged particles (except nuclei) which live long enough to be detected directly in the detector are pions, kaons, protons, electrons, and muons. Muons have signatures very similar to pions in the TPC and TOF detectors, requiring to use other means to detect them. Therefore, they are not relevant for this analysis, other than as a nuisance for the pion identification, which should be corrected for[2]. The other particle species can easily be identified in the central barrel, but since electrons have a huge background from pair production from photons interacting with the detector material, and are not very relevant for the physics case, they are not included in this study.

These particles can be identified through the specific energy loss in the TPC and their absolute velocity $\beta = v/c$, which is measured by the TOF. The average energy loss $\langle \mathrm{d}E/\mathrm{d}x \rangle$ of a particle traversing a medium is related to β through the Bethe-Bloch formula, which can be approximated as [10, p. 86]

$$\left\langle \frac{\mathrm{d}E}{\mathrm{d}x} \right\rangle = \frac{4\pi\alpha^2 h^2 Z^2 n_\mathrm{e}}{m_\mathrm{e}\beta^2} \left(\ln\left(\frac{2m_\mathrm{e}c^2\beta^2\gamma^2}{I} \right) - \beta^2 - \frac{\delta(\gamma)}{2} \right), \tag{8.6}$$

where $\alpha \simeq 1/137$ is the fine structure constant, h is Planck's constant, m_e is the electron mass, n_e is the electron density in the material, I is the ionisation potential, Z is the charge of the particle, $\gamma = 1/\sqrt{1-\beta^2}$ is its Lorentz factor, and δ is a factor entering from dielectric screening which is important only for energetic particles. Therefore, the energy deposit in a given material only depends on the velocity and charge of the particle. In the ALICE TPC however, one instead uses a parametrisation proposed by the ALEPH collaboration, which is on the form [100]

$$\frac{\mathrm{d}E}{\mathrm{d}x} = \frac{P_1}{\beta^{P_4}} \left(P_2 - \beta^{P_4} - \ln\left(P_3 + \frac{1}{(\beta\gamma)^{P_5}} \right) \right), \tag{8.7}$$

[2] As charged pions primarily decay into muons, and primary muons are created in electromagnetic processes (which are suppressed compared to strong processes in hadron collisions), most muons in the central barrel actually originate from pion decay, which makes this not strictly necessary, but should be done for increased accuracy.

where the parameters $P_1 - P_5$ are obtained by fitting the measured response to known particles being detected in the TPC. The quantities β and γ are related to the particle momentum p and mass m through

$$\beta = \frac{p}{\sqrt{p^2 + m^2}}, \qquad \gamma = \sqrt{1 + \left(\frac{p}{m}\right)^2},$$

where c has been set to 1. Finally, p/Z can be obtained by measuring the radius of curvature in the magnetic field.

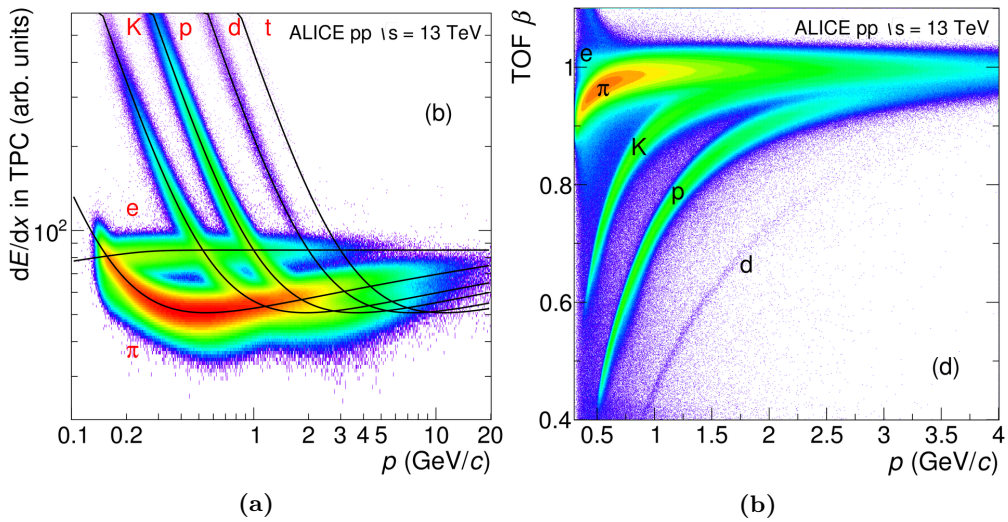

Figure 8.2 – **(a)** TPC signal (specific energy loss), along with lines over expected detector response for some of the most common particles, and **(b)** TOF β, for tracks originating from 13 TeV pp collisions, as a function of p. Figures taken from Ref. [128].

In practice, the measured $\mathrm{d}E/\mathrm{d}x$ will be smeared by the tail in the energy loss distribution[3] and the TOF β will be smeared by the detector resolution, resulting in a spread around the expected average. Yet, both the TPC and TOF provide clearly distinguishable particle bands when measuring the response as a function of momentum, as seen in Fig. 8.2. This is particularly the case at low momentum, but starting at $\sim 0.7\,\mathrm{GeV}/c$ for the TPC and $\sim 1.5\,\mathrm{GeV}/c$ for the TOF detector, the hadron bands start to mix. By combining information from both detectors, this limit can be extended somewhat, but already at $\sim 2\,\mathrm{GeV}/c$ kaon tracks are largely contaminated by pions. To obtain accurate correlation results, a high purity is required, so one needs to resolve this issue. There are basically three options:

- One can limit the analysis to a momentum range where the track purity is high, with consequently a large loss in statistics. Only about half of the tracks have

[3] $\mathrm{d}E/\mathrm{d}x$ follows a Landau distribution, which has a long tail towards large values. Therefore a few events may deposit significantly more energy than the average, limiting the resolution.

a successful TOF hit with a good matching between the TPC and TOF PID information[4], but this makes it possible to extend the momentum range somewhat.

- If desiring to reach a higher momentum and further increase the purity, a commonly used option is to apply a rejection cut, such that tracks where the PID information is overlapping with another species are rejected. Apart from the additional loss in efficiency, this approach is heavily reliant on the detector simulation, which has turned out to not simulate the distribution of the TOF β very well (and in particular the TOF tail, see below), resulting in a significant systematic uncertainty.

- A final option is to include all good tracks and subtract the contributions from misidentified tracks. As the correlation function is additive, this is possible if the misidentification fractions are known. This is the approach used in this analysis.

Before being used in any analysis, the tracking information is processed into ESD and AOD files (cf. Section 4.7), which are the actual analysis files. In the process, the detector responses are converted to the number of standard deviations $n\sigma^i$ from the expected detector response of each track type $i = \pi, \mathrm{K}, \mathrm{p}, \mathrm{e}$. The TPC response follows to a good approximation a Gaussian distribution, whereas the TOF distribution also has a tail in β, falling off as $\exp(-\beta)$ rather than the Gaussian behaviour of $\exp(-\beta^2)$. This is difficult to simulate, which is the reason why a rejection cut in $n\sigma^i$ is unreliable. For this analysis, an inclusion cut is instead applied, such that tracks with $|n\sigma^i| < 4$ for at least one i are included in the analysis. Here

$$n\sigma^i = \begin{cases} n\sigma^i_{\mathrm{TPC}}, & \text{no TOF hit}, \\ \sqrt{(n\sigma^i_{\mathrm{TPC}})^2 + (n\sigma^i_{\mathrm{TOF}})^2}, & \text{TOF hit}, \end{cases} \tag{8.8}$$

i.e. TOF information is used whenever available, otherwise just the TPC information is used. By using the TPC information alone, it would be impossible to disentagle the different track types where the detector response lines cross. Therefore, one needs to use both detectors, but including tracks also without a successful TOF hit increases the overall statistical significance considerably. By requiring that the combined $n\sigma^i$ of the TPC and the TOF is below 4, mismatched tracks are rejected. Moreover, the inclusion probability is not as dependent on the TOF tails as the rejection probability, making the Monte Carlo simulations more reliable. In Section 8.7, it is described how the contributions from different particle species are disentangled.

A summary of all track cuts used for the associated pions, kaons, and protons is given in Table 8.1. The reason for the higher minimum momentum for protons is that these are heavier and thus a higher momentum is required to escape the magnetic field. Moreover, at low momentum they are largely contaminated by material interactions, due to protons in the material being knocked out by other particles from the collision and entering the detector volume. The high-momentum limit is due to a poor spectra closure at higher momentum, cf. Section 8.9.1. A cut for the distance-of-closest-approach in the transverse plane (DCA_{xy}) to the primary vertex was used to suppress the contamination from secondaries. This was set to 7 times the resolution of this quantity [128], since it

[4]Sometimes the measured TOF β is associated with a TPC track with a specific energy loss corresponding to another particle species, which is known as TOF mismatch.

from statistical arguments is more or less certain that a track with a larger DCA_{xy} does not originate from the primary vertex. A cut in longitudinal DCA (DCA_z) was used to reduce out-of-bunch pileup. In experimental data, DCA_{xy} and DCA_z are measured by extrapolating the track back to the vicinity of the collision point. In addition, a couple of quality cuts based on the TPC clustering parameters were used to ensure a good PID quality and high resolution of the detector response. Finally, since in this analysis the tracks are used for measuring correlations with a Ξ baryon, it was required that all the tracks used for reconstruction of a Ξ candidate (see the next section) were rejected. For the trigger itself, this is a necessary constraint to avoid autocorrelations, but all such tracks were excluded to get the efficiency correction right.

Table 8.1 – Track cuts used for associated particles, $n\sigma^i$ is defined by Eq. (8.8).

Transverse momentum	$0.2 < p_\mathrm{T} < 3\,\mathrm{GeV}/c$ ($p_\mathrm{T} \geq 0.4\,\mathrm{GeV}/c$ for protons)
Pseudorapidity	$\lvert\eta\rvert < 0.8$
Number of TPC clusters (out of 159)	$N_\mathrm{cl} \geq 70$
Fraction of findable[a] TPC clusters	$N_\mathrm{cl}/N_\mathrm{cl}^f > 0.8$
Goodness-of-fit χ^2 per cluster (2 d.o.f./cluster)	$\chi^2/N_\mathrm{cl} < 4$
Number of associated SPD clusters	≥ 1
TPC and ITS refits	required
Longitudinal DCA	$\mathrm{DCA}_z \leq 2$ cm
Transverse DCA	$\mathrm{DCA}_{xy} < 0.0105 + 0.035 p_\mathrm{T}^{-1.1}$ cm
PID	$\lvert n\sigma^i\rvert < 4$ for either of $i = \pi$,K,p,e
Other constraints	not used in reconstruction of any Ξ candidates

[a] Clusters for which it is geometrically possible that they could be assigned to the track.

8.5.2 Identification of V^0 Particles and Cascades

Weakly decaying particles which can be identified through their decay topology are divided into V^0 particles and cascades. The former are neutral particles (which thus leave no direct tracks in the detector) which decay into a pair of oppositely charged long-lived daughter particles. Due to the applied magnetic field, these will bend in opposite directions, resulting in a "V" pattern in the detector. The strange particles which can be identified in this way are K^0_s (decays into two pions) and Λ (decays into a π^- and a proton). Due to time constraints, K^0_s have not been included in this analysis.

Cascades, on the other hand, are charged and decay into a long-lived charged particle (called bachelor) and a V^0 particle. In the detector, this will leave a bachelor track originating from a point quite close to the primary vertex, and a V^0 decay somewhat further out. A schematic figure of the detection pattern is shown in Fig. 8.3. There are two cascades which can be identified in this way, Ξ^- and Ω^- (and their antiparticles).

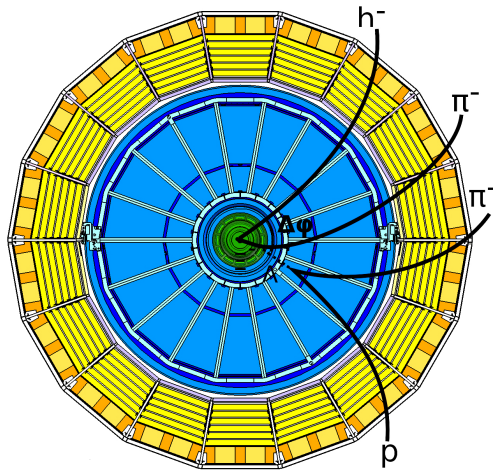

Figure 8.3 – Schematic view of a detection of a Ξ^- baryon through its decay products, along with another negative hadron h^-, projected on the transverse plane of the detector. The correlation between the two particles is measured from the angular distance. Projected in azimuthal angle, this is $\Delta\varphi$, as indicated in the figure. Modified version of figure taken from Ref. [129].

These decay primarily as [11]

$$\Xi^- \to \pi^- + \Lambda^0 \qquad \text{BR} = 99.9\%,$$
$$\Lambda^0 \to \pi^- + \text{p} \qquad \text{BR} = 63.9\%,$$

and

$$\Omega^- \to \text{K}^- + \Lambda^0 \qquad \text{BR} = 67.8\%,$$

where BR is the branching ratio. Due to statistical limitations, Ω baryons have been excluded from this analysis, which is the reason why it fully focuses on Ξ.

To identify the V^0s and cascades, tracks satisfying a selection of topological cuts following the expected decay topology are grouped together, which is done already in the central data processing stage. These were further refined in the analysis to reduce the background, as described henceforth. Moreover, a constraint on the deviation in TPC signal strength from the expected daughter particle of $|n\sigma| < 4$ was applied to each track. Apart from that, most track cuts listed in Table 8.1 were not applied, which for instance means that a successful refit is only required for the TPC. To complement this, it was required that for *at least one* of the daughter tracks, there is either a successful ITS refit or a TOF hit. This was done to ensure that the tracks indeed point back to the primary vertex, in order to suppress out-of-bunch pileup. If using the TPC alone, there may linger tracks from another bunch crossing in the detector, and to discard such tracks a faster detector is required. Finally, an invariant mass cut was applied to the decay products such that the deviation in invariant mass from the expected value

$|\Delta M^{\text{inv}}|$ should not be too large. The invariant mass is calculated as

$$M^{\text{inv}} = \sqrt{\left(\sum_i E_i\right)^2 - \left|\sum_i \mathbf{p}_i\right|^2}, \tag{8.9}$$

where E_i and $\mathbf{p}_i$ are the energies and 3-momenta of the cascade or V^0 daughter particles. The energy is calculated as

$$E_i = \sqrt{m_i^2 + p_i^2},$$

where the masses are obtained from the hypothesised decay particles.

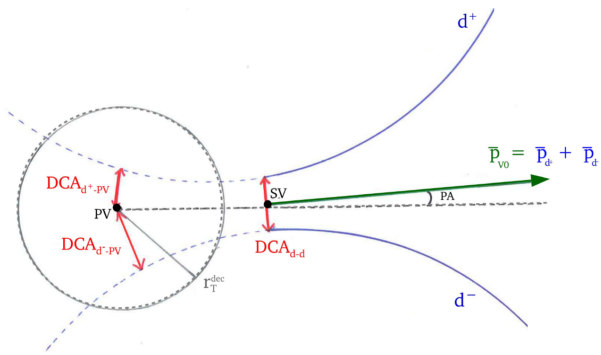

Figure 8.4 – Schematic drawing of a V^0 decay with definitions of all topological cut values. Legend: d^+: positive V^0 track, d^-: negative V^0 track, PV: primary vertex, SV: secondary vertex, DCA: distance-to-closest approach, see text for definitions, PA: pointing angle, $r_{\text{T}}^{\text{dec}}$: transverse decay radius; this is the minimum accepted distance between PV and SV. Figure taken from Ref. [130, p. 102].

Schematic drawings of a V^0 decay and a cascade decay are shown in Figs. 8.4 and 8.5, respectively, where all cut variables are defined. Three DCA variables are defined. These are the extrapolated distance-of-closest-approach between each daughter track and the primary vertex $\text{DCA}_{\text{d}-\text{PV}}$, the DCA between the V^0 daughter tracks $\text{DCA}_{\text{d}-\text{d}}$, and the DCA between the reconstructed V^0 (or cascade) in the momentum direction and the primary vertex $\text{DCA}_{\text{V0}-\text{PV}}$ ($\text{DCA}_{\text{casc}-\text{PV}}$). Here, a lower limit should be set on $\text{DCA}_{\text{d}-\text{PV}}$ in order to reduce the combinatorial background from primary tracks, and $\text{DCA}_{\text{d}-\text{d}}$ and $\text{DCA}_{\text{V0}-\text{PV}}$ ($\text{DCA}_{\text{casc}-\text{PV}}$) should be as small as possible. For the cascades, a lower cut on DCA between the V^0 daughter and the primary vertex is used in some analyses to suppress combinatorial background from primary Λs, but this was not done here since the simulations showed that this would not have the desired effect.

On top of this, there is a cut on the pointing angle, which is defined as the angular difference between the momentum sum of all daughters and the motion of the V^0 or cascade with respect to the primary vertex. In practice, one defines the cut in the cosine of this, $\cos(PA)$. A $\cos(PA)$ of 1 corresponds to full alignment between the momentum sum and the particle motion, which would be expected in an ideal detector. Therefore, the cut should be set close to unity. This is actually a very efficient way of cutting away combinatorial background, but since the MC description of it is quite poor, it should

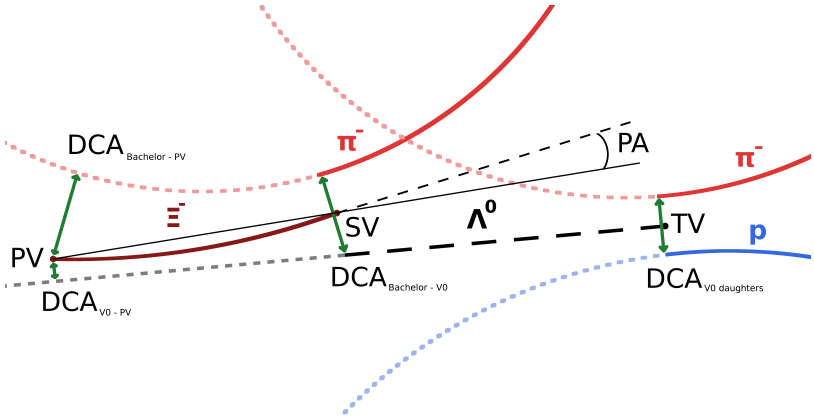

Figure 8.5 – Schematic drawing of a Ξ^- decay (cascade topology) with complementary topological cut definitions. Additionally, most cuts defined in Fig. 8.4 are applied also to the V^0 daughter and a r_T^{dec} cut is applied to the secondary vertex. $DCA_{Bachelor-V0}$ is not used in the analysis. New definitions in this figure: TV: tertiary vertex, $DCA_{V0\,daughters}$: same as DCA_{d-d} in the text. Figure taken from Ref. [131, p. 24].

not be set too tight, else the efficiency correction gets wrong. For cascades, a cut is also put on the daughter $\cos(PA)$ in many analyses, but due to the momentum kink received by the daughter particle upon the Ξ decay, this is actually not very accurate. In this analysis, it was shown that such a cut had little to no effect on the background reduction, and thus it was not applied.

The final cut is on the transverse decay radius, $r_T^{V^0}$ (r_T^{casc}), which is the distance between the primary and secondary vertices in the transverse plane. In particular at low momentum, there is a large risk of contamination from tracks originating from interactions with the material in the ITS (material background) and therefore a lower cut in $r_T^{V^0}$ or r_T^{casc} is used. Additionally, an *upper cut* should be applied at low momenta, since the V^0 or cascade will at most travel a few decay lengths before it decays, and therefore this may be used to reduce the combinatorial background. This can in principle be done by imposing a cut on the decay length, but since this requires taking into account the Lorentz boost, a simpler approach was to determine this from Monte Carlo simulations as described below.

To determine which cuts to use, a coarse invariant mass cut of $|\Delta M^{inv}| < 80\,\mathrm{MeV}/c$ was applied as a starting point, in combination with the $n\sigma$ cuts and a slightly tighter set of topological cuts than the standard ESD cuts. For the Λs, an additional cut of $|\Delta M^{inv}_{K^0_s}| > 10\,\mathrm{MeV}/c$ was applied, since otherwise there is quite a large contamination from K^0_s, which are much more common than Λ. The MC signal and background distributions were then studied for each cut (including invariant mass cut of the Λ daughter to the Ξ candidate, but not mother invariant mass cuts) in cut-p_T space. Regions with a large background and/or low signal-to-background ratio (signal/background $\lesssim 0.1$, but slightly higher tolerance at low p_T where the signal is low) were cut away. If there was a strong p_T dependence on these parameters, a p_T dependent cut was selected manually, and a power law was fitted to the data points. Otherwise, a flat cut was used. The

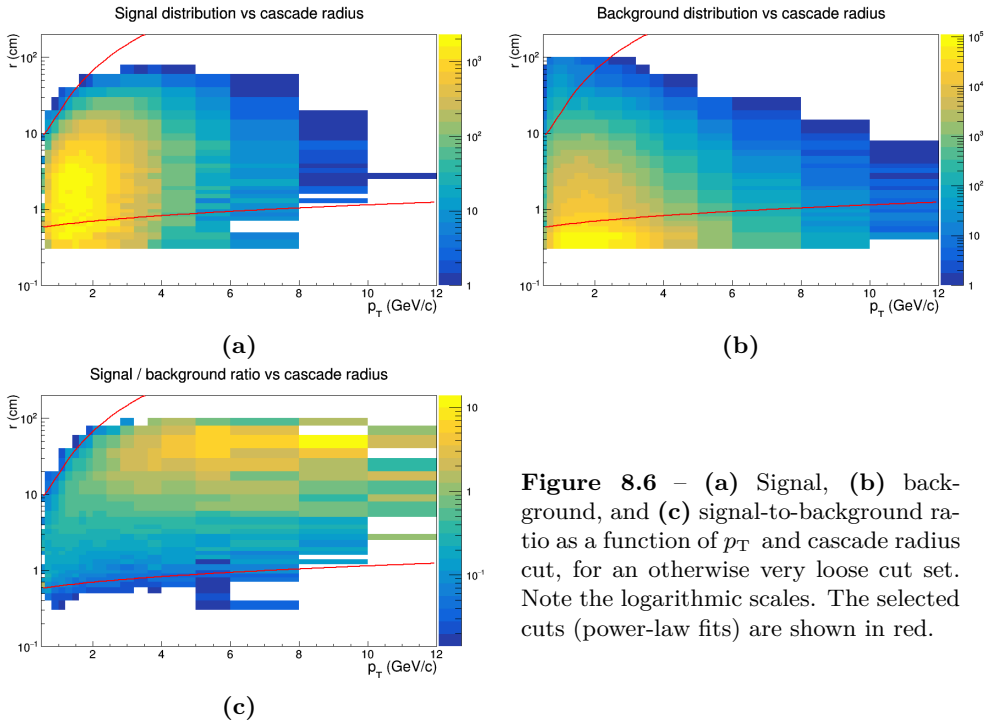

Figure 8.6 – **(a)** Signal, **(b)** background, and **(c)** signal-to-background ratio as a function of p_{T} and cascade radius cut, for an otherwise very loose cut set. Note the logarithmic scales. The selected cuts (power-law fits) are shown in red.

selection procedure is demonstrated for cascade radius in Fig. 8.6, with complementary figures for other cuts in Appendix D.2.

To refine the invariant mass cuts, a different approach was used. Here, a two-dimensional histogram in $p_{\mathrm{T}} - \Delta M^{\mathrm{inv}}$ space was filled with Ξ or Λ candidates from the ALICE data set, using the obtained topological cuts. Each slice in p_{T} was fitted by a double Gaussian to describe the signal[5], plus a background function. For Ξ and most of the p_{T} range in Λ, a one-dimensional polynomial turned out to describe the background adequately, but for $p_{\mathrm{T}} < 1\,\mathrm{GeV}/c$, there is a strong invariant-mass dependence of the background originating from material interactions. This required an additional term on the form

$$B(x) = c_1 \cdot \exp\left(-4\left(\frac{x - c_2}{c_3} - 1\right)\right) \cdot \left(\frac{x - c_2}{c_3}\right)^4, \tag{8.10}$$

where $c_1 - c_3$ are fit parameters (i.e. the same function as is used in Eq. (5.2)), to accurately describe the background. The final invariant-mass cut was then set as the 3σ limits of the wider Gaussian component, except the two lowest p_{T} bins ($0.4 < p_{\mathrm{T}} \leq 0.6\,\mathrm{GeV}/c$, used for feed-down corrections) for Λ, where the narrower component was used. In these bins the wider component sometimes had very low amplitude, yielding unstable fit results. To be able to use this in the correlation measurement, a power law was fitted to σ in the lower-p_{T} part of the spectrum, which was matched to a constant

[5]The invariant-mass peak has an approximately Gaussian shape due to the detector smearing, but the tails are generally longer than what is obtained from a Gaussian alone. This is remedied by superposing two Gaussian distributions, which has proven to describe the data very well.

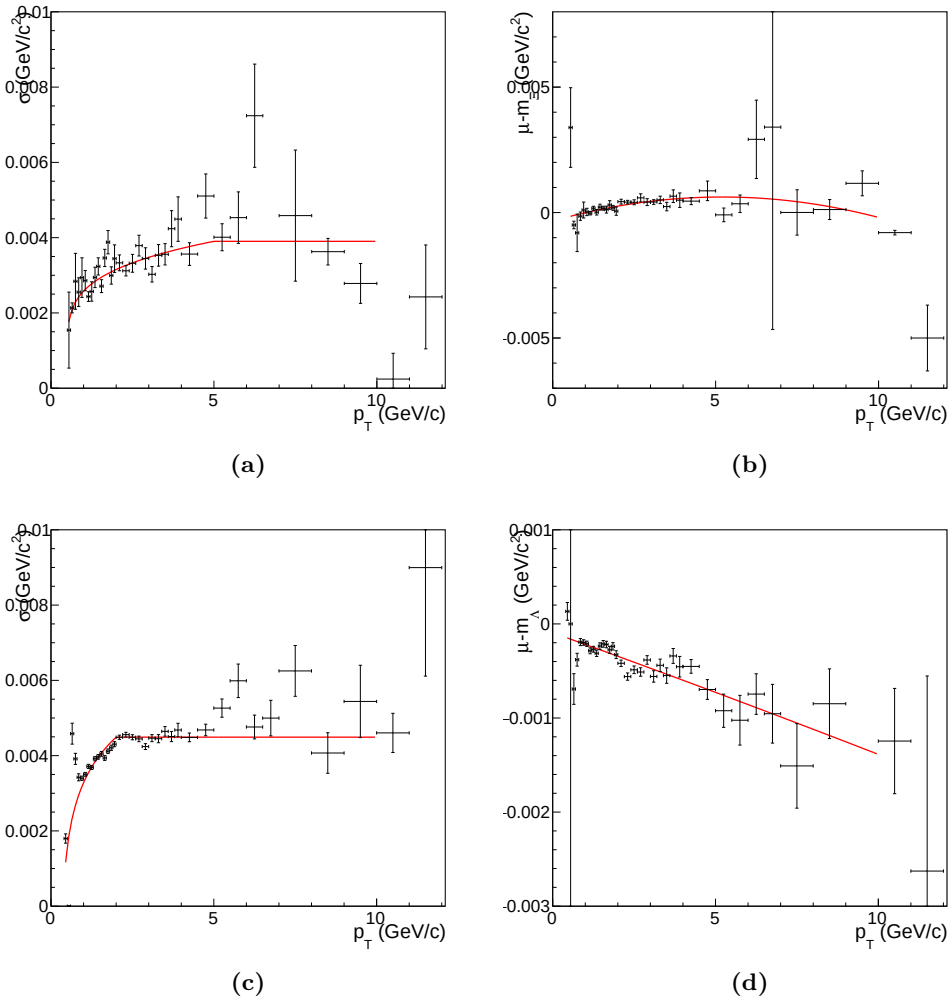

Figure 8.7 – Fitted values of **(a)** the width σ and **(b)** the mean μ of the invariant-mass peak for Ξ^-, and **(c)** and **(d)** the same quantities for Λ. The solid curves are fits to the parameters, which were used when selecting these baryons in the analysis.

value at the upper range (mostly done due to large fitting uncertainties in this range). The resulting fits for Ξ and Λ are shown in Figs. 8.7a and 8.7c, respectively. A similar fit was applied to the mean value μ of the same Gaussian component (this was not assumed to be centered at zero), but this was fitted by a two-dimensional polynomial instead, as shown in Figs. 8.7b and 8.7d.

Following reconstruction of all Ξ candidates in the event, it was just as for the tracks themselves ensured that the Λ baryons used as associated particles were not among the candidates for Ξ daughters. Here, this also helped in reducing the feed-down

Table 8.2 – Cuts used for the selection of Λ candidates. The topological cuts are defined in Fig. 8.4. For the p_{T}-dependent cuts, p_{T} is in GeV/c.

Daughter track cuts	
General	either ITS refit or TOF hit for *at least one* daughter track, not included in Ξ reconstruction
Transverse momentum	$0.15 < p_{\mathrm{T}} < 20\,\mathrm{GeV}/c$
Pseudorapidity	$\lvert\eta\rvert < 0.8$
PID selection	$\lvert n\sigma_{\mathrm{TPC\,PID}}\rvert < 4$
Pion daughter DCA	$\mathrm{DCA}_{\mathrm{d-PV}} > 0.10\,\mathrm{cm}$
Proton daughter DCA	$\mathrm{DCA}_{\mathrm{d-PV}} > 0.03\,\mathrm{cm}$
V^0 cuts	
Transverse momentum	$0.6 < p_{\mathrm{T}} < 12\,\mathrm{GeV}/c$
Pseudorapidity	$\lvert\eta\rvert < 0.72$
Invariant mass	$\lvert\Delta M^{\mathrm{inv}} + (0.05 + 0.13 p_{\mathrm{T}} - 0.016\,p_{\mathrm{T}}^2)\mathrm{MeV}/c^2\rvert$
	$< \begin{cases} 11.7 + 210\,((p_{\mathrm{T}} - 0.3)^{0.0017} - 1)\,\mathrm{MeV}/c^2, & p_{\mathrm{T}} < 2\mathrm{GeV}/c, \\ 13.6\,\mathrm{MeV}/c^2, & p_{\mathrm{T}} \geq 2\mathrm{GeV}/c \end{cases}$
DCA V^0-daughters	$\mathrm{DCA}_{\mathrm{d-d}} < 1.0\,\mathrm{cm}$
V^0 DCA	not applied
V^0 radius	$\max(0.2, -1.1 + 1.2(p_{\mathrm{T}} - 0.35)^{0.56}) < r_{\mathrm{T}}^{\mathrm{V0}} < 83 p_{\mathrm{T}} - 22\,\mathrm{cm}$
Cosine of pointing angle	$\cos(PA) > 0.995$
$\mathrm{K}^0_{\mathrm{s}}$ rejection	$\lvert\Delta M^{\mathrm{inv}}_{\mathrm{K}^0_{\mathrm{s}}}\rvert > 10\,\mathrm{MeV}/c$

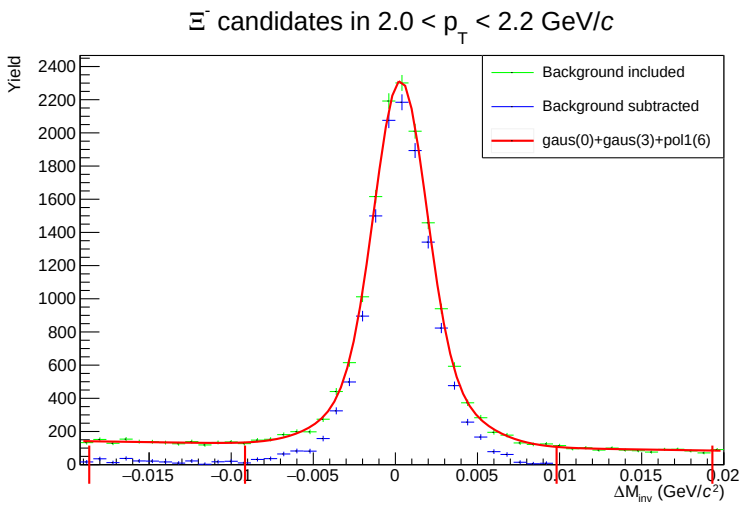

Figure 8.8 – Invariant-mass distribution of Ξ^- candidates in the transverse-momentum region $2.0 < p_{\mathrm{T}} \leq 2.2\,\mathrm{GeV}/c$, including a double-Gaussian fit + one-dimensional polynomial background. The red ticks mark the 3σ limits of the wider Gaussian component, which are used as limits for the signal region. The background is estimated from the sideband region, located between 3 and 6σ away from the mean.

Table 8.3 – Cuts used for the selection of Ξ candidates. The topological cuts are defined in Figs. 8.4 and 8.5. Throughout this table, p_{T} is in GeV/c. The cascade pseudorapidity cut is applied due to a very low reconstruction efficiency close to the detector edges, cf. Section 8.6.

Daughter track cuts	
General	either ITS refit or TOF hit for *at least one* daughter track
Transverse momentum	$0.15 < p_{\mathrm{T}} < 20\,\mathrm{GeV}/c$
Pseudorapidity	$\lvert\eta\rvert < 0.8$
PID selection	$\lvert n\sigma_{\mathrm{TPC\,PID}}\rvert < 4$
V^0 pion daughter DCA	$\mathrm{DCA_{d-PV}} > \max(0.03, -0.11 + 0.18\,(p_{\mathrm{T}} - 0.45)^{-0.36})\,\mathrm{cm}$
V^0 proton daughter DCA	$\mathrm{DCA_{d-PV}} > \max(0.03, -3.085 + 3.159\,(p_{\mathrm{T}} - 0.45)^{-0.019})\,\mathrm{cm}$
Bachelor DCA	$\mathrm{DCA_{bach-PV}} > 0.021 + 0.034\,(p_{\mathrm{T}} - 0.45)^{-0.73}\,\mathrm{cm}$
V^0 cuts	
Invariant mass	$\lvert\Delta M^{\mathrm{inv}}\rvert < 2.6 + 2.5\,p_{\mathrm{T}}\,\mathrm{MeV}/c^2$
DCA V^0-daughters	$1.5\,\mathrm{cm}$
V^0 radius	$1.2 < r_{\mathrm{T}}^{\mathrm{V0}} < 16 + 57\,(p_{\mathrm{T}} - 0.45)^{1.1}\,\mathrm{cm}$
Cosine of pointing angle	not applied
DCA V^0-primary vertex	not applied
Cascade cuts	
Transverse momentum	$1.2 < p_{\mathrm{T}} < 12\,\mathrm{GeV}/c$
Pseudorapidity	$\lvert\eta\rvert < 0.7$
Invariant mass	$\lvert\Delta M^{\mathrm{inv}} + (0.41 - 0.47 p_{\mathrm{T}} + 0.056\,p_{\mathrm{T}}^2)\mathrm{MeV}/c^2\rvert$ $< \begin{cases} 6.6\,(p_{\mathrm{T}} - 0.5)^{0.24} + 2.1\,\mathrm{MeV}/c^2, & p_{\mathrm{T}} < 5\mathrm{GeV}/c, \\ 11.7\,\mathrm{MeV}/c^2, & p_{\mathrm{T}} \geq 5\mathrm{GeV}/c \end{cases}$
Cascade DCA	$\mathrm{DCA_{casc-PV}} < \min(2.0, 0.007 + 1.34\,(p_{\mathrm{T}} - 0.45)^{0.68})\,\mathrm{cm}$
Cascade radius	$0.57 + 0.09(\,p_{\mathrm{T}} - 0.45)^{0.81}\,\mathrm{cm} < r_{\mathrm{T}}^{\mathrm{casc}} < 9 + 27\,(p_{\mathrm{T}} - 0.45)^{1.7}\,\mathrm{cm}$
Cosine of pointing angle	$\cos(PA) > \max(0.993, 0.9983 - 3.2 \cdot 10^{-3}(p_{\mathrm{T}} - 0.45)^{-2.25})$

contamination, cf. Section 8.8. Summaries of all cuts – topological as well as others – for the (primary) Λ and Ξ selections are given in Table 8.2 and 8.3, respectively.

With all cuts applied, some background still remains, which needs to be corrected for. While the spectra can be obtained directly from the invariant-mass fits, this can in practice not be done when measuring the correlation function. Instead, the background contribution needs to be subtracted from the full correlation function within the invariant-mass window in order to retain the signal correlation. This was done through sideband subtraction, which means that one selects data from regions on either side of the invariant-mass peak and uses these regions to approximate the background under the peak. These should be as close to the peak as possible to make sure that the background correlations do not change too much, but be sufficiently far away as to (almost) not contain any signal. For Ξ, it turned out that background regions between $3 - 6\sigma$ away from the peak was enough, but for Λ some signal still remained in this window, and thus the region $4 - 7\sigma$ was used instead. If the background is linear, and the sum of the background windows is of the same width as the signal window, the sum of the background in the sideband region will be the same as in the signal region.

For the non-linear part, i.e. low-p_{T} Λ, the sidebands need to be scaled such that the integral of the background in the signal region is the same as for the sideband regions. Since the non-linearity originates from material interactions, whereas the linear

part is mostly combinatorial, with possibly different correlations to the trigger, the sidebands should be combined in such a way that the relative proportion between the two components in the signal region is preserved. The procedure for this is described in Appendix A.2, where also the obtained coefficients are listed. In the very lowest bin, there was an issue with the invariant-mass distribution for $\overline{\Lambda}$, and since this part of the spectrum was only used for scaling purposes for the feed-down corrections (cf. Section 8.8), it was decided to only use the Λ result here. A demonstration of the sideband subtraction for Ξ in the region $2.0 < p_{\mathrm{T}} \leq 2.2\,\mathrm{GeV}/c$ is shown in Figure 8.8. Analogous plots for a few other transverse-momentum regions for both Λ and Ξ are given in Appendix D.2.

8.6 Efficiency Corrections

When measuring correlations in data, one needs to correct for the efficiency loss in the detector, which is done through Monte Carlo simulations. For the trigger particles, this is done on-the-fly, with each trigger divided by the reconstruction efficiency $\epsilon_{\mathrm{trig}}(p_{\mathrm{T}}, \eta)$. When measuring correlations with Λ or other Ξ baryons, a similar factor was applied also to the associated particles. However, for the method used here, which is described in the next section, the correlation function for the associated particles detected as tracks in the detector was sliced in associated-particle momenta, and the efficiency corrections were applied just upon merging.

For both trigger and associated particles, efficiency corrections were calculated from Monte Carlo simulations, by dividing the reconstructed yields of each particle propaged through the detector simulation, by the generated yields from the event generator. For the generated particles, only primary particles were included, whereas for reconstructed, most particles originating from decay of secondary particles were included as well. This contamination is known as feed-down, and for minor sources the most efficient way to compensate for it is to include it in the efficiency corrections. For a correlation analysis, however, one should be careful about this, and therefore a more accurate approach was used for the largest feed-down sources. This is described in Section 8.8.

The efficiencies are momentum dependent, both due to the magnetic field and momentum dependent interactions with the materials. The former depends on transverse momentum p_{T}, particularly since low-p_{T} particles may be trapped in the magnetic field, while the latter rather depends on the total momentum, p. Moreover, longer tracks (larger $|\eta|$) activate more clusters in the TPC, and are thus more likely to satisfy the track cuts. On top of that, detector non-uniformities add an additional dependence in η and φ, but this is mostly compensated for by the event mixing, and thus it was concluded that no efficiency corrections in φ were necessary. For pseudorapidity on the other hand, it turned out that corrections were necessary for Ξ and Λ baryons – which have a large η dependence due to the increasing fraction of daughter tracks escaping when approaching the detector edges – but not for tracks that can be detected directly in the detector.

The efficiencies Ξ and Λ as a function of p_{T} and η are shown in Fig. 8.9 and 8.10, respectively. Efficiencies for pions, kaons, and protons were measured in p, using the same binning as used for the slicing of the correlation function (cf. Section 8.7). It turned out that these differed slightly between run periods; therefore they were

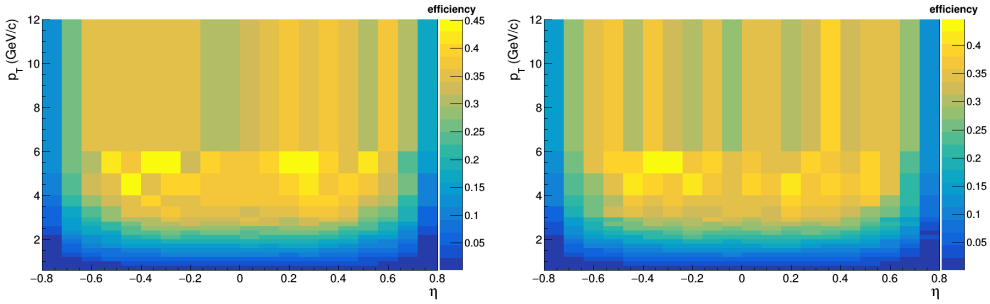

(a) **(b)**

Figure 8.9 – Efficiency in (η, p_{T}) space for **(a)** Ξ^- and **(b)** $\overline{\Xi}^+$.

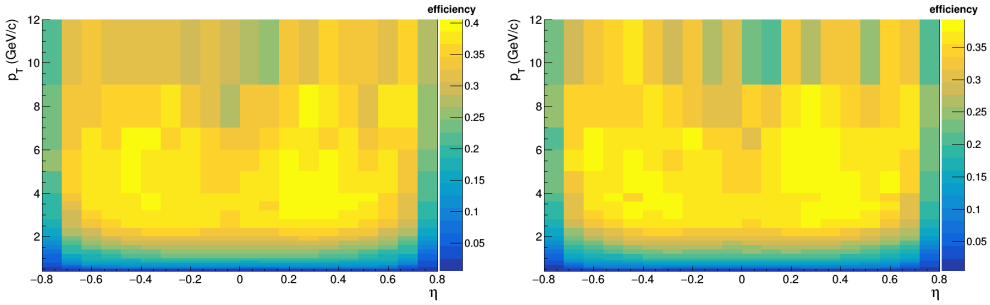

(a) **(b)**

Figure 8.10 – Efficiency in (η, p_{T}) space for **(a)** Λ and **(b)** $\overline{\Lambda}$. Secondary particles (feed-down) have been excluded from this calculation.

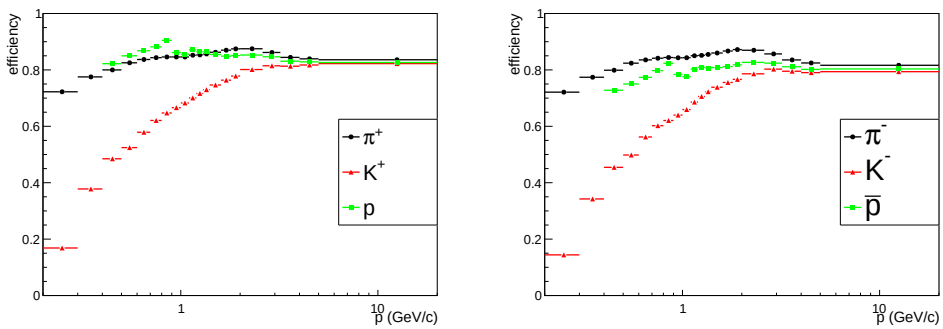

(a) **(b)**

Figure 8.11 – Particle efficiencies in the run period pp16k for **(a)** positive particles and **(b)** negative particles, not including contamination from secondary particles.

calculated for each run period separately. The results for the pp16k dataset are shown in Fig. 8.11.

8.7 Disentangling Contamination From Misidentified Tracks: the Linear Algebra Method

As mentioned in Section 8.5.1, all tracks satisfying the track cuts summarised in Table 8.1 were used in the analysis and a subtraction scheme was applied to remove the contamination from misidentified tracks. More specifically, linear algebra techniques were used to extract the *per-particle* correlations from the *per-track* correlations. In this section, it is described how this was done.

To begin with, the tracks were separated in those having a TOF hit and those without, since the purity gets significantly higher if using the combined TPC+TOF information compared to the TPC alone (but tracks without TOF information are used nevertheless due to the low TOF efficiency). At low momentum, however, the distinction power of the TPC is high and the TOF efficiency is very low. Therefore, for particles with low momentum, only the TPC information was used. Different momentum limits for the TOF were used for different track types. These are summarised in Table 8.4.

Table 8.4 – Momentum (p) limits for enabling the TOF detector for each particle type.

Particle type	**Momentum limit (GeV/c)**
π, e	0.4
K	0.6
p	0.9

To begin with, each track was associated with the particle species with the smallest $|n\sigma^i|$, as defined by Eq. (8.8), except that a pion veto was applied against electrons, due to the much larger abundance of pions. This yields eight different track classes, four for tracks with a TOF hit and four for those without. For each class, the misidentification fraction of each particle type was first measured in MC. This was done as a function of momentum, p, since the detector responses are dependent on p rather than p_{T}. Moreover, when measuring the misidentification fractions, as well as the efficiencies of associated tracks, the track yields were multiplicity weighted to mimic the presence of a Ξ baryon. This was done to assure that the fraction of tracks with a successful TOF hit, as well as the total efficiency, was correct, since these quantities may have a multiplicity bias. To obtain the misidentification fractions in data, it was assumed that the detector simulation is correct to first order, which is a reasonable assumption in the selected momentum region if not applying any strong rejection cuts (cf. Section 8.5.1), and is confirmed in Fig. 8.19. However, the particle ratios may be different at generator level compared to data, so this needs to be corrected for.

If the fraction of particle species j in MC is x_j^{MC}, and in data it is x_j^{data}, the contamination fraction in data will scale by $x_j^{\mathrm{data}}/x_j^{\mathrm{MC}}$, save for a normalisation factor, simply because there were more particles of this species to begin with. With normalisation, the misidentification fraction of species j for track type i in data becomes

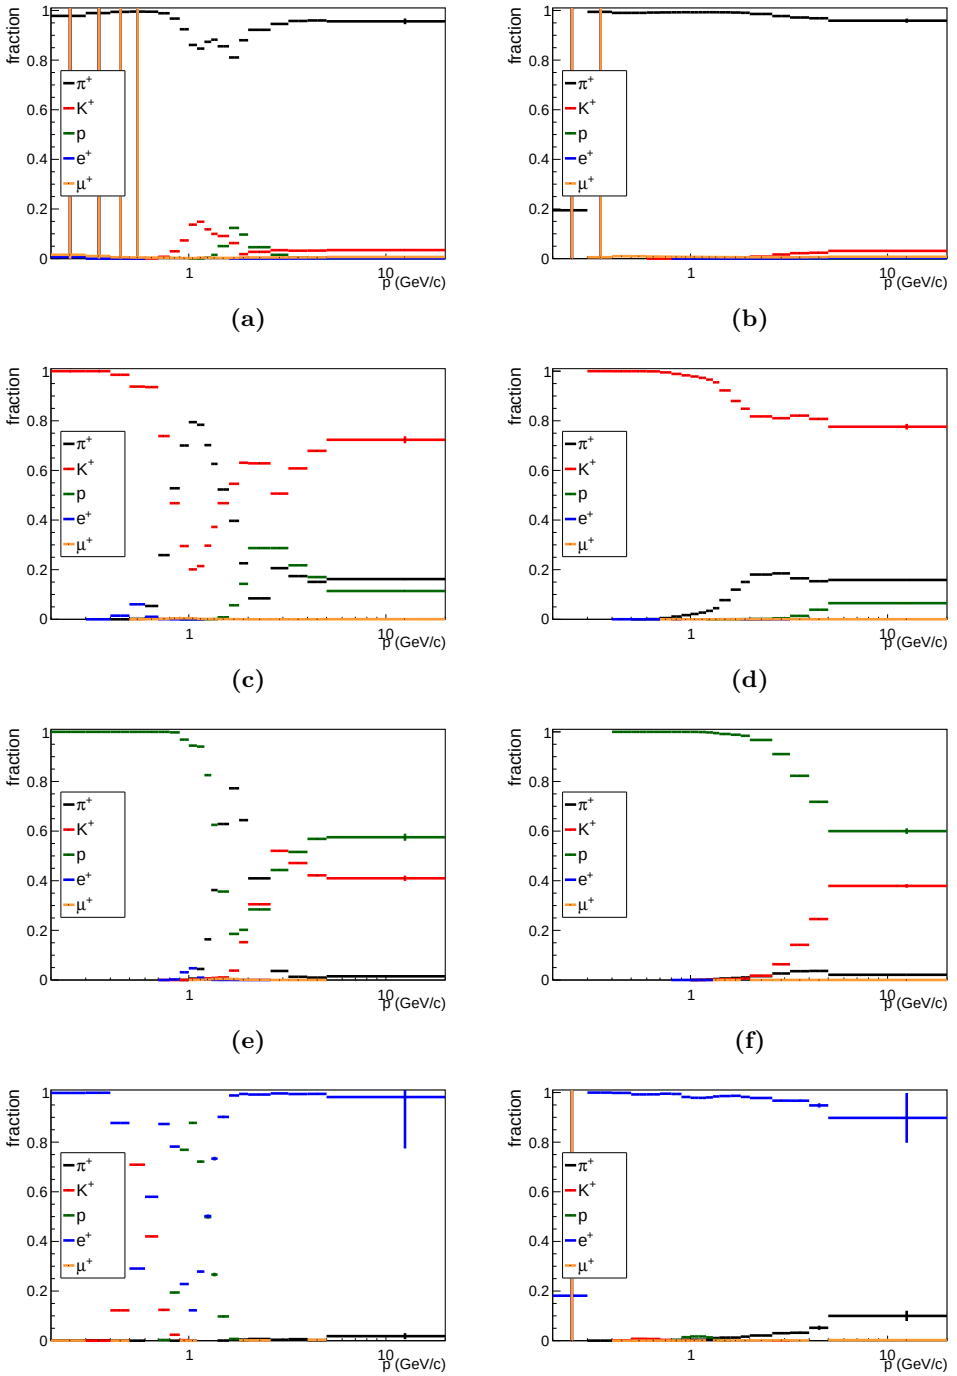

Figure 8.12 – Misidentification fraction in data for positive tracks, as a function of momentum. Left: TPC only tracks. Right: Combined TPC+TOF tracks. **(a-b)**: Pion tracks, **(c-d)**: kaon tracks, **(e-f)**: proton tracks, **(g-h)**: electron tracks.

$$a_{ij}^{\mathrm{data}} = a_{ij}^{\mathrm{MC}} \cdot \frac{x_j^{\mathrm{data}}/x_j^{\mathrm{MC}}}{\sum_k a_{ik}^{\mathrm{MC}} x_k^{\mathrm{data}}/x_k^{\mathrm{MC}}}, \tag{8.11}$$

where a_{ij}^{MC} is the misidentification fraction in MC. If letting y_i be the fraction of tracks of type i, one can form the equation system

$$\mathbf{A}^{\mathrm{T}}\mathbf{y} = \mathbf{x}, \tag{8.12}$$

where $\mathbf{A}$ is an 8×4 matrix with all a_{ij}, $\mathbf{y}$ is a column vector with all y_i and $\mathbf{x}$ is a column vector with all x_j. This identity is valid both in MC and in data, and Eq. (8.11) forms the link between the two systems. Hence, Eqs. (8.11) + (8.12) can be combined into a single system. Since $\mathbf{A}^{\mathrm{MC}}$ and $\mathbf{x}^{\mathrm{MC}}$ are known, by putting $\mathbf{y} \equiv \mathbf{y}^{\mathrm{data}}$, this combined system can be solved iteratively to obtain $\mathbf{A}^{\mathrm{data}}$ and $\mathbf{x}^{\mathrm{data}}$, i.e. the misidentification fractions and particle ratios in data. This was repeated in each momentum bin, and only for tracks in events containing a Ξ trigger since the correlation with a Ξ trigger will also affect the total yield. The obtained misidentification fractions for each track type are shown in Fig. 8.12. To account for muon contamination, it was assumed that the μ/π ratio is correct in MC, which is reasonable since most muons originate from pion decay. Since muons are so strongly mixed with pions, no attempt was made to separate these particle species at track level. Instead, the muon fraction was just subtracted off from the pion yield.

Here it has been assumed that the misidentification fraction is independent of the distance to the trigger particle. Following the same logic as above, if the trigger is giving rise to strong correlations, this is not the case if the desired particle is correlated differently with the trigger than other particles contaminating the track. Therefore, this correction should ideally be done in $(\Delta y, \Delta\varphi)$ space, which would require fitting and would thus be quite challenging. This was not attempted due to time constraints. The results from the MC closure test (cf. Section 8.9.1) indicate that this is probably not necessary for minimum-bias correlations, but may be so if pursuing a p_{T} dependent analysis.

To measure the correlation function, this was first separated into track types (defined by the most likely particle species), and calculated separately in each momentum bin. Moreover, the calculations were repeated for each mass hypothesis $((\pi, \mathrm{K}, \mathrm{p}, \mathrm{e}))$, since this affects Δy (cf. Eq. (8.1)), and the rapidity distribution needs to be correct at the final merging in order to give the correlation function the correct shape. To obtain the correlations to each particle species, one needs to solve the system

$$\mathbf{A}\tilde{\mathbf{S}}^{\mathrm{particle}} = \tilde{\mathbf{S}}^{\mathrm{track}},$$

where $\tilde{\mathbf{S}}$ is a column vector with either the same-event or mixed-event correlations to all tracks/particles, normalised to the number of tracks/particles[6]. Since $\mathbf{A}$ is an 8×4 matrix, this system is overdetermined and a least-squares solution is required in order to maximise the statistical significance. This has the solution [114, p. 93]

$$\tilde{\mathbf{S}}^{\mathrm{particle}} = (\mathbf{A}^{\mathrm{T}}\mathbf{W}\mathbf{A})^{-1}\mathbf{A}^{\mathrm{T}}\mathbf{W}\tilde{\mathbf{S}}^{\mathrm{track}}, \tag{8.13}$$

[6]This works, since the correlation function is additive, at least locally. Ideally, this system should be solved separately in each $(\Delta y, \Delta\varphi)$ bin, but this might not be possible due to time constraints. If not, most errors caused by only doing this globally should enter into the MC closure test, cf. Section 8.9.1.

where the weight-matrix $\mathbf{W} = \mathrm{diag}(\mathbf{N}^{\mathrm{track}})$, where $\mathbf{N}^{\mathrm{track}}$ is a vector with the number of tracks of each track type. These were not efficiency corrected, since this is used for determining the variance only. Moreover, the muon and pion fractions were merged exclusively in this step, since otherwise the efficiency correction would not be correct[7]. In this final step, the correlations obtained using the invariant mass of the desired particle species were used. The obtained weights for combining each track type are summarised in Fig. 8.13. Since a subtraction is sometimes needed, negative weights are possible.

Finally, in order to obtain the unnormalised correlation functions S and B used in

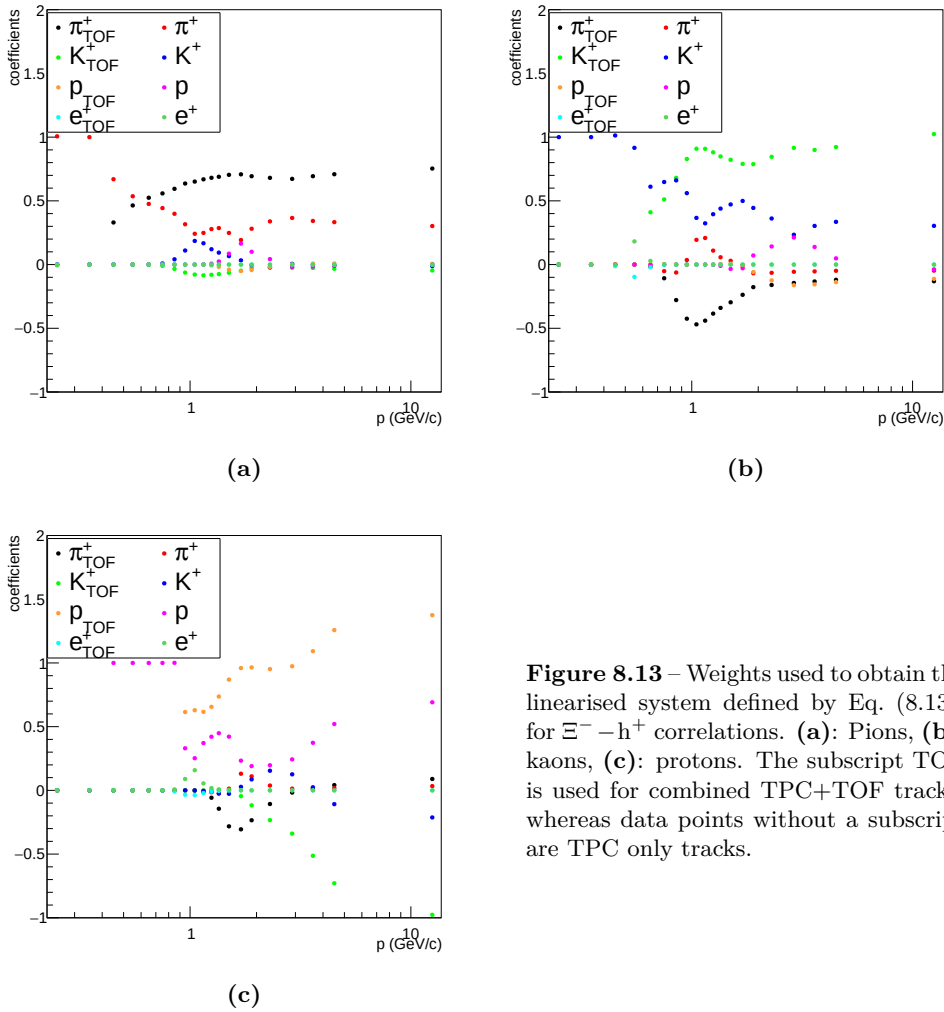

(a) **(b)**

(c)

Figure 8.13 – Weights used to obtain the linearised system defined by Eq. (8.13), for $\Xi^- - \mathrm{h}^+$ correlations. **(a)**: Pions, **(b)**: kaons, **(c)**: protons. The subscript TOF is used for combined TPC+TOF tracks, whereas data points without a subscript are TPC only tracks.

[7]Since the muon and pion tracks are merged, we are forced to measure the combined muon plus pion correlation function. The muon fraction was subtracted when evaluating Eq. 8.14.

Eq. (8.4), $\tilde{\mathbf{S}}^{\text{particle}}$ was multiplied by

$$\mathbf{N}^{\text{particle}} = \mathcal{E}(p)\mathbf{A}^{\mathrm{T}}\mathbf{N}^{\text{track}}, \tag{8.14}$$

where $\mathcal{E}(p)$ is a diagonal matrix with all particle efficiencies, which are summarised in Fig. 8.11. The results from the different momentum bins were merged before dividing the same-event and mixed-event correlation functions.

8.8 Feed-down corrections

Not all particles which leave tracks in the detector originate from the collision itself, but a significant fraction originates from decays of more short-lived particles. This is known as feed-down, and since the purpose of this analysis is to study the production mechanisms in the collisions, this needs to be corrected for.

There are two types of particles which contribute to the feed-down: weakly decaying particles and resonances. The latter decay via either the strong or the electromagnetic interaction and are not affected by any topological cuts. In particular strongly decaying particles can be considered to be part of the collision itself. These will affect the correlation function in various ways, but their interactions are rather studied from their combined contribution to the full correlation function than to try to isolate them. A possible exception is the ϕ meson, which to a good approximation is a pure $\mathrm{s\bar{s}}$ state [11, p. 281], and therefore can be used to study strangeness production, e.g. through $\phi - h$ correlations. The ϕ meson decays primarily into two kaons, but it has still been decided to include its contribution in the $\Xi - \mathrm{K}$ correlation functions.

The weakly decaying particles on the other hand, are long-lived enough to decay relatively far from the collision vertex (typically in the cm $-$ dm range) and therefore any contamination from these in the correlation function will probe different processes than the ones we are interested in. All weakly decaying particles which contribute significantly to the feed-down are included in the list of hadrons in Appendix C, along with their main decay channels. The particles with a decay length $> 1\,\mathrm{m}$ will mostly decay in the TPC or outside the detector volume and will thus not significantly contribute to the feed-down, which is ensured by the DCA cut listed in Table 8.1. Instead, the particles with decay lengths $c\tau$ in the cm range are the ones that contribute to the feed-down.

To make matters worse, particles from the collision may interact with material in the detector, in particular the ITS, where they may either scatter, knock out nucleons (or electrons) from the material, or produce new particles through various reactions (e.g. antimatter annihilation). This is known as material feed-down. Protons and pions from material feed-down will also contaminate the V^0 selection, but here it is corrected for by sideband subtraction (cf. Section 8.5.2). Both feed-down from weak decays and material interaction are reduced by the DCA cut in Table 8.1, but since it was decided not to apply a similar cut on V^0s due to its impact on the signal loss, the decay feed-down contribution to the Λ data set is quite significant.

To determine which feed-down sources are likely to significantly affect the results, the fraction of tracks originating from each of these sources as a function of p_{T} was extracted from MC simulations. The only significant sources ($\gtrsim 1\%$) were $\mathrm{K}^0_{\mathrm{s}}$ decay

into pions ($\sim 1.5\%$ of the tracks), Λ ($\sim 5\%$) and Σ^+ ($\sim 2\%$) decay into protons, material feed-down to protons ($\sim 1.5\%$; these are knocked out from the material, so this does not include antiprotons), and Ξ^0 ($\sim 6\%$) and Ξ^- ($\sim 5\%$) decay into Λ baryons. All other sources – including all contributions to kaons or Ξ – are small or negligible and were therefore assumed to be well enough compensated for by including them in the efficiency corrections. While being larger sources, there were no strong indications (in MC) that either K^0_s, Σ^+, or material interactions (which mainly originate from pions) would contribute substantially different to the correlation function than the particles subject to the contamination. While their correlation functions obviously have a different shape, the contribution from these sources is still small enough that it would require nearly an order-of-magnitude difference in correlation strength to significantly affect the results. Therefore, the same procedure was used also here, and any effect from it should enter into the systematic uncertainty through the MC closure test (cf. Section 8.9.1)[8].

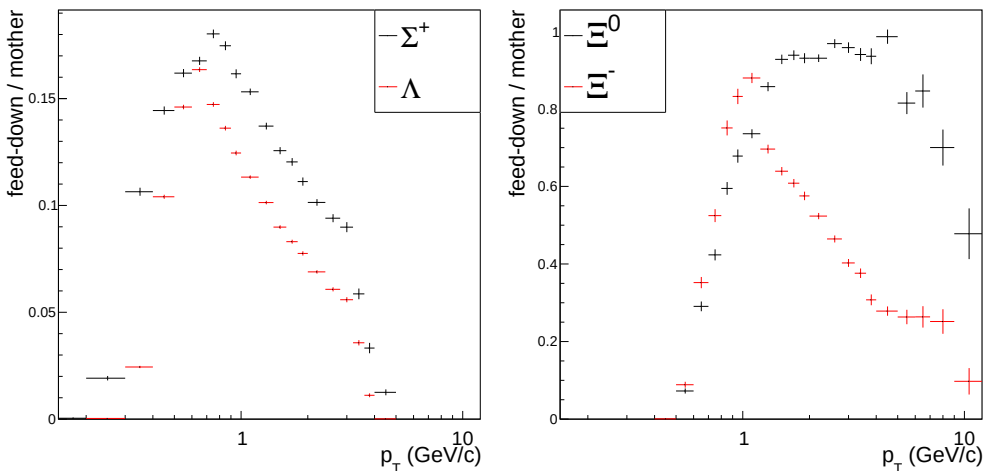

Figure 8.14 – Efficiency-corrected probabilities of detecting **(a)** a proton with $0.4 < p_{\mathrm{T}} < 3\,\mathrm{GeV}/c$ and **(b)** a Λ baryon with $0.6 < p_{\mathrm{T}} < 12\,\mathrm{GeV}/c$, originating from a selection of secondary particles, as a function of the transverse momentum of the mother particle.

This leaves the contributions from Λ decaying into protons, and Ξ^0 and Ξ^- decaying into Λ baryons, which are large enough to require a different approach. The $\Xi - \Lambda$ and

[8] One may argue that this puts too much trust on the MC model, which is relevant criticism. The dominant error source for the material contribution is the GEANT3 detector simulation, which is known to be imperfect, but since the dominant contribution is pions knocking out protons and the $\Xi - \pi$ and $\Xi -$ p correlation functions (cf. Chapter 9) are reasonably well described by PYTHIA, this simulation is probably good enough. Moreover, K^0_s yields are well-reproduced by PYTHIA, so a similar conclusion can be drawn also here. The contribution from Σ^+ on the other hand is more uncertain, but since this is very challenging to access experimentally, this is probably as close as one can get. The $\Xi^- - \Sigma^+$ correlation function is suppressed by only sharing a single quark, so this contamination should indeed be much smaller than the one from $\Xi^- - \Lambda$. The Σ^+ yield is more uncertain, and is probably underestimated in PYTHIA, but this is not experimentally accessible either.

$\Xi - \Xi$ correlation functions are expected to be significantly different from the ones for $\Xi - \mathrm{p}$ and $\Xi - \Lambda$, meaning that their contribution from feed-down needs to be measured separately and subtracted from the full correlation function. Moreover, the Λ and Ξ spectra are not very well reproduced by the MC generator. To account for these constraints, the MC generator was solely used to measure the average number of (efficiency corrected) detected daughter particles in the target p_{T} interval per generated mother particle, as a function of p_{T} and η for the mother (as it turned out to be a quite strong η dependency). The resulting distributions projected in p_{T} are shown in Fig. 8.14a for protons and Fig. 8.14b for Λ baryons. The full (η, p_{T}) maps can be found in Appendix D.2. The reason for the decrease in contamination from charged Ξ baryons to Λ, which is not present for Ξ^0, is the rejection of Λ baryons used to reconstruct Ξ candidates. These coefficients were then used as weights when measuring the correlation functions for feed-down correction.

Since there is no viable way of reconstructing Ξ^0s in the ALICE detector, one needs to make use of Monte Carlo methods to estimate the feed-down contribution from this particle, which was done by comparing the $\Xi^{\pm} - \Xi^0$ correlation functions with the ones for $\Xi^{\pm} - \Xi^{\pm}$ in PYTHIA8 and EPOS LHC. In both models, the number of generated neutral and charged baryons was the same within statistical uncertainties, and same-baryon number correlations were consistent with each other for the two species. Therefore it was assumed that this is the case also in data. In EPOS, this assumption holds also for opposite-baryon number pairs, but this is not the case for PYTHIA, where the correlation with Ξ^0 is much weaker, as is shown in Fig. 8.15. Here, the $\Xi^- - \overline{\Xi}^0$ correlation is consistent with the $\Xi^- - \overline{\Xi}^+$ correlation at the away side, where it is dominated by the underlying event. The contribution from direct interactions with the trigger, which was estimated by subtracting the same-baryon number correlation (cf. Chapter 9 and 10), was on average 70% lower than for $\Xi^{\pm}$ and consistent with a flat ratio.

In data, most of the results fall somewhere in between these two models, cf. Chapter 9.

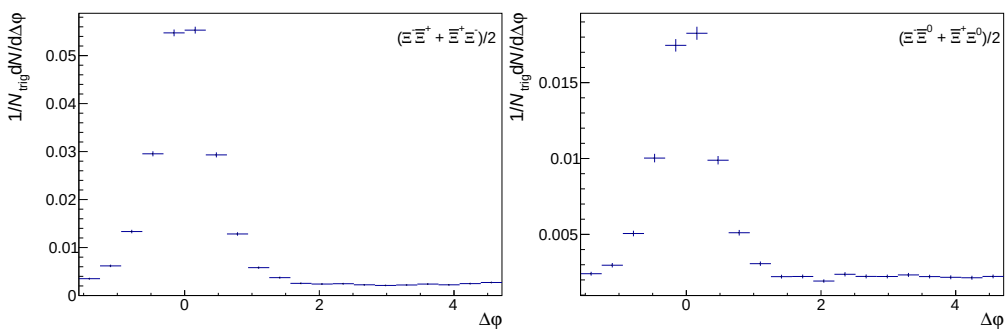

Figure 8.15 – Correlations between charged Ξ baryons and **(a)** other charged Ξ baryons and **(b)** neutral Ξ baryons of opposite baryon number, projected in $\Delta\varphi$, simulated by PYTHIA8, Monash tune.

Therefore, it was assumed that this is the case also here[9], and therefore the suppression of opposite-baryon number Ξ^0 was assumed to be 36% (except when evaluating the Monte Carlo closure test, cf. Section 8.9.1, where the PYTHIA value was used instead) in order to cover the two extremes within the systematic uncertainties. In practice, the feed-down contribution from Ξ baryons to Λ was measured as

$$N_{\Xi}^{\text{feeddown}} = \begin{cases} (f_{\Xi^-} + f_{\Xi^0})N_{\Xi}^{SB} & \text{same baryon number,} \\ (f_{\Xi^-} + (1-r)f_{\Xi^0})N_{\Xi}^{OB} + r f_{\Xi^0} N_{\Xi}^{SB} & \text{opposite baryon number,} \end{cases} \tag{8.15}$$

where N_{Ξ}^{SB} and N_{Ξ}^{OB} are the number of same- and opposite-baryon number detected $\Xi - \Xi$ pairs, respectively, f_{Ξ^-} and f_{Ξ^0} are the probabilities for the Λ daughter particle of a charged and neutral Ξ, respectively, to appear within the Λ sample, and r is the reduction factor for Ξ^0.

Due to the low efficiencies at low p_{T}, the lowest-p_{T} weights were merged (weighted with the spectra) into a single weight, which was applied to the lowest bin used for measuring the correlation function. Finally, the weighted correlation function was measured following the procedure in Section 8.3. The final correlation function was then convoluted by the autocorrelation function between mother and daughter, i.e. the expected smearing from the decay, which was simulated in MC but weighted with the differences in p_{T} spectra. Since this could only be done (at least without introducing any biases) to the final correlation function, the contribution from feed-down was not subtracted from the correlation function until the very end.

8.9 Monte Carlo Simulations

Apart from being used as a tool for e.g. efficiency and feed-down corrections, Monte Carlo simulations are important for validating the method and for testing model predictions. The procedures for this are described below.

8.9.1 Monte Carlo Closure Test

The validation of a method through MC simulations is known as a *Monte Carlo closure test*, and is simply a comparison between the results obtained by running the full method on the data set following the detector simulation, with the obtained result directly from the generator. The full reconstruction is obtained by using the same code as in data, whereas at generator level, the correlation function is measured from known particle species by only using the procedure described in Section 8.3 and the same kinematic cuts (in p_{T} and η) as for the reconstructed tracks. The ratios between reconstructed and generated MC, projected onto $\Delta\varphi$, are shown in Fig. 8.16 for $\Xi-\pi$ and $\Xi-$K correlations, Fig. 8.17 for Ξ − p and $\Xi - \Lambda$ correlations, and Fig. 8.18 for $\Xi - \Xi$ correlations. The agreement is generally very good, $< 2\%$ deviation for most measurements, although the

[9]From local conservation of quantum numbers alone, which is not implemented in EPOS LHC, one can expect that the correlation is stronger for charged Ξ baryons than for Ξ^0. The string breaking picture of PYTHIA seems to to some extent, but not fully, explain the correlations seen in data (cf. Chapter 10), meaning that the suppression of $\Xi^{\pm} - \Xi^0$ probably is substantial, but lower in data than in PYTHIA.

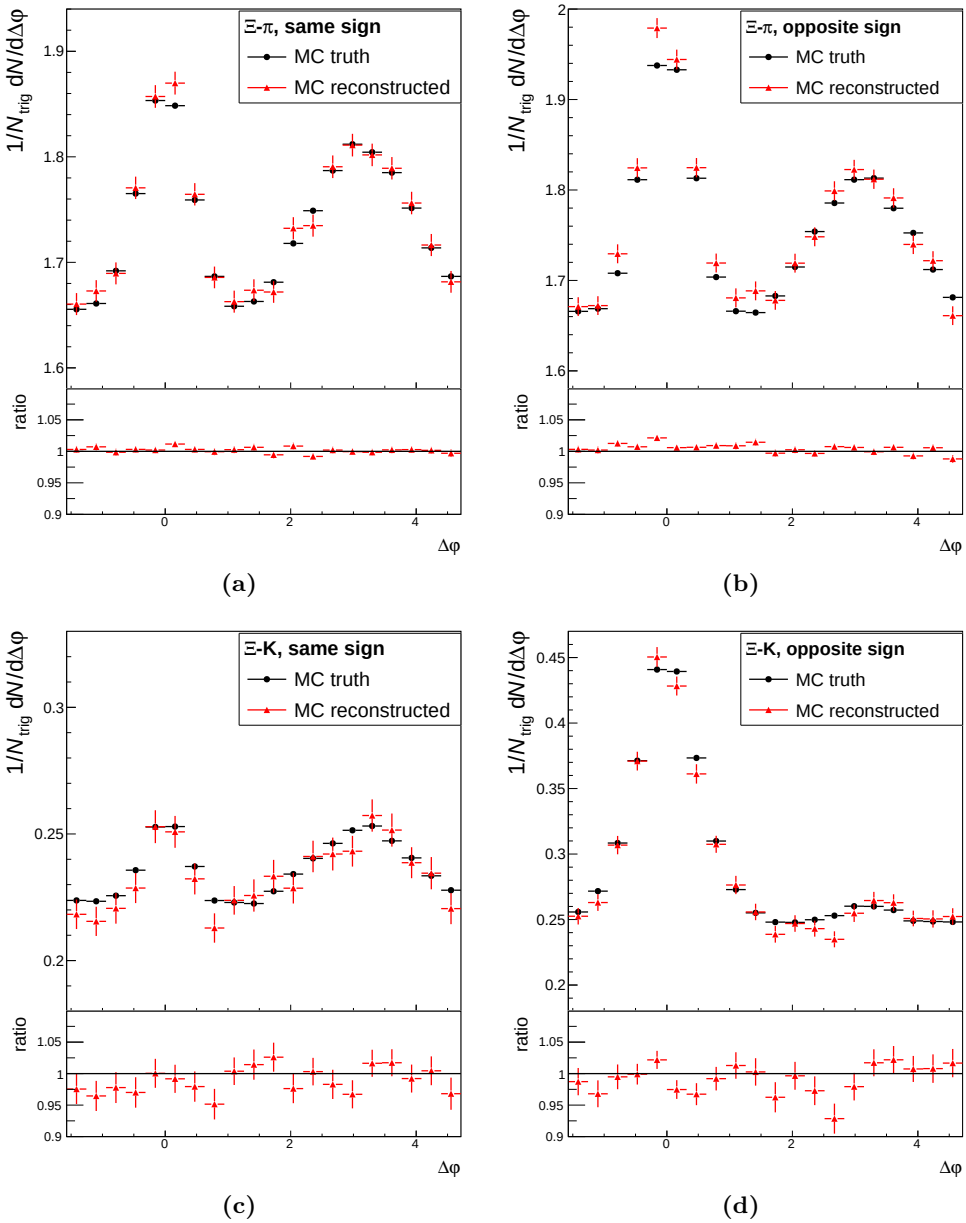

Figure 8.16 – Closure test for $\Xi - \pi$ and $\Xi - \mathrm{K}$ correlations, projected in $\Delta\varphi$ for **(a)** same-sign $\Xi - \pi$ correlations, **(b)** opposite-sign $\Xi - \pi$ correlations, **(c)** same-sign $\Xi - \mathrm{K}$ correlations, and **(d)** opposite-sign $\Xi - \mathrm{K}$ correlations. The lower panels show ratios between reconstructed and generated MC.

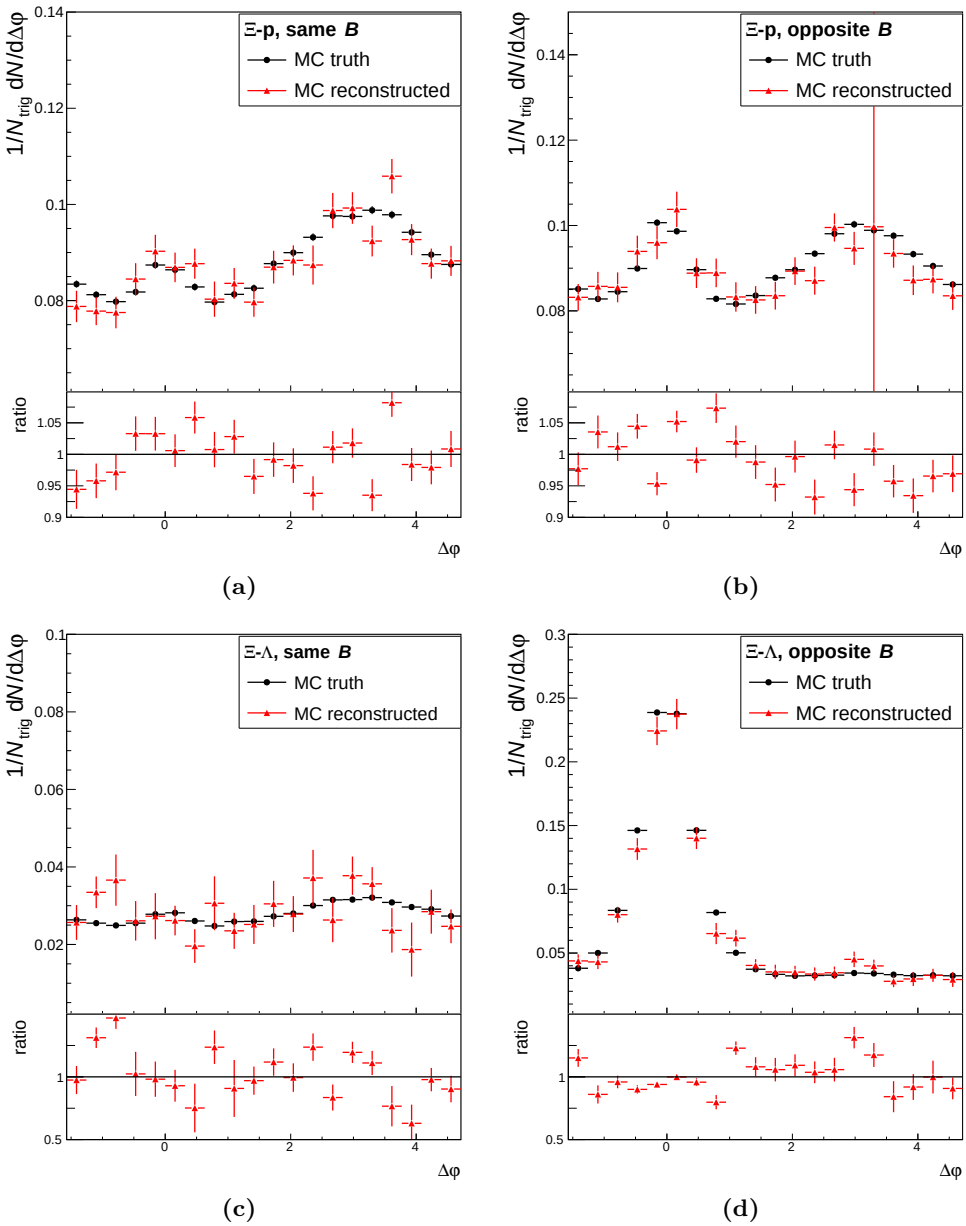

Figure 8.17 – Closure test for $\Xi - \mathrm{p}$ and $\Xi - \Lambda$ correlations, projected in $\Delta\varphi$ for **(a)** same-baryon number $\Xi - \mathrm{p}$ correlations, **(b)** opposite-baryon number $\Xi - \mathrm{p}$ correlations, **(c)** same-baryon number $\Xi - \Lambda$ correlations, and **(d)** opposite-baryon number $\Xi - \Lambda$ correlations. The lower panels show ratios between reconstructed and generated MC.

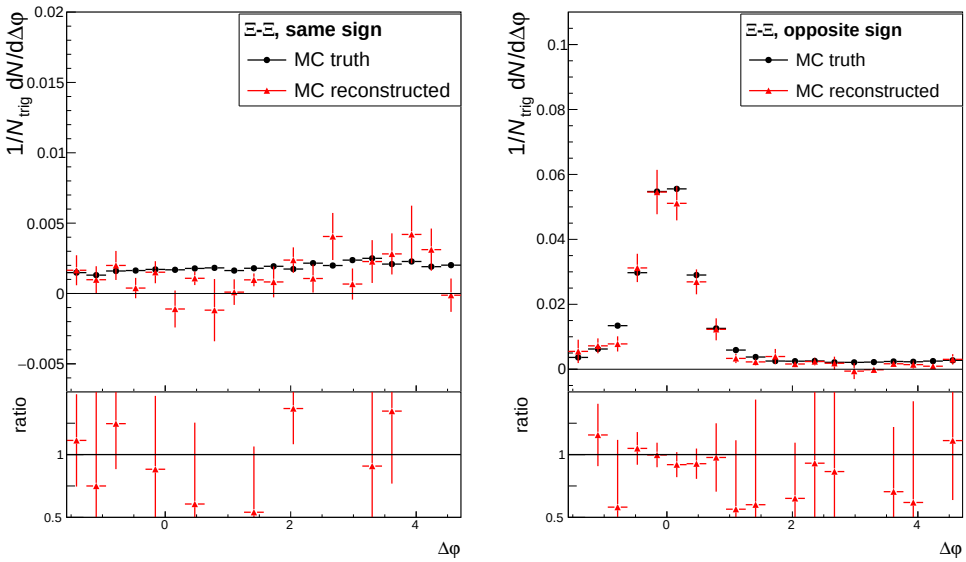

(a) **(b)**

Figure 8.18 – Closure test for $\Xi - \Xi$ correlations, projected in $\Delta\varphi$ for **(a)** same-baryon number correlations, and **(b)** opposite-baryon number correlations. The lower panels show ratios between reconstructed and generated MC.

statistical sample is a bit too small for in particular $\Xi - \Xi$ correlations, but to some extent also for $\Xi - \Lambda$ correlations, to be able to fully evaluate the uncertainty of the method.

Apart from a MC closure test of the correlations, it is important to validate the reconstruction methods for the particles, and in particular the tracks. This is best done by making a similar test for the spectra. For the pions and kaons, it has been confirmed that the agreement in MC is perfect, whereas a tiny deviation is seen for the protons due to uncertainties in the Λ spectrum. Therefore, a more conclusive way to test this is to compare the spectra with published ALICE results (which are not so reliant on the MC description as the approach used here), i.e. a data closure test rather than a MC closure test. This is done in Fig. 8.19, where it is shown that the agreement is good for pions, kaons, and protons for $p_{\mathrm{T}} \lesssim 3\,\mathrm{GeV}/c$, which is the onset of the relativistic rise in the TPC response (roughly corresponding to the δ term in Eq. (8.6)), which is likely not described well enough in the detector simulation for the method to work properly. This is the reason for the upper cut in p_{T} used in Table 8.1. For Ξ and Λ baryons, there is a slight disagreement, but considered good enough for our purposes[10].

[10]The data point with the largest disagreement in the Λ spectrum, corresponding to $0.5 < p_{\mathrm{T}} \leq 0.6\,\mathrm{GeV}/c$ is only used for feed-down subtraction. The discrepancy here is due to an improper description of the background, but the cause of the relatively large discrepancies for both Ξ and Λ at intermediate to high p_{T} is unknown.

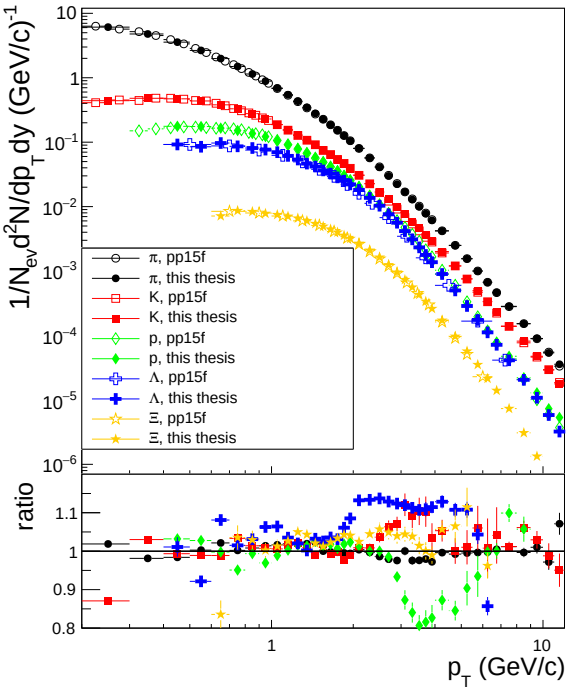

Figure 8.19 – Spectra of $\pi, \mathrm{K}, \mathrm{p}, \Lambda$, and Ξ as a function of p_{T}, compared with published results analysed on the pp15f data set [128]. Positive and negative particles have been combined to obtain these results. Note that the first data point for kaons in the lower panel deviates significantly from unity due to an interpolation error, and not to actual differences between the obtained spectra. This is also the reason for the seemingly abrupt dip in the Λ ratio at high p_{T}.

8.9.2 Model Comparisons

Model comparisons have been done to three different flavours of PYTHIA8 [74], namely the Monash tune (the same as used for the closure test), Junctions Mode 0 [87], and a yet unofficial tune with rope hadronisation. These models are all described in Section 3.8.2. Moreover a single implementation of EPOS, described in Section 3.8.4, has been tested, which is the now a bit dated event generator EPOS LHC [69]. The reason for using this, is that the newer generator, EPOS3, is not yet publicly available. The approach is the same as used for the reference in the MC closure test described in the previous section. A total number of 191 million events were used for PYTHIA8, 40 million events for the junction extension, 20 million for the rope extension, and 14 million for EPOS LHC. Without the efficiency losses associated with the reconstruction in the detector, this turned out to be well enough to measure the correlation function, although a larger sample would have been preferred for EPOS LHC. The results are presented in Chapter 9, along with comparisons to data.

8.10 Systematic Uncertainties

The uncertainty on any measured quantity can be divided into a statistical part and a systematic one. The statistical uncertainty is due to limitations in the statistical sample, and will get reduced by repeating the experiment. The systematic uncertainty on the other hand, is the uncertainty which will remain no matter how many times the experiment is repeated, i.e. this is the uncertainty imposed by the method itself. While there are well-established methods to calculate the statistical uncertainty, the same cannot be said about the systematic part. Nevertheless, it is at least as important to estimate this. The ideal way to estimate the systematic uncertainty is to repeat the measurement using a completely different experiment and also make sure to use different procedures. There exists however only one accelerator such as the LHC and one detector with the capabilities of ALICE, so in our field this is unfortunately impossible. Instead, the established procedure is to repeat the analysis with variations to all cuts and input parameters, and any variation which results in a deviation beyond the statistical uncertainty, is considered as a systematic uncertainty. Otherwise, the contribution from this particular variation is set to zero. Each variation should be set up in such a way that it is reasonable to assume that the variations are independent from each other, and the final uncertainty is obtained by summing them in quadrature. In this analysis, this is for instance assured by varying all trigger cuts simultaneously.

The following systematic variations were studied:

- Use a tighter cut set for the Ξ and Λ baryons, defined in Tables 8.5 and 8.6.
- Change some of the cuts related to the TPC clustering for the associated particles, summarised in Table 8.7.
- Use a tighter PID cut for the associated particles, $|n\sigma| < 3$.
- Use a fixed momentum value at $p = 0.6\,\mathrm{GeV}/c$ for all particle types as a starting point for using TOF information for the associated tracks.
- Use a more narrow vtx_z range, $|\mathrm{vtx}|_z < 8\,\mathrm{cm}$.
- Use a coarser momentum binning when calculating the misidentification fractions, 10 bins instead of 20 (each bin is merged into two).
- Use a tighter pileup cut for Ξ, requiring two daughter tracks with either a successful ITS refit or a TOF hit.
- Change the sideband definition to $4-7\sigma$ from the mean for Ξ and $5-8\sigma$ for Λ, instead of $3-6\sigma$ and $4-7\sigma$, respectively.
- Turn off the multiplicity weighting used when calculating efficiencies of associated particles and misidentification fractions.
- Calculate the coefficients used for weighting the correlation function used for feed-down corrections only as a function of p_{T} instead of in (p_{T}, η) space.

- When measuring the feed-down correction, scale the full correlation function instead of just the lowest-p_{T} bin when correcting for the low-p_{T} tail of the mother correlation, which cannot be incorporated for statistical reasons (motivated mostly since it smoothens the p_{T}-dependence and reduces the statistical error).
- Vary the reduction factor used for estimating the feed-down from Ξ^0 to Λ, as defined in Eq. (8.15), to $r = 0.72$. Since this is designed to be symmetric, one only needs to vary this in one direction.

Table 8.5 – Cuts used as a tight variation for the selection of Ξ candidates. Only cuts which are different from the ones listed in Table 8.3 are listed here. Throughout this table, p_{T} is in GeV/c.

Daughter track cuts	
V^0 pion daughter DCA	$\mathrm{DCA}_{\mathrm{d-PV}} > \max(0.05, -0.021 + 0.104\,(p_{\mathrm{T}} - 0.45)^{-0.7})\,\mathrm{cm}$
V^0 proton daughter DCA	$\mathrm{DCA}_{\mathrm{d-PV}} > \max(0.05, -2.22 + 2.34\,(p_{\mathrm{T}} - 0.45)^{-0.0165})\,\mathrm{cm}$
Bachelor DCA	$\mathrm{DCA}_{\mathrm{bach-PV}} > 0.037 + 0.033\,(p_{\mathrm{T}} - 0.45)^{-1.14}\,\mathrm{cm}$
V^0 cuts	
Invariant mass	$\lvert\Delta M^{\mathrm{inv}}\rvert < \min(8, 2.75 + 2.5\,p_{\mathrm{T}})\,\mathrm{MeV}/c^2$
DCA V^0-daughters	$\mathrm{DCA}_{\mathrm{d-d}} < \min(1.6, 0.2 + 1.1\,(p_{\mathrm{T}} - 0.45)^{0.6})\,\mathrm{cm}$
V^0 radius	$1.6 < r_{\mathrm{T}}^{\mathrm{V0}} < \max(10, 73\,p_{\mathrm{T}} - 34)\,\mathrm{cm}$
Cascade cuts	
Cascade DCA	$\mathrm{DCA}_{\mathrm{casc-PV}} < \min(2.0, 0.13 + 0.74\,(p_{\mathrm{T}} - 0.45)^{0.68})\,\mathrm{cm}$
Cascade radius	$0.6 + 0.16\,p_{\mathrm{T}}\,\mathrm{cm} < r_{\mathrm{T}}^{\mathrm{casc}} < 9 + 27\,(p_{\mathrm{T}} - 0.45)^{1.7}\,\mathrm{cm}$
Cosine of pointing angle	$\cos(PA) > \max(0.997, 0.9998 - 2.45 \cdot 10^{-3}(p_{\mathrm{T}} - 0.45)^{-1.1})$

Table 8.6 – Cuts used as a tight variation for the selection of Λ candidates. Only cuts which are different from the ones listed in Table 8.2 are listed here. For the p_{T}-dependent cut, p_{T} is in GeV/c.

Daughter track cuts	
Pion daughter DCA	$\mathrm{DCA}_{\mathrm{d-PV}} > 0.16\,\mathrm{cm}$
Proton daughter DCA	$\mathrm{DCA}_{\mathrm{d-PV}} > 0.07\,\mathrm{cm}$
V^0 cuts	
DCA V^0-daughters	$\mathrm{DCA}_{\mathrm{d-d}} < 0.6\,\mathrm{cm}$
V^0 radius	$0.58 + 0.23(p_{\mathrm{T}} - 0.35)^{1.25}) < r_{\mathrm{T}}^{\mathrm{V0}} < 10 + 52(p_{\mathrm{T}} - 0.35)^{1.5}\,\mathrm{cm}$
Cosine of pointing angle	$\cos(PA) > 0.998$

Table 8.7 – Cut variations used for the associated tracks as a systematic variation.

Cut	Default cuts	Systematic variation
Minimum number of TPC clusters	70	60
Maximum χ^2 per TPC cluster	4	5
Maximum DCA to z vertex	2 cm	3 cm

Additionally, the uncertainty on the fit value used for the event mixing normalisation (obtained from Eq. (8.5)) was assigned as a systematic uncertainty. The disagreement between reconstructed and generated MC from the Monte Carlo closure test was assigned as another source, since this should quantify the uncertainty imposed by the method itself. This will however not cover uncertainties from the propagation of MC to data, so in order to estimate this, the deviation between the measured spectra of the associated particles[11] and published results derived from Fig. 8.19 was used as an additional (flat) uncertainty. Moreover, as already has been mentioned, the detector simulation is not perfect, which adds an additional uncertainty from the tracking efficiencies which has been estimated as 0.7% for pions, 0.5% for kaons, and 1.5% for protons[12].

To quantify the effect of each systematic variation, fits were applied to the ratio between the variation and the default configuration. For projections in Δy, this ratio was always consistent with a constant, but for some variations there was a relatively large difference between near- and away-side correlations. Consequently, a fit of them form

$$A + B\cos(\Delta\varphi), \tag{8.16}$$

was applied to projections in $\Delta\varphi$, where A and B are fit parameters, since this follows the periodicity of the correlation function. For differences between opposite- and same-sign correlations, the corresponding fits were applied directly to the differences instead and not to the ratios. Moreover, since there for some of the correlations is a quite large difference between simulation output and data, a scaling was applied to the difference between opposite- and same-sign correlations when evaluating the impact from the MC closure test.

To check whether each change is statistically significant, the two correlation functions were assumed maximally correlated, i.e. the correlation coefficient is

$$\rho = \frac{\sigma_{\text{var}}}{\sigma_{\text{def}}},$$

where σ_{var} and σ_{def} are the standard deviations of the correlation function measured from the variation and default configuration, respectively. This correlation coefficient was used to assign the error at each data point in the ratio. Consequently, the error on the sum of fit parameters determined whether the systematic difference was significant or not. If the error was larger than the sum of parameters, the variation was rejected.

Summary of systematic uncertainties

Below, a summary of the statistically significant variations after applying the smoothening described above is presented. Summaries of the systematic uncertainties and the difference between opposite and same-sign correlations for the minimum-bias selection are shown in Fig. 8.20 for $\Xi - \pi$ correlations, Fig. 8.21 for $\Xi - \mathrm{K}$ and $\Xi - \mathrm{p}$ correlations, and Fig. 8.22 for $\Xi - \Lambda$ and $\Xi - \Xi$ correlations. Analogous plots for the high- and low-multiplicity sample are shown in Appendix D.2. Only results for projections in

[11] Unless the correlation function has a strong p_{T} dependence, which does not seem to be the case, any uncertainty in the trigger yields will be cancelled by the normalisation.

[12] Recommended values by the ALICE Data Preparation Group (DPG).

$\Delta\varphi$ are shown in the figures. The systematic uncertainties for projections in Δy were calculated separately. The full range of variations of projections in $\Delta\varphi$, near-side Δy, and away-side Δy is summarised in Table 8.8.

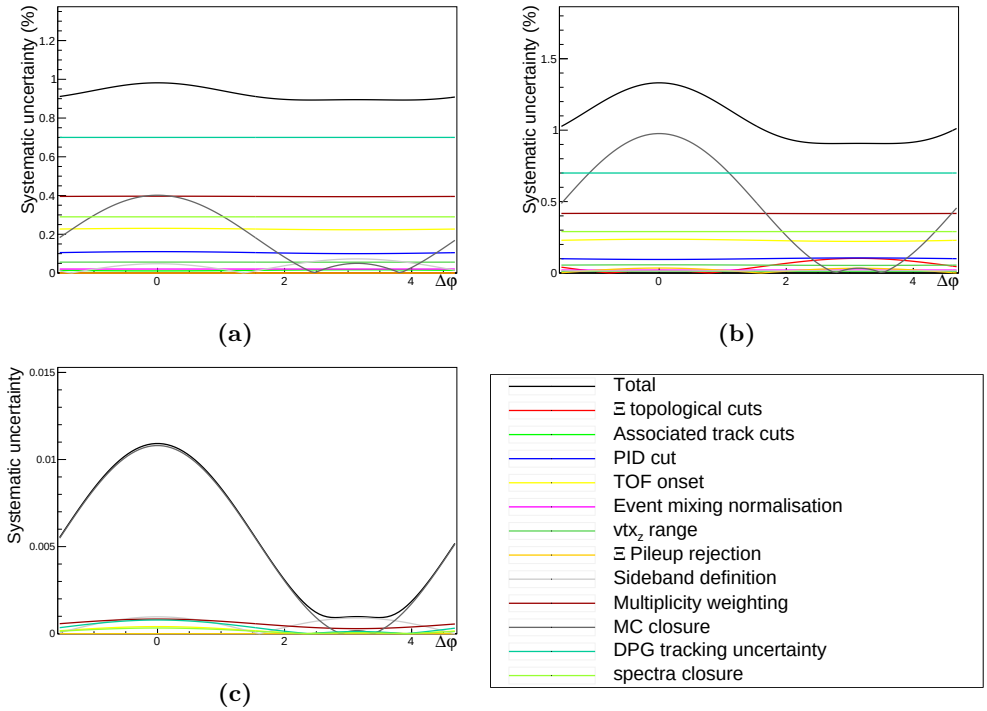

Figure 8.20 – Summary of systematic uncertainties from each source (insignificant sources are set to zero, but still included in the legend) for the projection in $\Delta\varphi$ of minimum-bias $\Xi - \pi$ correlations for **(a)** same-sign correlations, **(b)** opposite-sign correlations, and **(c)** differences between the two.

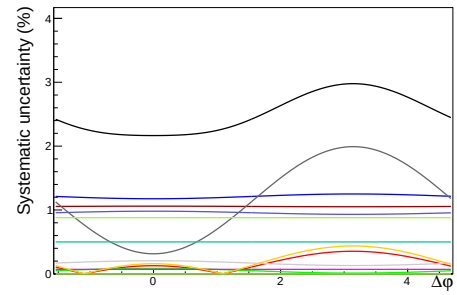

(a)

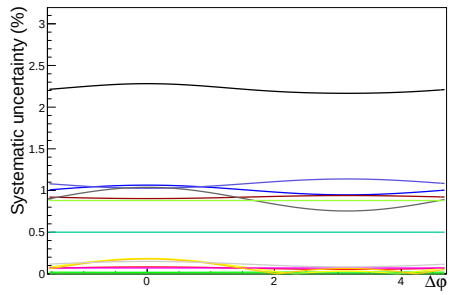

(b)

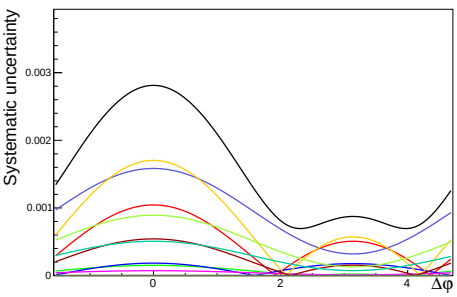

(c)

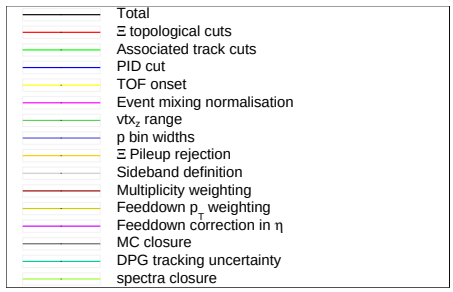

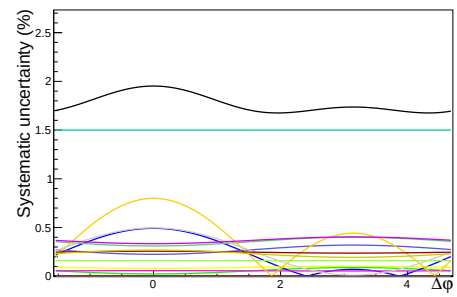

(d)

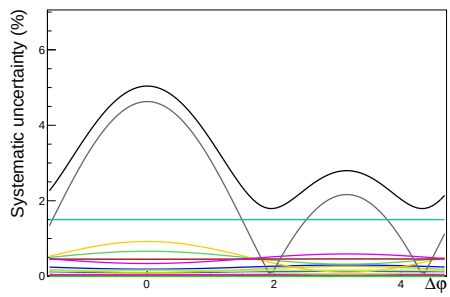

(e)

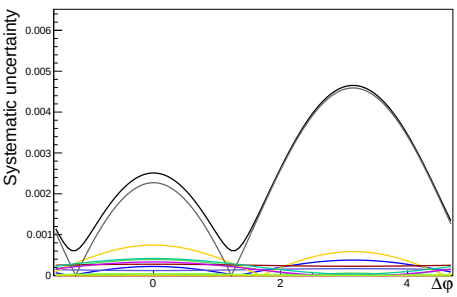

(f)

Figure 8.21 – Summary of systematic uncertainties from each source (insignificant sources are set to zero, but still included in the legend) for the projection in $\Delta\varphi$ in minimum-bias events, for **(a)** same-sign $\Xi - \mathrm{K}$ correlations, **(b)** opposite-sign $\Xi - \mathrm{K}$ correlations, **(c)** differences between the two, **(d)** same-baryon number $\Xi - \mathrm{p}$ correlations, **(e)** opposite-baryon number $\Xi - \mathrm{p}$ correlations, and **(f)** differences between the two.

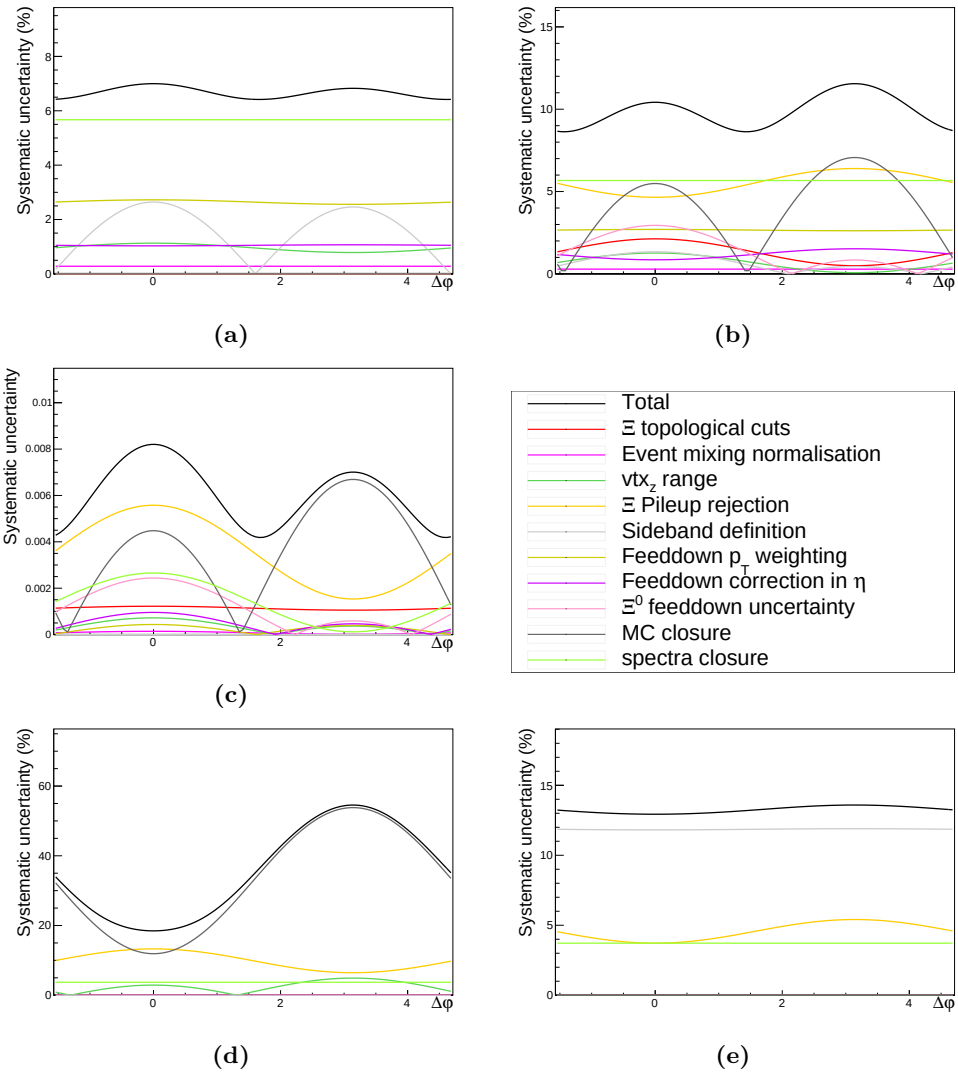

Figure 8.22 – Summary of systematic uncertainties from each source (insignificant sources are set to zero, but still included in the legend) for the projection in $\Delta\varphi$ in minimum-bias events, for **(a)** same-baryon number $\Xi - \Lambda$ correlations, **(b)** opposite-baryon number $\Xi - \Lambda$ correlations, **(c)** differences between the two, **(d)** same-baryon number $\Xi - \Xi$ correlations, **(e)** opposite-baryon number $\Xi - \Xi$ correlations, and **(f)** differences between the two.

Table 8.8 – Summary of all systematic uncertainties. The ranges are the minimum and maximum values of each quantity across all multiplicity configurations. For the opposite-same sign differences, absolute uncertainties are given. Statistically insignificant effects are denoted by "neg.". Legend: SS = same sign, OS = opposite sign, SB = same baryon number, OB = opposite baryon number.

Systematic effect	SS π	OS π	SS K	OS K	SB p
Ξ/Λ topological cuts	$0-0.6\%$	$0-0.4\%$	$0-1.2\%$	$0-0.8\%$	$0-5.9\%$
$\pi, \mathrm{K}, \mathrm{p}$ ITS-TPC cuts	$< 0.1\%$	$< 0.1\%$	$0-0.1\%$	$0-0.2\%$	$0-0.2\%$
PID cuts	$0-0.3\%$	$0.1-0.3\%$	$0.8-1.6\%$	$0.7-1.3\%$	$0-0.5\%$
TOF onset	$0.1-0.2\%$	$0.2-0.3\%$	$0-0.2\%$	$0-0.4\%$	$0-0.7\%$
Event mixing normalisation	$< 0.1\%$	$< 0.1\%$	$0.1-0.3\%$	$0.1-0.3\%$	$0-0.2\%$
vtx_z range	$0-0.2\%$	$0-0.2\%$	$0-1.5\%$	$0-0.3\%$	$0-1.1\%$
Momentum binning	[a]	[a]	$0.3-1.1\%$	$0.8-1.1\%$	$0.2-0.4\%$
Pileup rejection	$0-0.3\%$	$0-0.4\%$	$0-1.5\%$	$0-0.2\%$	$0-3.5\%$
Sideband region	$0-0.3\%$	$0-0.4\%$	$0-0.9\%$	$0-0.7\%$	$0-1.3\%$
Multiplicity weighting	$0.2-1.3\%$	$0.3-1.3\%$	$0.7-2.8\%$	$0.6-2.8\%$	$0-0.4\%$
Feed-down p_{T} weighting				$0-0.6\%$	
1D feed-down correction				$0-0.4\%$	
Ξ^0 uncertainty					
MC closure	$0-0.4\%$	$0-1.4\%$	$0-3.1\%$	$0.8-2.1\%$	$0-1.5\%$
Spectra closure	0.3%	0.3%	0.9%	0.9%	0.2%
DPG efficiency uncertainty	0.7%	0.7%	0.5%	0.5%	1.5%
Total	$0.9-1.6\%$	$0.9-2.1\%$	$1.9-4.6\%$	$2.1-3.8\%$	$1.7-6.4\%$

Systematic effect	OB p	SB Λ	OB Λ	SB Ξ	OB Ξ
Ξ/Λ topological cuts	$0-1.0\%$	$0-24\%$	$0-6.7\%$	$0-120\%$	$0-21\%$
$\pi, \mathrm{K}, \mathrm{p}$ ITS-TPC cuts	$0-0.2\%$				
PID cuts	$0-0.5\%$				
TOF onset	$0-0.4\%$				
Event mixing normalisation	$0-0.2\%$	$0-1.1\%$	$0.3-1.2\%$	neg.	neg.
vtx_z range	$0-1.4\%$	$0-1.9\%$	$0-2.2\%$	$0-4.9\%$	$0-3.5\%$
Momentum binning	$0-0.2\%$ [b]				
Pileup rejection	$0.1-1.5\%$	$0-7.8\%$	$2.9-10.6\%$	$0-51\%$	$0-20\%$
Sideband region	$0-1.7\%$	$0-5.9\%$	$0-4.4\%$	$0-17\%$	$0-12\%$
Multiplicity weighting	$0.1-0.9\%$				

[a] Rejected due to a bug in the code.
[b] Rejected for the low-multiplicity configuration due to a bug in the code.

Systematic effect	OB p	SB Λ	OB Λ	SB Ξ	OB Ξ
Feed-down p_{T} weighting	$0-1.2\%$	$1.4-4.9\%$	$1.1-6.3\%$		
1D feed-down correction	$0.3-1.0\%$	$0-1.1\%$	$0-1.7\%$		
Ξ^0 uncertainty			$0-4.6\%$		
MC closure	$0-4.6\%$	neg.	$0-7.1\%$	$12-78\%$	$0-60\%$
Spectra closure	0.2%	5.7%	5.7%	3.7%	3.7%
DPG efficiency uncertainty	1.5%				
Total	$1.8-5.5\%$	$6.4-26\%$	$7.2-15\%$	$14-130\%$	$6.5-67\%$

Systematic effect	OS-SS π ($\times10^{-3}$)	OS-SS K ($\times10^{-3}$)	OB-SB p ($\times10^{-3}$)	OB-SB Λ ($\times10^{-3}$)	OB-SB Ξ ($\times10^{-3}$)
Ξ/Λ topological cuts	$0-3.7$	$0-5.3$	$0-2.9$	$0-3.8$	$0-2.8$
$\pi, \mathrm{K}, \mathrm{p}$ ITS-TPC cuts	$0-1.6$	<0.5	$0-1.0$		
PID cuts	$0-2.9$	$0-0.9$	$0-0.6$		
TOF onset	$0-1.7$	$0-2.1$	<0.5		
Event mixing normalisation	neg.	<0.5	neg.	$0-0.5$	neg.
vtx_z range	$0-1.9$	$0-1.0$	$0-0.9$	$0-2.2$	$0-0.5$
Momentum binning	[a]	$0-1.5$	<0.5[b]		
Pileup rejection	$0-4.6$	$0-2.8$	$0-2.9$	$0-8.9$	$0-3.1$
Sideband region	$0-3.7$	$0-2.1$	$0-1.8$	$0-7.4$	$0-2.0$
Multiplicity weighting	$0-1.3$	$0-2.7$	$0-0.5$		
Feed-down p_{T} weighting			<0.5	$0-3.2$	
1D feed-down correction			<0.5	$0-1.3$	
Ξ^0 uncertainty				$0-4.6$	
MC closure	$0-32$	neg.	$0-7.1$	$0-7.0$	$0-4.8$
Spectra closure	<0.5	$0.1-1.0$	<0.5	$0.1-3.3$	$0-0.4$
DPG efficiency uncertainty	$0-1.1$	$0.1-0.6$	$0-0.5$		
Total	$0.9-33$	$0.5-5.7$	$0.6-7.9$	$3.7-10$	$0.1-5.1$

[a] Rejected due to a bug in the code.
[b] Rejected for the low-multiplicity configuration due to a bug in the code.

Chapter 9

Correlation Results

9.1 ALICE Results

Angular $\Xi - \pi$ and $\Xi - \mathrm{K}$ correlations in $(\Delta y, \Delta\varphi)$ space for minimum-bias events

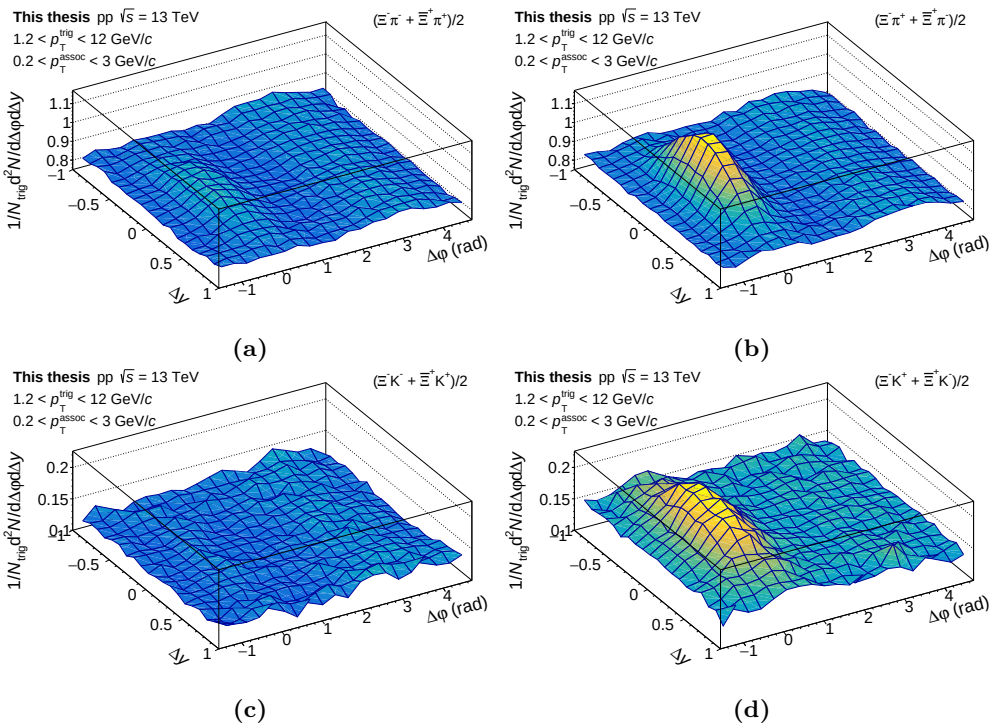

Figure 9.1 – Correlations in $(\Delta y, \Delta\varphi)$ space for minimum-bias events between Ξ baryons and **(a)** same-sign pions, **(b)** opposite-sign pions, **(c)** same-sign kaons, and **(d)** opposite-sign kaons.

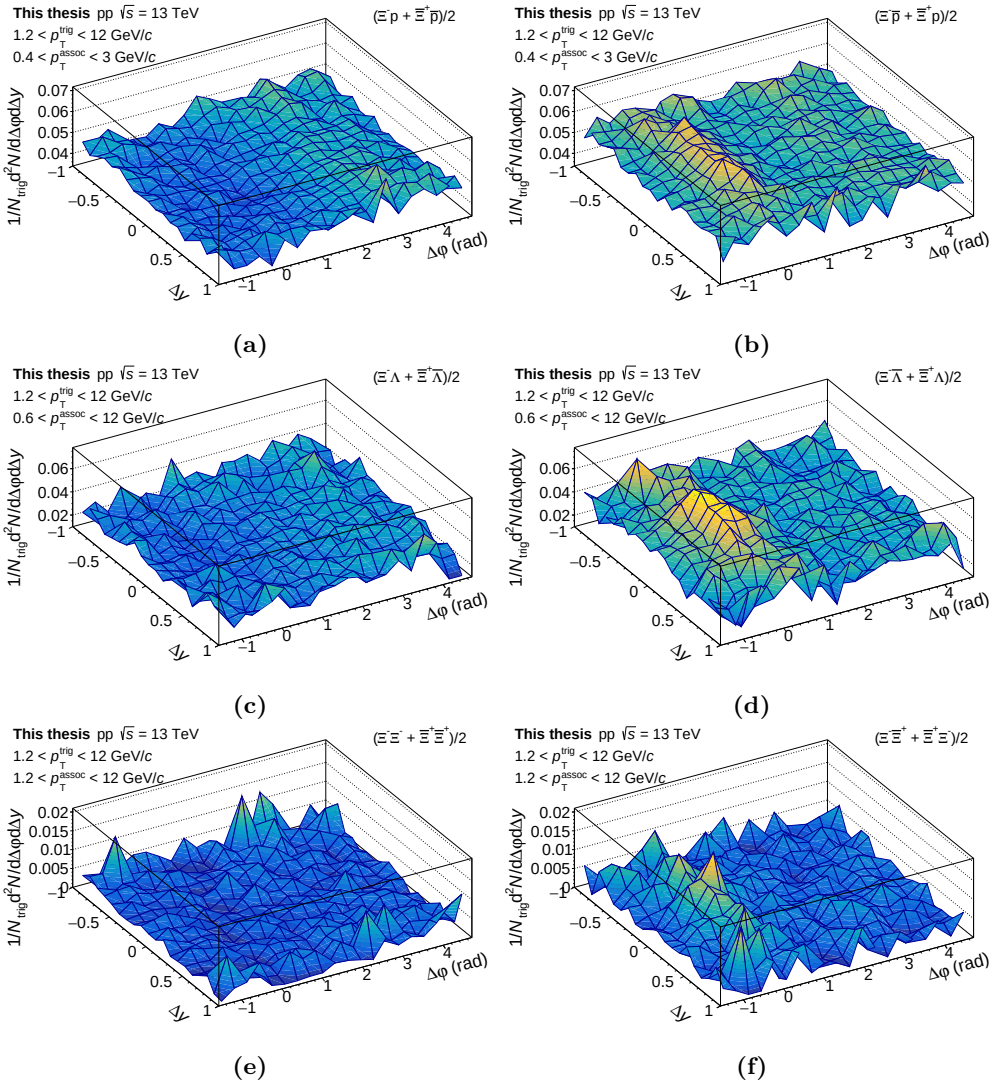

Figure 9.2 – Correlations in $(\Delta y, \Delta\varphi)$ space between Ξ baryons and **(a)** same-B (baryon number) protons, **(b)** opposite-B protons, **(c)** same-B Λ baryons, **(d)** opposite-B Λ baryons, **(e)** same-B Ξ baryons, and **(f)** opposite-B Ξ baryons.

are shown in Fig. 9.1 and $\Xi - \mathrm{p}$, $\Xi - \Lambda$, and $\Xi - \Xi$ correlations are shown in Fig. 9.2. The near side ($|\Delta\varphi| < \pi/2$) of the correlation function is shown on the left side of each figure and the away side on the right side. Common features are a near-side peak and an away-side ridge, which extends far in Δy. The near-side peak is a signature of strong correlations between trigger and associated particles produced close to each other in phase space, which for $\Xi-$meson pairs is mainly present for oppositely charged particles. For $\Xi - \pi$ and $\Xi - \mathrm{K}$ correlations, this effect is also present for same-charge

pairs, but is much weaker (note that for each particle species, the scales are the same for the left and right panels of Figs. 9.1 and 9.2. For different particle species, different scales are used). If this is the case also for $\Xi - \mathrm{K}$ correlations, it is not statistically significant given the limited sample.

For $\Xi-$baryon correlations, the near-side peak is instead only present for pairs of particles with opposite baryon number, indicating that baryon number affects the correlation function more than charge. For pairs between particles with the same baryon number, there is rather an anti-correlation on the near side, which is present even for protons, which have opposite charge. This means that production of other baryons of the same baryon number is disfavoured near a Ξ baryon.

Due to momentum conservation, the near-side peak must be balanced by an away-side ridge for the *total* particle yield. Its presence in all $\Xi - h$ correlation functions, regardless of charge, indicates that this is not very much dependent on the particle species, but rather a global effect due to the underlying event. Similarly, one can expect that the underlying event will affect also the near-side correlations, e.g. from elliptic flow (since v_2 describes expansion of the medium both on the near and away side, cf. Section 3.7, this shows up as a near-side ridge) and various jet correlations. Since the Ξ baryons are expected to correlate with particles of opposite quantum numbers, same-sign correlations are a signature of the underlying event. This also explains the relative flatness of same-charge or baryon number correlations. Therefore, a way to quantify the effect the Ξ has on particle production is to subtract same-quantum number correlations from opposite-quantum number correlations. This is only shown for projections (see Section 9.1.1) and not the full correlation function.

The major difference between $\Xi - \pi$ and $\Xi - \mathrm{K}$ correlations is that the near-side peak in $\Xi - \mathrm{K}$ correlations is wider and stronger relative to the background. This is likely related to strangeness production, and will be discussed in Chapter 10. A similar difference in peak amplitude is observed when comparing $\Xi - \mathrm{p}$ correlations with $\Xi - \Lambda$ correlations, although both of these correlations have a wide near-side peak. For $\Xi - \Xi$ correlations, the statistical errors are a bit too large to be able to fully describe the shape of the near-side peak, but it seems to be similar to the other $\Xi-$baryon correlation functions.

9.1.1 Multiplicity Dependent Results

The shapes of the correlation functions measured in the high- and low-multiplicity samples are all qualitatively similar to the minimum-bias results, and thus these results are shown in Appendix D.3. To quantify the differences between each multiplicity selection, projections of all these correlations functions in $\Delta\varphi$ and Δy, where the latter have been separated in near-side correlations ($|\Delta\varphi| < 3\pi/10$, which is chosen to roughly cover the peak) and away-side correlations ($|\Delta\varphi - \pi| < \pi/2$). Moreover, differences between same- and opposite-quantum number (charge or baryon number) correlations are included. The results are shown in Figs. 9.3, 9.4, 9.5, 9.6, and 9.7 for correlations between Ξ baryons and pions, kaons, protons, Λ baryons, and other Ξ baryons, respectively. The observant reader may notice that for $\Xi - \Xi$ correlations, a few of the data points have negative values, but these are just artefacts of the background subtraction from a very limited statistical sample, and do not make physical sense.

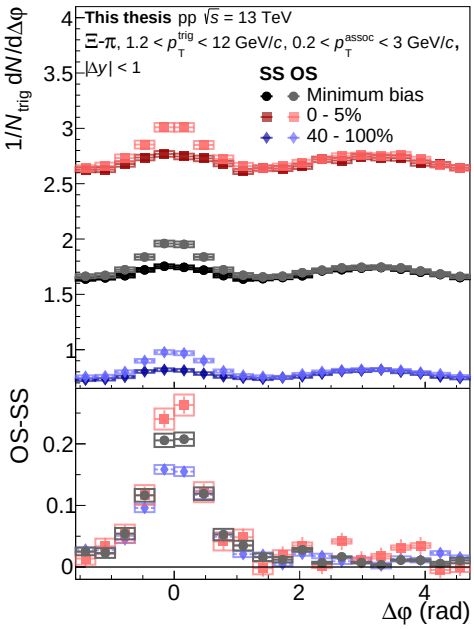

(a)

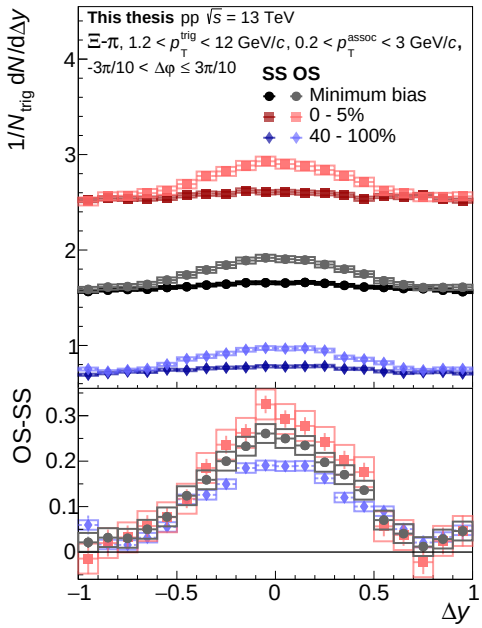

(b)

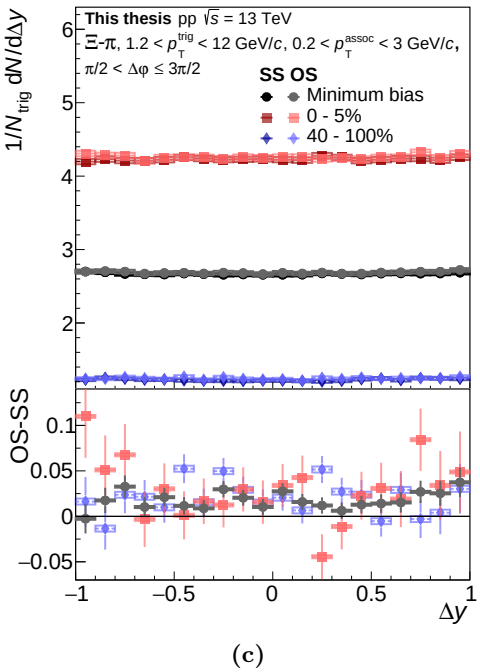

(c)

Figure 9.3 – Correlations between Ξ baryons and pions for different multiplicity classes. The bottom panels show differences between opposite- and same-sign correlations for each selection. Statistical and systematic errors are marked with bars and boxes, respectively. **(a)** Projection in $\Delta\varphi$, **(b)** projection in Δy on the near side ($-3\pi/10 < \Delta\varphi \leq 3\pi/10$), **(c)** projection in Δy on the away side ($\pi/2 < \Delta\varphi \leq 3\pi/2$). Legend: SS = same sign, OS = opposite sign.

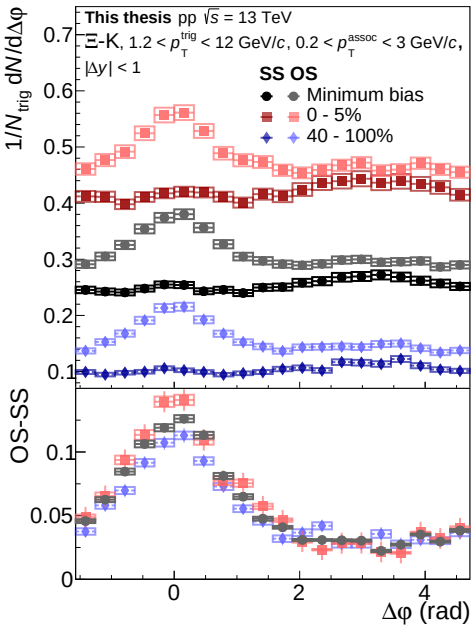

(a)

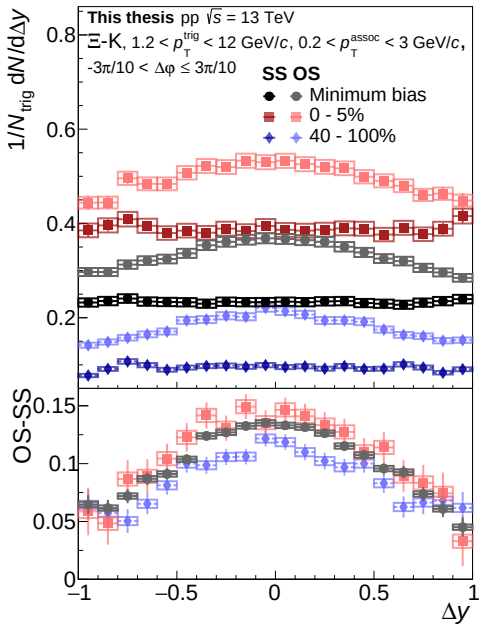

(b)

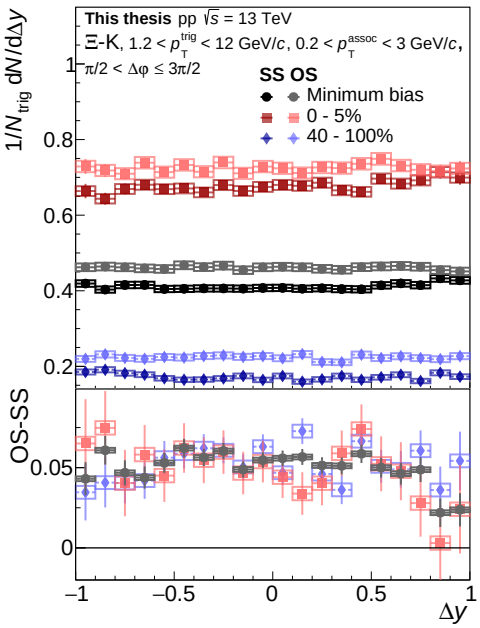

(c)

Figure 9.4 – Correlations between Ξ baryons and kaons for different multiplicity classes. The bottom panels show differences between opposite- and same-sign correlations for each selection. Statistical and systematic errors are marked with bars and boxes, respectively. **(a)** Projection in $\Delta\varphi$, **(b)** projection in Δy on the near side ($-3\pi/10 < \Delta\varphi \leq 3\pi/10$), **(c)** projection in Δy on the away side ($\pi/2 < \Delta\varphi \leq 3\pi/2$). Legend: SS = same sign, OS = opposite sign.

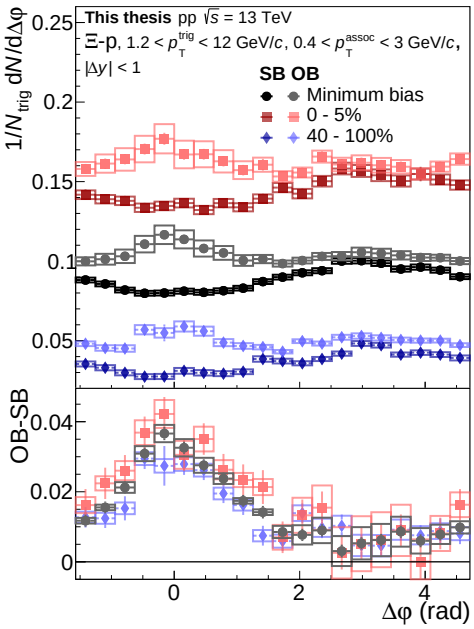

(a)

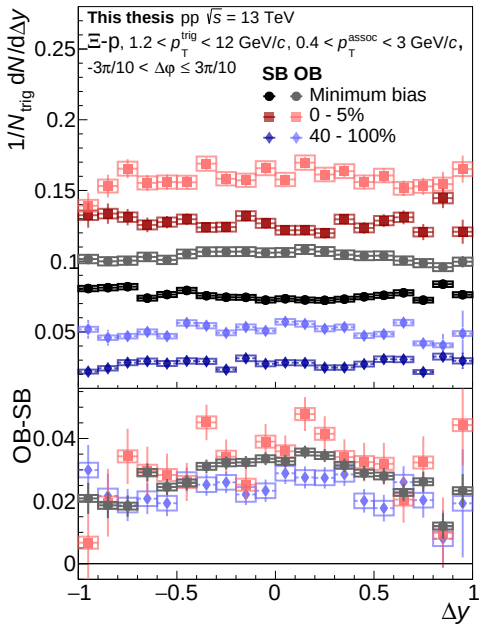

(b)

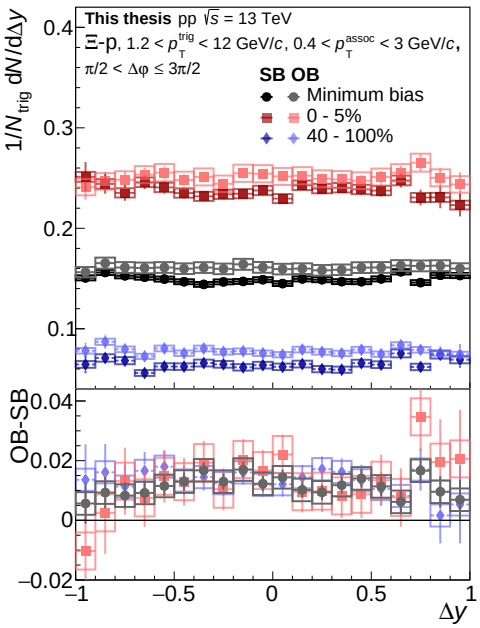

(c)

Figure 9.5 – Correlations between Ξ baryons and protons for different multiplicity classes. The bottom panels show differences between opposite- and same-sign correlations for each selection. Statistical and systematic errors are marked with bars and boxes, respectively. **(a)** Projection in $\Delta\varphi$, **(b)** projection in Δy on the near side ($-3\pi/10 < \Delta\varphi \leq 3\pi/10$), **(c)** projection in Δy on the away side ($\pi/2 < \Delta\varphi \leq 3\pi/2$). Legend: SB = same baryon number, OB = opposite baryon number.

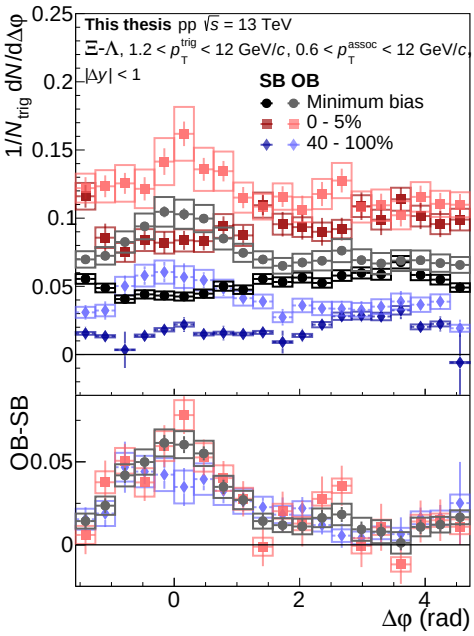

(a)

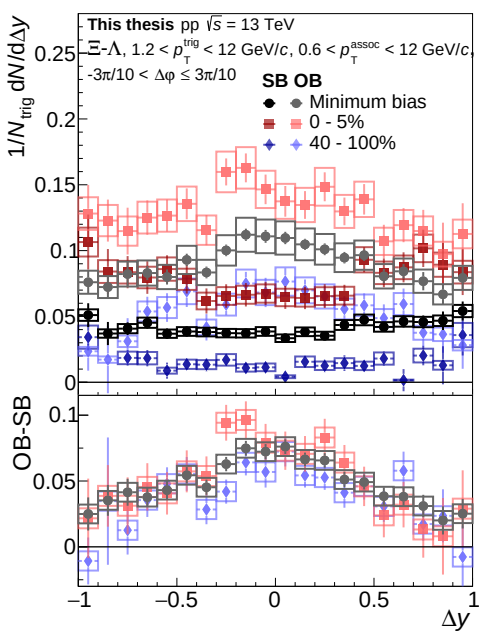

(b)

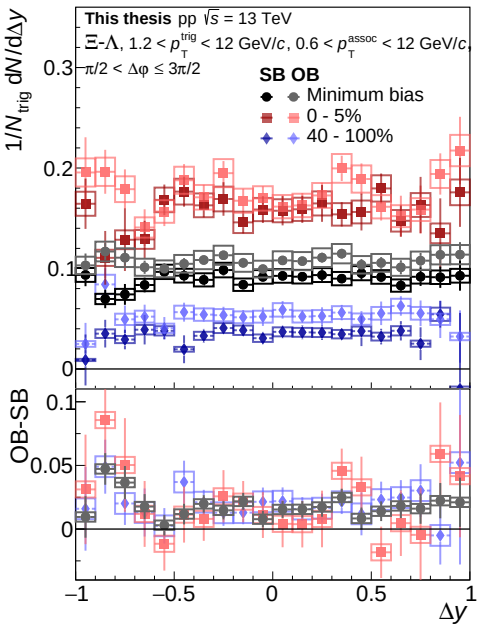

(c)

Figure 9.6 – Correlations between Ξ and Λ baryons for different multiplicity classes. The bottom panels show differences between opposite- and same-sign correlations for each selection. Statistical and systematic errors are marked with bars and boxes, respectively. **(a)** Projection in $\Delta\varphi$, **(b)** projection in Δy on the near side ($-3\pi/10 < \Delta\varphi \leq 3\pi/10$), **(c)** projection in Δy on the away side ($\pi/2 < \Delta\varphi \leq 3\pi/2$). Legend: SB = same baryon number, OB = opposite baryon number.

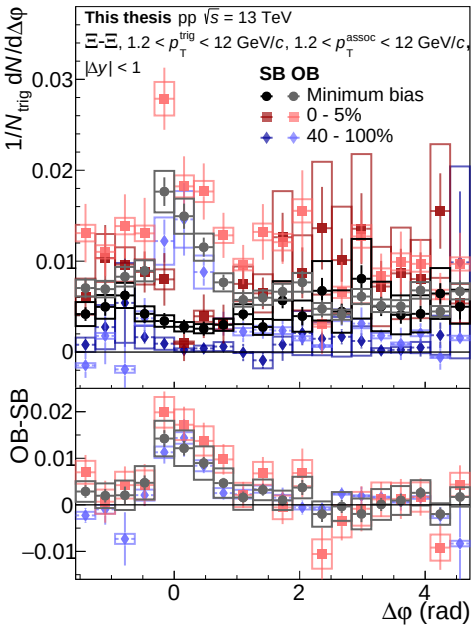

(a)

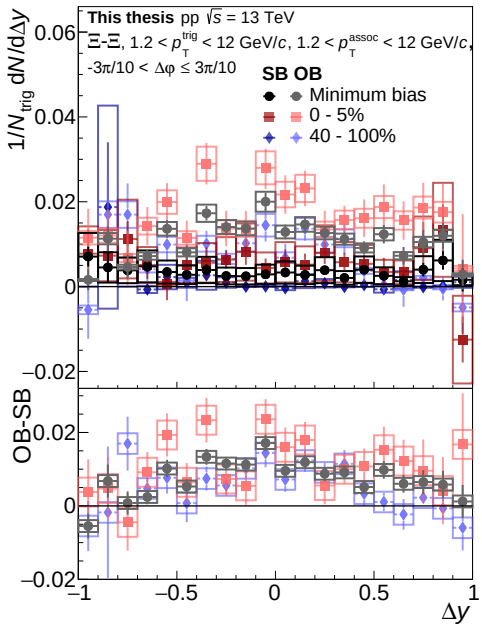

(b)

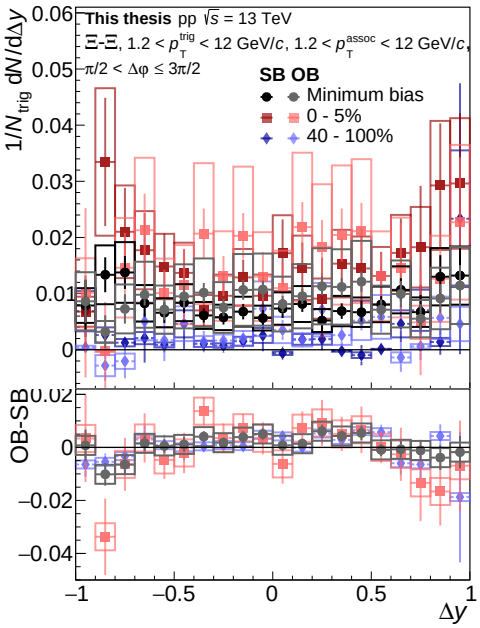

(c)

Figure 9.7 – Correlations between pairs of Ξ baryons for different multiplicity classes. The bottom panels show differences between opposite- and same-sign correlations for each selection. Statistical and systematic errors are marked with bars and boxes, respectively. **(a)** Projection in $\Delta\varphi$, **(b)** projection in Δy on the near side ($-3\pi/10 < \Delta\varphi \leq 3\pi/10$), **(c)** projection in Δy on the away side ($\pi/2 < \Delta\varphi \leq 3\pi/2$). Legend: SB = same baryon number, OB = opposite baryon number.

For all correlation functions, the results are very similar between multiplicity classes. The overall baseline increases thanks to the increasing multiplicity, and the relative difference between opposite- and same-sign correlations increases, which is a consequence of that the absolute difference does not change very much. There is a slight increase in peak amplitude when going to higher multiplicities, which may be compensated for by a possible narrowing in Δy, but it is difficult to tell given the large statistical errors at large $|\Delta y|$.

9.2 Simulation Results

Simulations have been carried out using the PYTHIA8 [74] and EPOS LHC [69] event generators. For PYTHIA8, the standard tune (Monash; cf. Section 3.8.2) is complemented with the junction and rope extensions. For junctions, the official tune Mode 0 is being used [87], whereas an unofficial tune largely based on the same parameters as Mode 0 is being used for ropes[1]. Angular $\Xi - \pi$, $\Xi - \mathrm{K}$, $\Xi - \mathrm{p}$, $\Xi - \Lambda$, and $\Xi - \Xi$ correlations, simulated by all of these event generators, are given in Figs. 9.8, 9.9, 9.10, 9.11, and 9.12, respectively.

For $\Xi - \pi$ correlations, there are not any major differences between PYTHIA and EPOS, nor between the different flavours of PYTHIA, at least qualitatively. There is a near-side peak present in all models for all charge combinations, but it is stronger for opposite-sign correlations, just as in data. The biggest difference is that the near-side peak is stronger and narrower in PYTHIA than in EPOS for same-charge correlations.

For $\Xi - \mathrm{K}$ correlations, the situation is different, with all PYTHIA flavours giving very similar results with a sharp near-side peak for opposite-sign correlations, which is both higher and narrower than in data, and a weak peak for same-sign correlations. EPOS gives a very different picture with a very wide near-side peak and strong ridges both on the near and away side. The magnitude indicates a large probability of producing kaons together with Ξ baryons, but unlike pions, they are very weakly correlated in phase space. Moreover, there are strong correlations far away in Δy, which are neither seen in PYTHIA nor in data.

For PYTHIA, $\Xi - \mathrm{p}$ correlations are qualitatively quite similar to what is observed in data, with a near-side peak for opposite-baryon number correlations and a tendency of anti-correlations between pairs of the same baryon number. The magnitude is quite different between different PYTHIA flavours, on the other hand, with a much stronger near-side peak in the junction and rope models, as well as stronger correlations overall. For EPOS, the difference to PYTHIA is smaller than for $\Xi - \mathrm{K}$ correlations, but again the near-side peak is wider and weaker, and there are strong correlation far away in rapidity. Moreover, for same-baryon number correlations, there is a tendency of anti-correlations on the *away* side, rather than on the near side.

For $\Xi - \Lambda$ and $\Xi - \Xi$ correlations, there is a very strong near-side peak in PYTHIA which dominates the entire correlation function, although the magnitude is slightly smaller for the rope model and even more so for the junction model. This peak is much

[1]The rope model includes junction interactions as implemented by the junction extension, but further includes more interactions from rope hadronisation. The tune that is used was provided by PYTHIA developer Christian Bierlich.

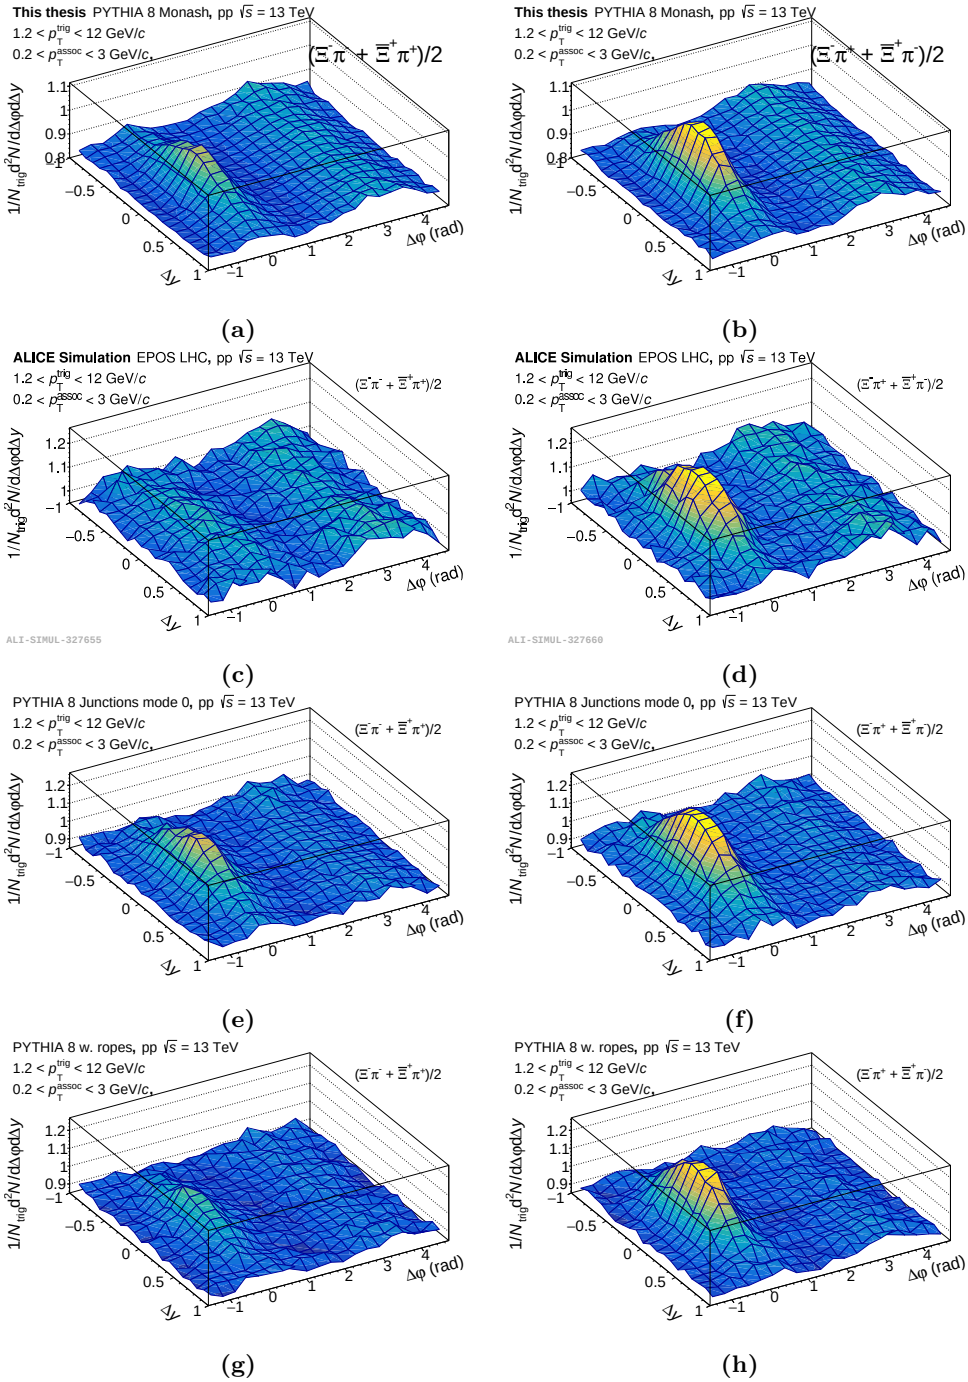

Figure 9.8 – Simulated 2D-correlations between Ξ baryons and pions, left: same-sign correlations, right: opposite-sign correlations. **(a)**-**(b)** PYTHIA8, **(c)**-**(d)** EPOS LHC, **(e)**-**(f)** PYTHIA8 with junctions, and **(g)**-**(h)** PYTHIA8 with ropes.

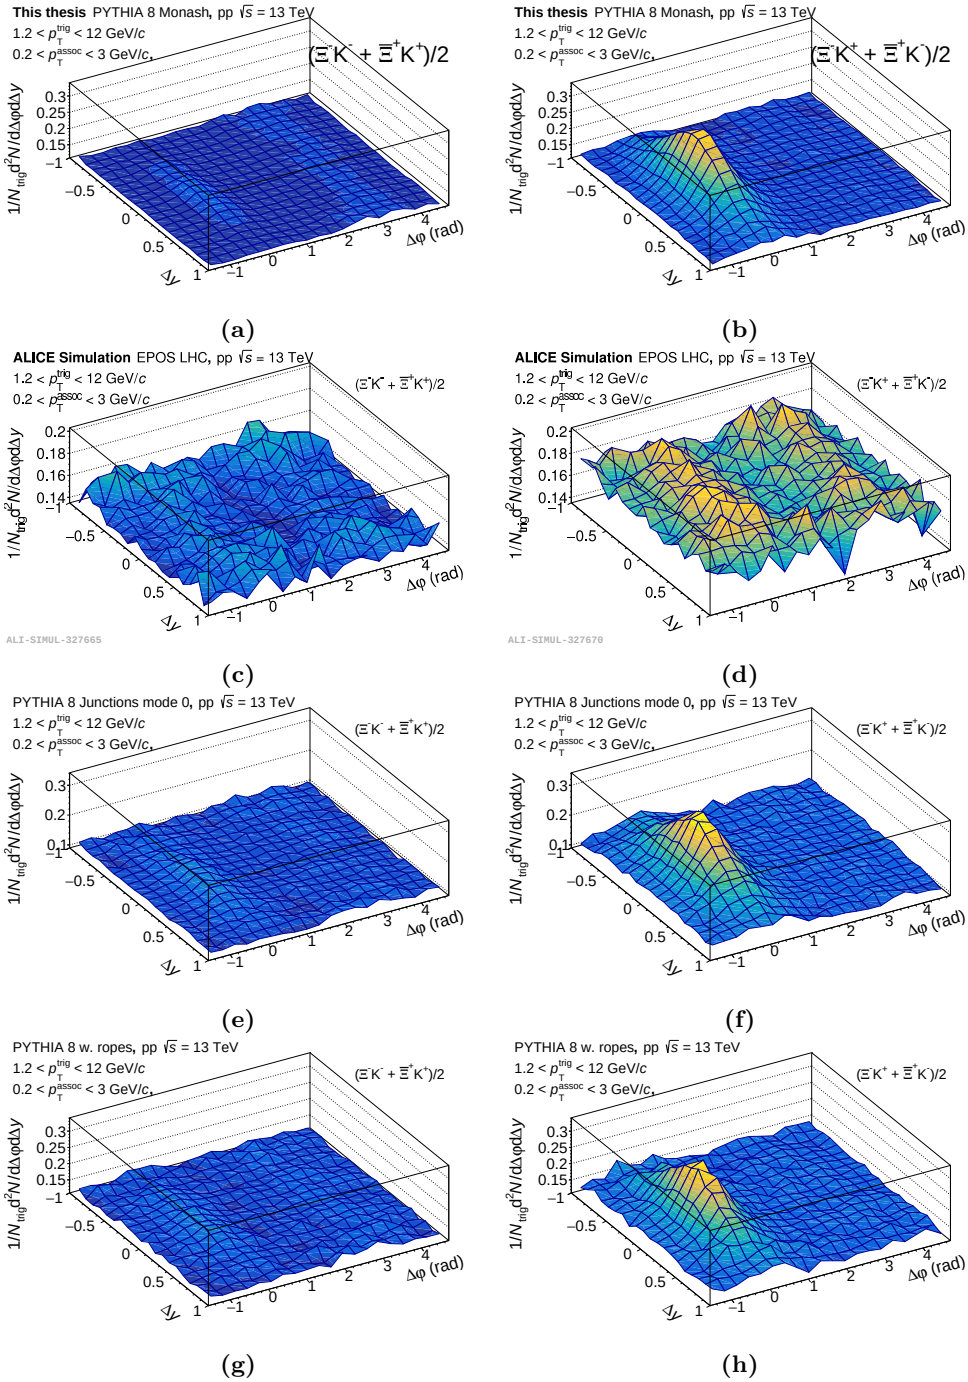

Figure 9.9 – Simulated 2D-correlations between Ξ baryons and kaons, left: same-sign correlations, right: opposite-sign correlations. **(a)**-**(b)** PYTHIA8, **(c)**-**(d)** EPOS LHC, **(e)**-**(f)** PYTHIA8 with junctions, and **(g)**-**(h)** PYTHIA8 with ropes.

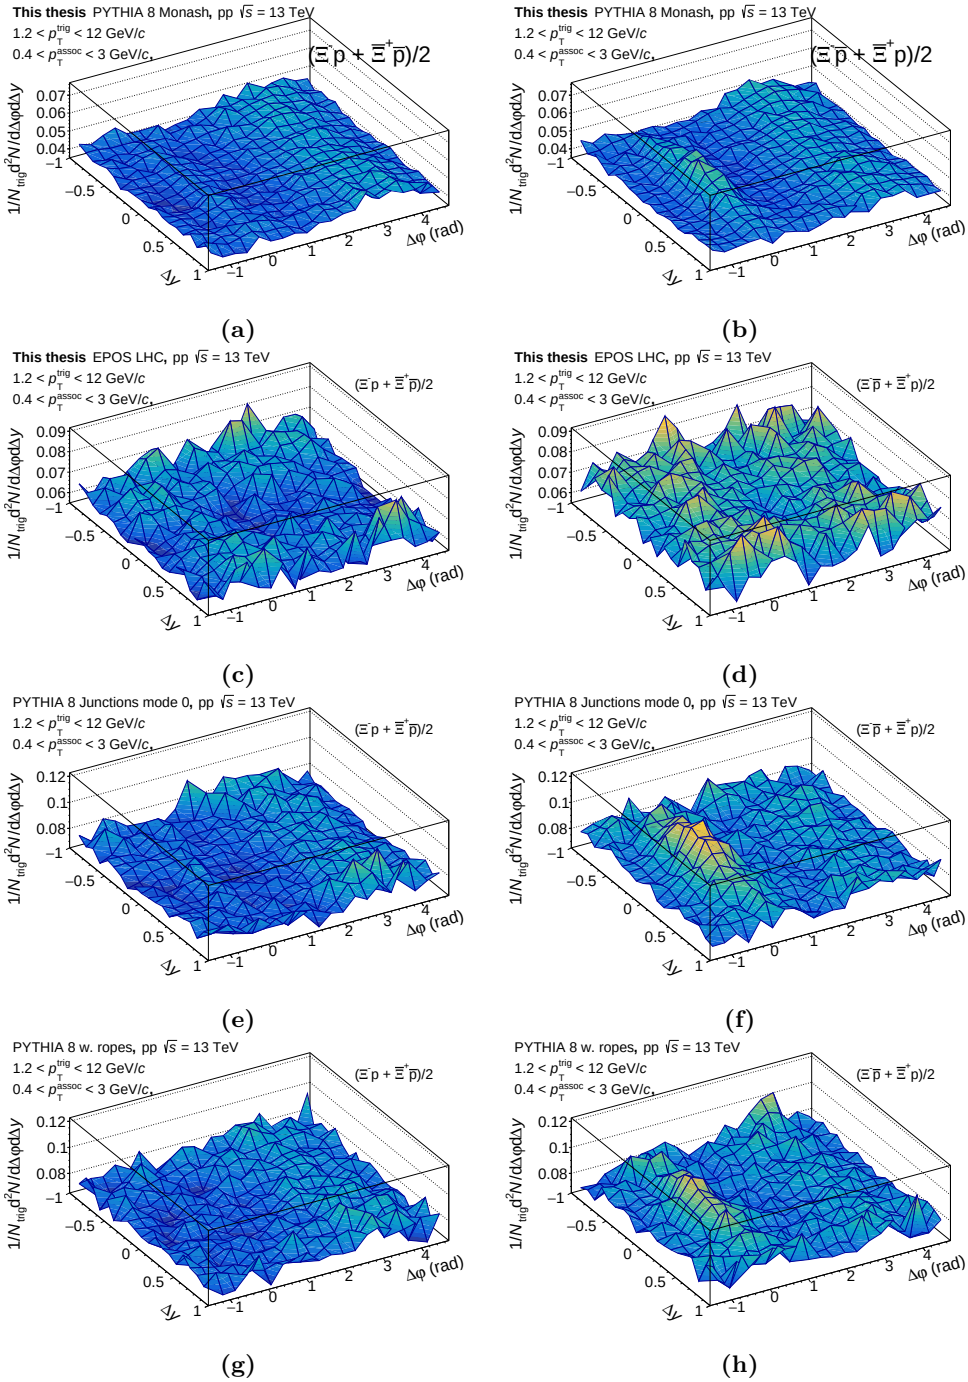

Figure 9.10 – Simulated 2D-correlations between Ξ baryons and protons, left: same-B correlations, right: opposite-B correlations. **(a)**-**(b)** PYTHIA8, **(c)**-**(d)** EPOS LHC, **(e)**-**(f)** PYTHIA8 with junctions, and **(g)**-**(h)** PYTHIA8 with ropes.

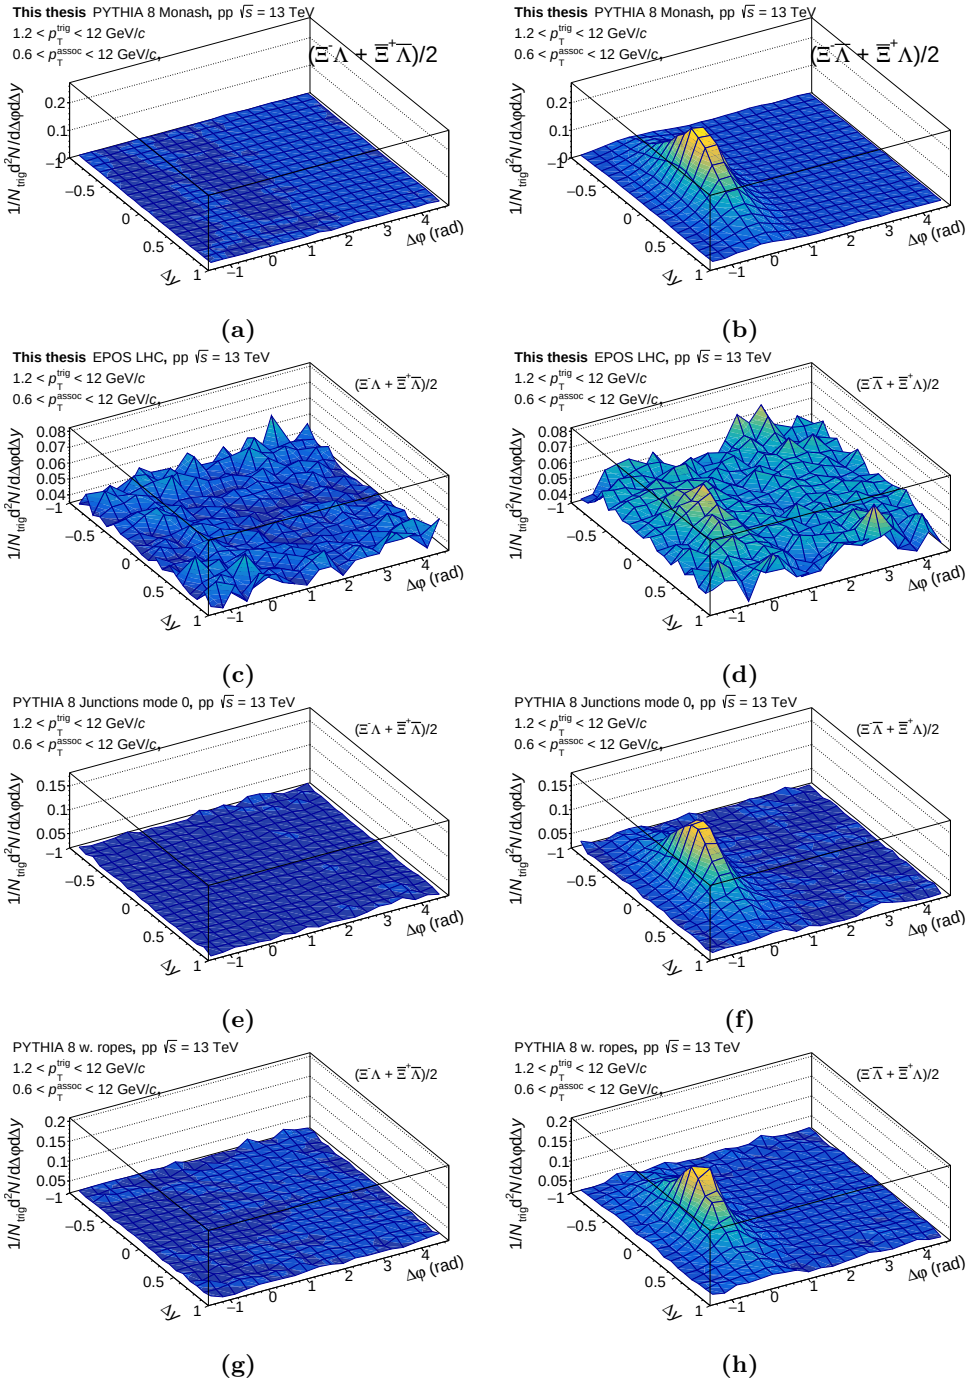

Figure 9.11 – Simulated 2D-correlations between Ξ and Λ baryons, left: same-B correlations, right: opposite-B correlations. **(a)**-**(b)** PYTHIA8, **(c)**-**(d)** EPOS LHC, **(e)**-**(f)** PYTHIA8 with junctions, and **(g)**-**(h)** PYTHIA8 with ropes.

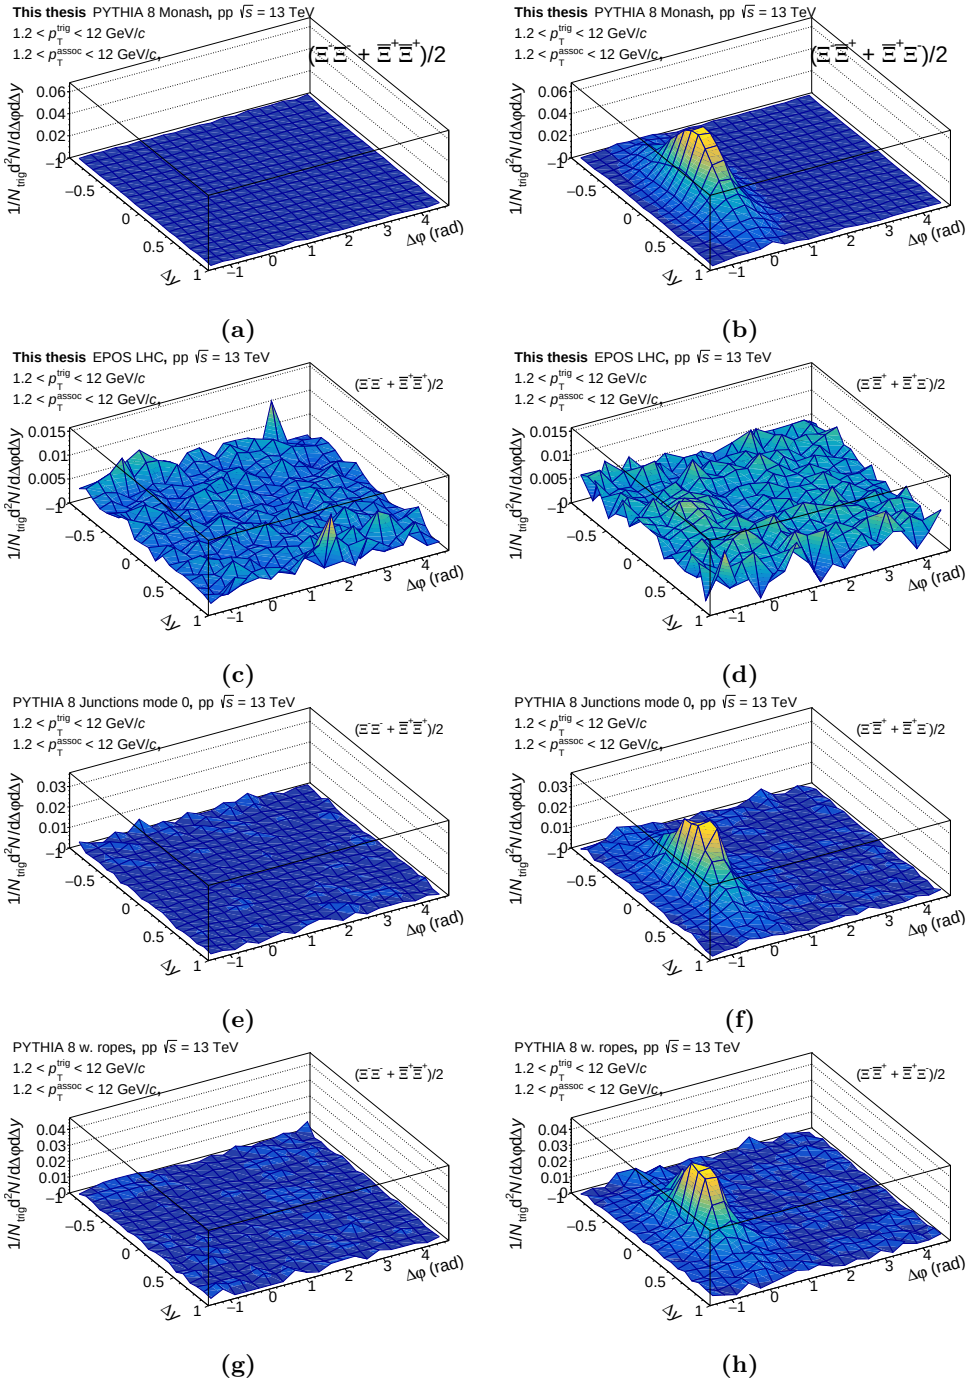

Figure 9.12 – Simulated 2D-correlations between pairs of Ξ baryons in the same event, left: same-B correlations, right: opposite-B correlations. **(a)**-**(b)** PYTHIA8, **(c)**-**(d)** EPOS LHC, **(e)**-**(f)** PYTHIA8 with junctions, and **(g)**-**(h)** PYTHIA8 with ropes.

narrower than what is observed in data. For EPOS, the correlation function is instead quite flat and similar to $\Xi - \mathrm{p}$ correlations, with only a weak near-side peak for correlations between pairs of opposite baryon number, which in particular is present for Λ. This is wider than what is seen in PYTHIA, but the sample used for the simulations is a bit too small to make a quantitive statement.

9.3 Comparisons Between Models and Data

To make quantitive comparisons between models and data, projections have been made in the same manner as in Section 9.1.1, including differences between opposite- and same-quantum number correlations. These results – only including the minimum-bias sample in data – are shown in Figs. 9.13, 9.14, 9.15, 9.16, and 9.17 for correlations between Ξ baryons and pions, kaons, protons, Λ baryons, and other Ξ baryons, respectively.

On the away side, the difference between same-sign and opposite-sign correlations is in each case – for models as well as data – consistent with a flat value. Moreover, this value is for most of the measurements of the same order of magnitude for different particle species. Therefore, weighted averages of the difference in away-side per-trigger yields between same- and opposite-quantum number particles are drawn in the same plot in Fig. 9.18, for both data and the various MC models. This is done in order to facilitate comparisons both between model and data, different particle species, and near-side versus away-side correlations, although one should be aware that the kinematic ranges are slightly different for different measurements. In order to quantify global correlations within the detector acceptance, a similar integration has also been done over the full range in $\Delta\varphi$ and going as far out in Δy as the detector allows (i.e. $|\Delta y| < 1$) for both data and all models. The results are shown in Fig. 9.19.

For $\Xi - \pi$ correlations, the shape in data is well-described by all models, but the Monash tune describes the underlying event better (the other models overestimate the baseline), whereas EPOS and the other PYTHIA flavours better quantify the difference between opposite- and same-sign correlations. However, even these models underestimate this difference on the near side.

For $\Xi - \mathrm{K}$ correlations, all models fail in describing the data. While PYTHIA in general, and the rope tune in particular, succeeds in describing the same-sign correlations outside the peak region, the near-side peak is too large, and also too narrow for opposite-sign correlations. Consequently, the difference between correlations of different charge combinations gets overestimated on the near side, ranging from about 40% for the rope model to about 60% for the junction model, and even more at the centre of the peak in the near-side projection. On the away side, the Monash tune underestimates the difference, but the rope and junction models do a better job. For EPOS, the shape is very different from data, and the same-sign correlations are too strong. Moreover, there is a very flat difference between opposite- and same-sign correlations, with only a slight tendency of a near-side peak. On the other hand, the integral of the difference is just as for the PYTHIA Monash tune and junction models similar to what is observed in data.

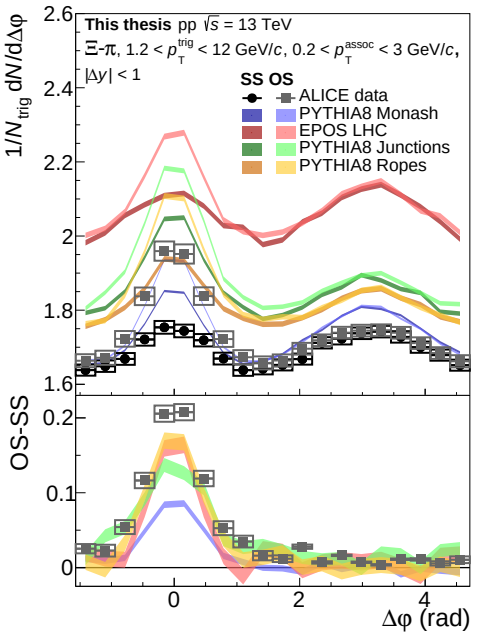

(a)

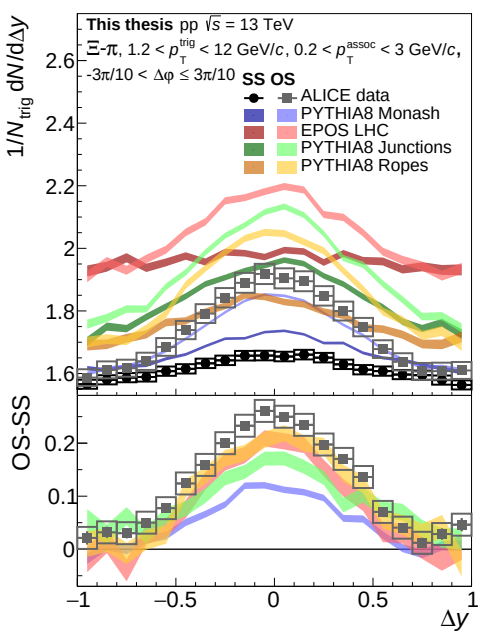

(b)

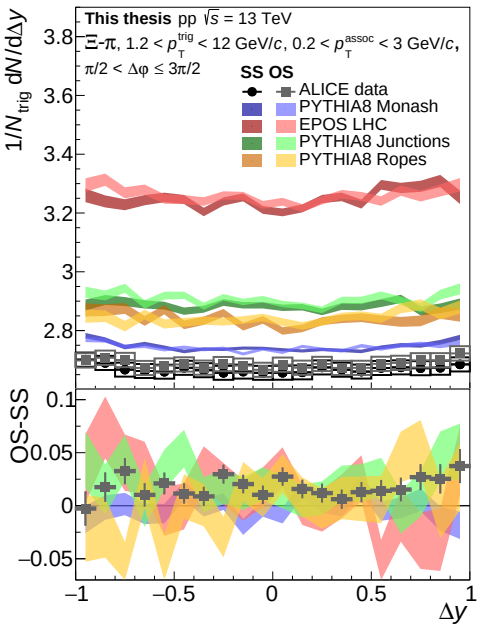

(c)

Figure 9.13 – Correlations between Ξ baryons and pions, comparing the results from EPOS and the various flavours of PYTHIA with ALICE data. The bottom panels show differences between opposite- and same-sign correlations for data, as well as each of the models. Statistical and systematic errors are marked with bars and boxes, respectively. **(a)** Projection in $\Delta\varphi$, **(b)** projection in Δy on the near side ($-3\pi/10 < \Delta\varphi \leq 3\pi/10$), **(c)** projection in Δy on the away side ($\pi/2 < \Delta\varphi \leq 3\pi/2$). Legend: SS = same sign, OS = opposite sign.

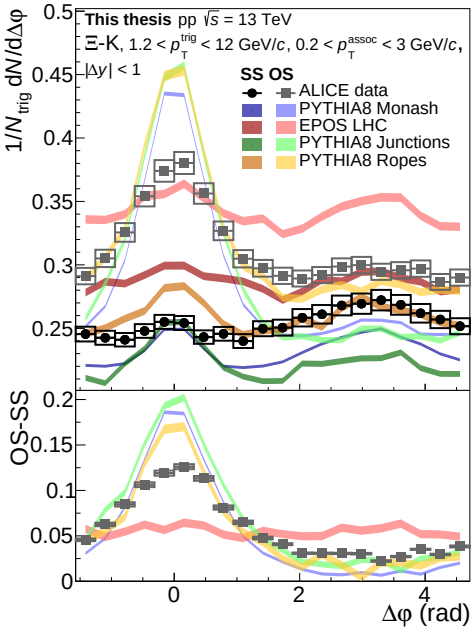

(a)

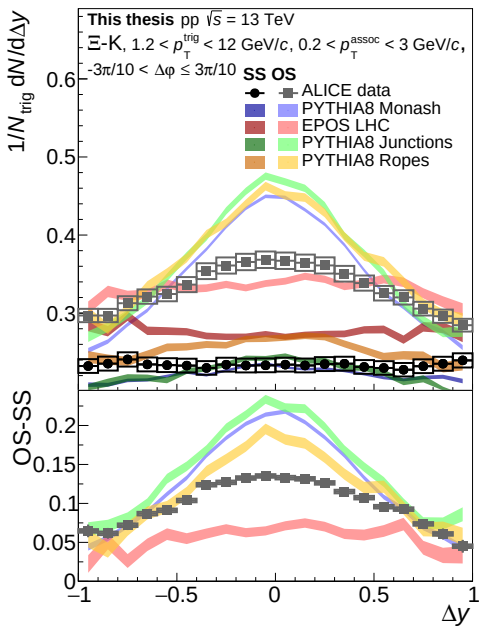

(b)

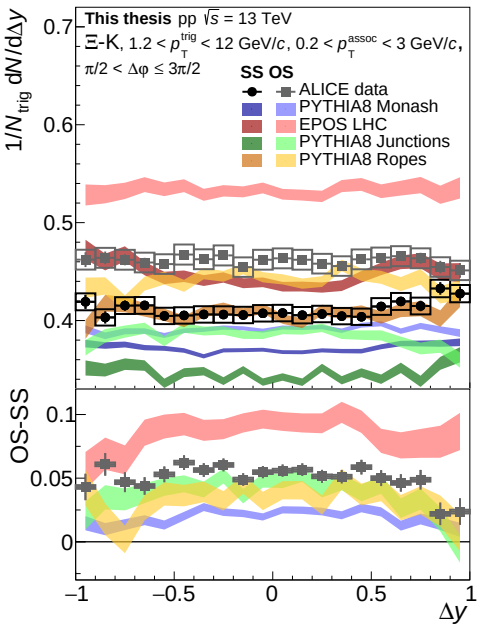

(c)

Figure 9.14 – Correlations between Ξ baryons and kaons, comparing the results from EPOS and the various flavours of PYTHIA with ALICE data. The bottom panels show differences between opposite- and same-sign correlations for data, as well as each of the models. Statistical and systematic errors are marked with bars and boxes, respectively. **(a)** Projection in $\Delta\varphi$, **(b)** projection in Δy on the near side ($-3\pi/10 < \Delta\varphi \leq 3\pi/10$), **(c)** projection in Δy on the away side ($\pi/2 < \Delta\varphi \leq 3\pi/2$). Legend: SS = same sign, OS = opposite sign.

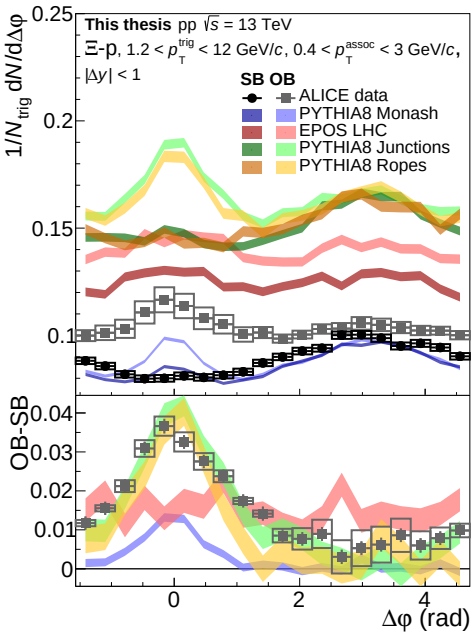

(a)

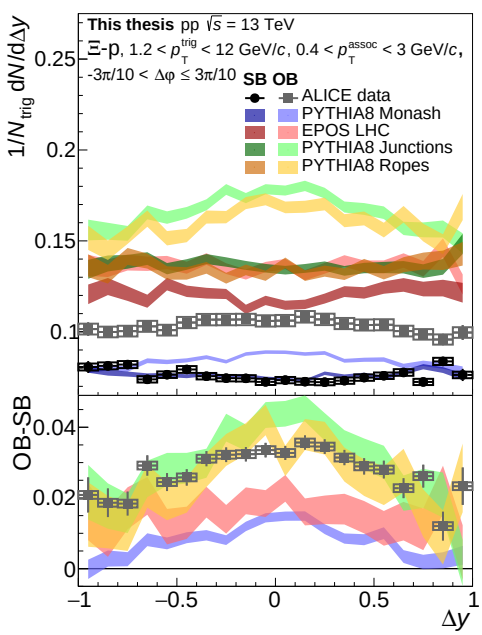

(b)

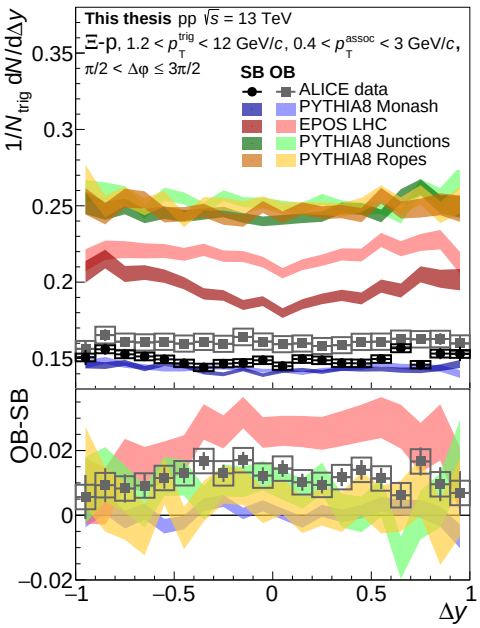

(c)

Figure 9.15 – Correlations between Ξ baryons and protons, comparing the results from EPOS and the various flavours of PYTHIA with ALICE data. The bottom panels show differences between opposite- and same-B correlations for data, as well as each of the models. Statistical and systematic errors are marked with bars and boxes, respectively. **(a)** Projection in $\Delta\varphi$, **(b)** projection in Δy on the near side ($-3\pi/10 < \Delta\varphi \leq 3\pi/10$), **(c)** projection in Δy on the away side ($\pi/2 < \Delta\varphi \leq 3\pi/2$). Legend: SB = same baryon number, OB = opposite baryon number.

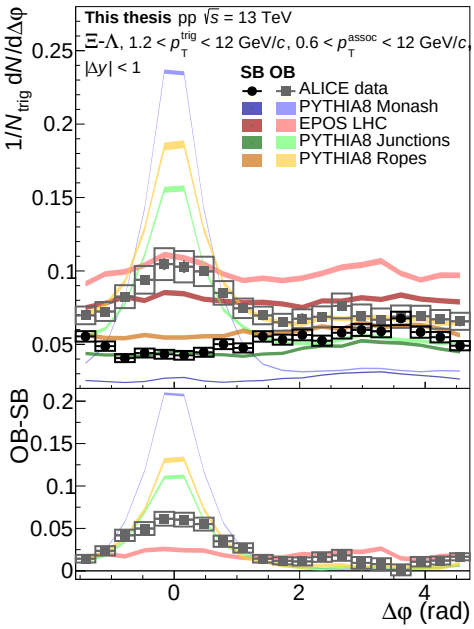

(a)

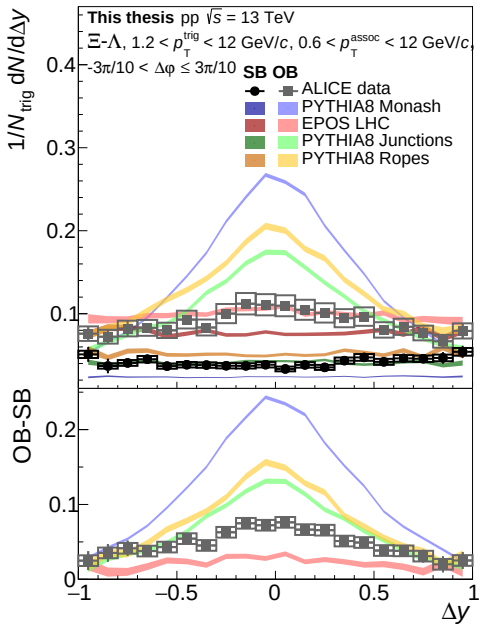

(b)

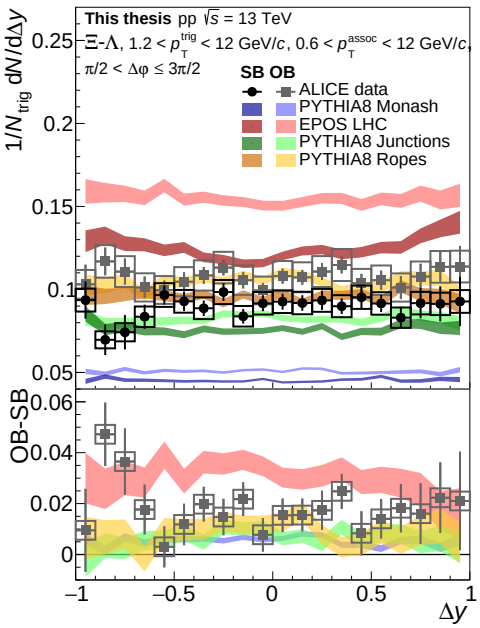

(c)

Figure 9.16 – Correlations between Ξ and Λ baryons, comparing the results from EPOS and the various flavours of PYTHIA with ALICE data. The bottom panels show differences between opposite- and same-B correlations for data, as well as each of the models. Statistical and systematic errors are marked with bars and boxes, respectively. **(a)** Projection in $\Delta\varphi$, **(b)** projection in Δy on the near side ($-3\pi/10 < \Delta\varphi \leq 3\pi/10$), **(c)** projection in Δy on the away side ($\pi/2 < \Delta\varphi \leq 3\pi/2$). Legend: SB = same baryon number, OB = opposite baryon number.

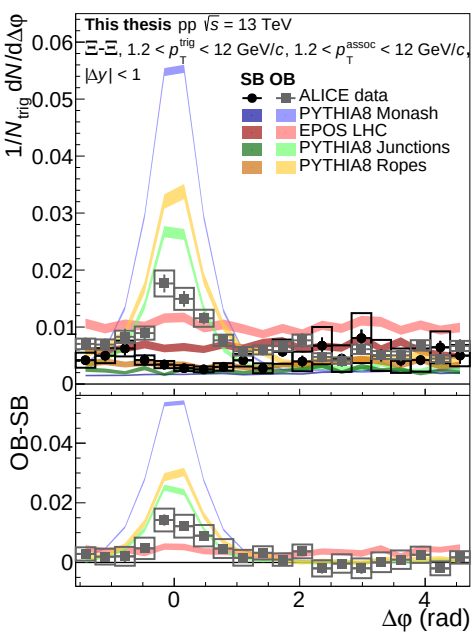

(a)

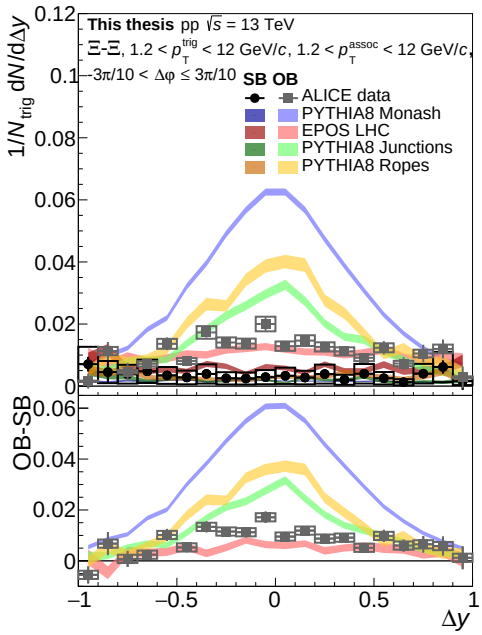

(b)

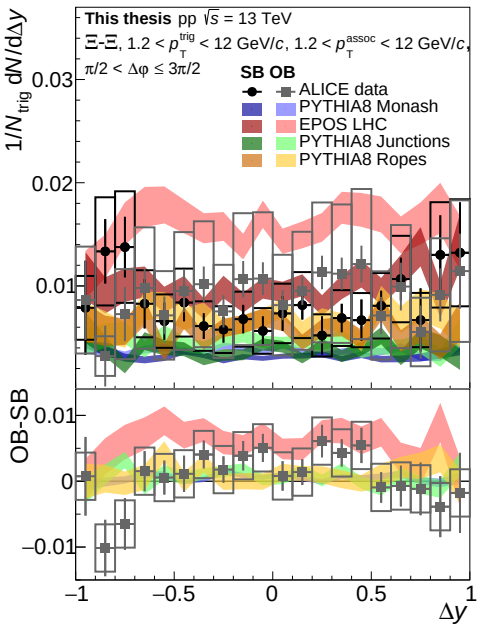

(c)

Figure 9.17 – Correlations between pairs of Ξ baryons, comparing the results from EPOS and the various flavours of PYTHIA with ALICE data. The bottom panels show differences between opposite- and same-B correlations for data, as well as each of the models. Statistical and systematic errors are marked with bars and boxes, respectively. **(a)** Projection in $\Delta\varphi$, **(b)** projection in Δy on the near side ($-3\pi/10 < \Delta\varphi \leq 3\pi/10$), **(c)** projection in Δy on the away side ($\pi/2 < \Delta\varphi \leq 3\pi/2$). Legend: SB = same baryon number, OB = opposite baryon number.

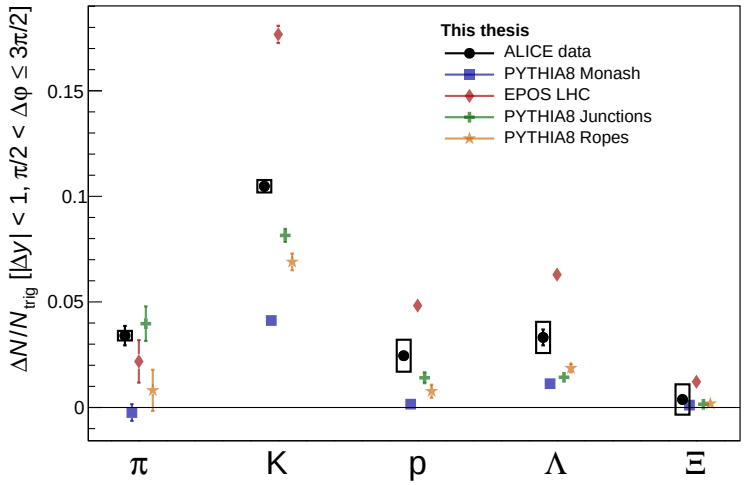

Figure 9.18 – Differences in integrated per-trigger yields between opposite- and same-sign or baryon number particles, integrated over the away side ($\pi/2 < \Delta\varphi \leq 3\pi/2$), for all particles and models included in this analysis. Statistical and systematic errors are marked with bars and boxes, respectively. Note that since these yields are integrated, they need to be scaled by $1/\pi \approx 0.32$ and 0.5, respectively, to match the normalisation used in the $\Delta\varphi$ and away-side Δy projections in Figs. 9.13–9.17.

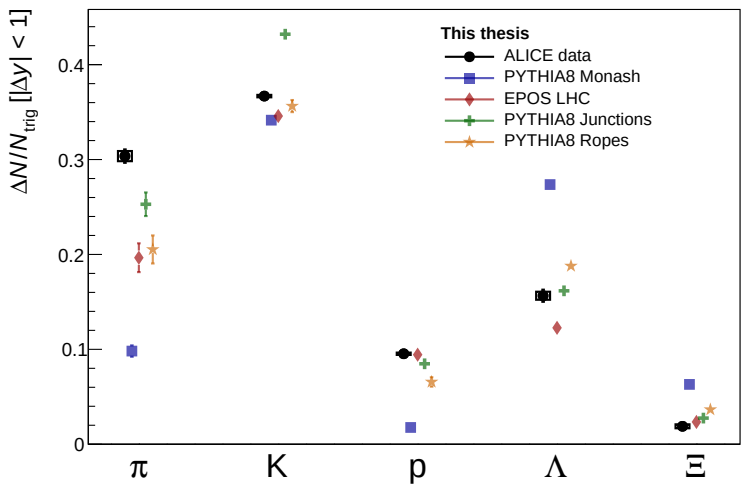

Figure 9.19 – Differences in per-trigger yields between opposite- and same-sign or baryon number particles, integrated over the volume within $|\Delta y| < 1$, for all particles and models included in this analysis. Statistical and systematic errors are marked with bars and boxes, respectively.

For $\Xi-$p correlations, the Monash tune again describes the baseline very well, apart from a small near-side peak not present in data, but strongly underestimates the correlation difference. The other PYTHIA flavours overestimate proton production in general, with about an 80% difference to data for same-sign pairs, but the difference between opposite- and same-sign particles is very well reproduced for both of these models. Again, EPOS describes the shape of the correlation function poorly and overestimates the baseline, but the integral of the difference gets very close to what is observed in data.

For $\Xi-\Lambda$ correlations, all PYTHIA flavours strongly overestimate the near-side peak in opposite-baryon number correlations. Despite a large difference between the flavours, the correlation difference is overestimated by about 70% in the best model (junctions). Moreover, the Monash tune underestimates the baseline with nearly 50%, but this is remedied by the rope and junction models, which get close to what is observed in data, although they largely fail to describe the near-side depletion. Also EPOS does a poor job in describing the correlation function, despite that it is very wide in Δy in data. Just as for the other results, the baseline is overestimated and the difference is too flat.

For $\Xi-\Xi$ correlations, the shape in both PYTHIA and EPOS is quite similar to their $\Xi-\Lambda$ correlation results, although the difference between different PYTHIA extensions is even more pronounced and the near-side peak is relatively stronger in PYTHIA than in EPOS. Also in data, the shape is similar as for $\Xi-\Lambda$ correlations, as well as the relative difference to the MC models.

Chapter 10

Discussion of Correlations and Conclusions

There are two aspects of the correlation function that can be studied to increase the understanding of particle production mechanism: the *shape* and the *magnitude*. In principle, the correlation function can be decomposed into several different parts originating from each of the various production mechanisms involved. The overall shape will be determined by the relative strength of each production mechanism, whereas the overall magnitude depends on the combined strength of all these mechanisms. In practice, it is not that simple, since several different mechanisms may affect the correlation function in a similar way, the same mechanism may affect several different features of the correlation function, and there is usually a strong interplay between different mechanisms. Moreover, aspects from different models are generally not very easy to combine, and tuning model parameters may affect other results, such as the particle spectra. Nevertheless, by studying the correlation function, we can give input to the theorists on which mechanisms are likely involved in particle production and which are less so, and also if some aspects are not well described by either model. In extension, this will get us closer to understanding what is happening in a pp collision, and more specifically what is the origin of the strangeness production in small systems.

Most of the results presented in Chapter 9 are minimum-bias results, meaning that results are combined across multiplicity and momentum of trigger and associated particles. While this maximises the statistical sample, it also means that different mechanisms may be involved. This is the reason for the multiplicity dependent study, which is discussed in Section 10.3. Further separation could have been done by also studying p_{T}-dependent correlations, but there was unfortunately not enough time for this.

10.1 Correlations Originating from the Underlying Event

The measured correlation function is a combination of correlations from the production of the Ξ baryon and correlations appearing from the event itself. The latter are due to that the Ξ is more likely to appear in certain regions in phase space, and have already been introduced as underlying-event correlations.

In the time frame of a proton collision, the total quantum numbers shared by the two protons are conserved, i.e. the total charge $Q = 2$, the quark number balances

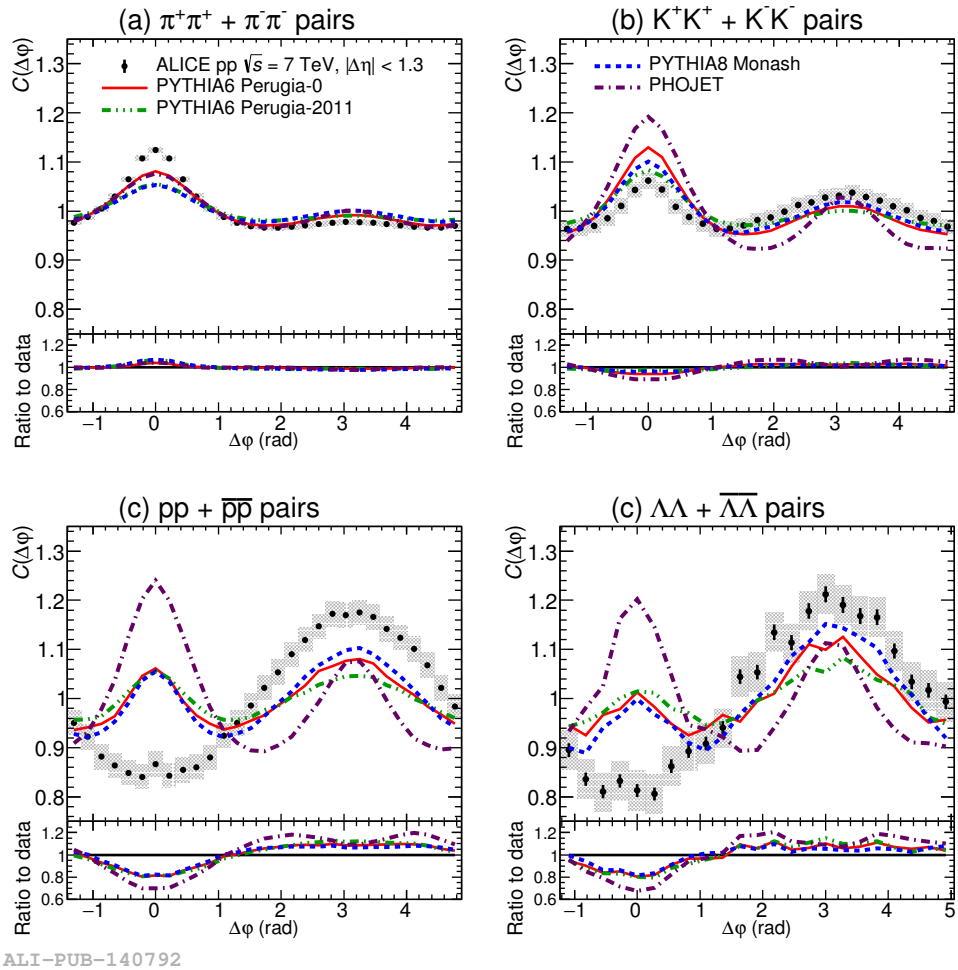

Figure 10.1 – Symmetric correlation functions (defined by Eq. (8.2)), projected in $\Delta\varphi$, measured by ALICE at $\sqrt{s} = 7\,\text{TeV}$, along with model predictions by various PYTHIA tunes and PHOJET (not discussed in the text) for pairs of like-sign **(a)** pions, **(b)** kaons, **(c)** protons, and **(d)** Λ baryons. Figure taken from Ref. [132].

$n_u - n_{\bar{u}} = 4$, $n_d - n_{\bar{d}} = 2$, baryon number $B = 2$, and strangeness $S = 0$. Therefore, any excess of these quantum numbers needs to be balanced by the opposite quantum number and consequently, two particles not sharing any opposite quantum numbers are not likely to be produced from the same interaction. Therefore, any such correlations must originate from the underlying event. This is the case for same-charge $\Xi-$meson correlations and same-baryon number $\Xi - \Lambda$ and $\Xi - \Xi$ correlations. While $\Xi - p$ pairs have opposite charge, which in principle could put them closer in phase space, there still needs to be some intermediate mechanism, since they do not share any quark-antiquark pair. Therefore, such correlations are likely mostly due to the underlying event as well.

The best probe for studying the underlying event is through pions, since these are by far the most abundant particles produced in the event. The same-sign $\Xi - \pi$ correlation function (Fig. 9.1a) is quite flat, indicating that the correlation to the underlying event is quite weak, but the presence of the near-side jet peak and away-side ridge indicates that Ξ baryons are more likely to be produced within the jet cone and the observed peak is due to minijet fragmentation. This is well described by both PYTHIA and EPOS, although the magnitude of the near-side peak is somewhat different between the models (including different PYTHIA extensions). A similar effect has been observed also in $\pi - \pi$ and $K - K$ correlations in pp collisions, which similarly are well described by PYTHIA, at least qualitatively [132]. This is shown in the upper panels of Fig. 10.1.

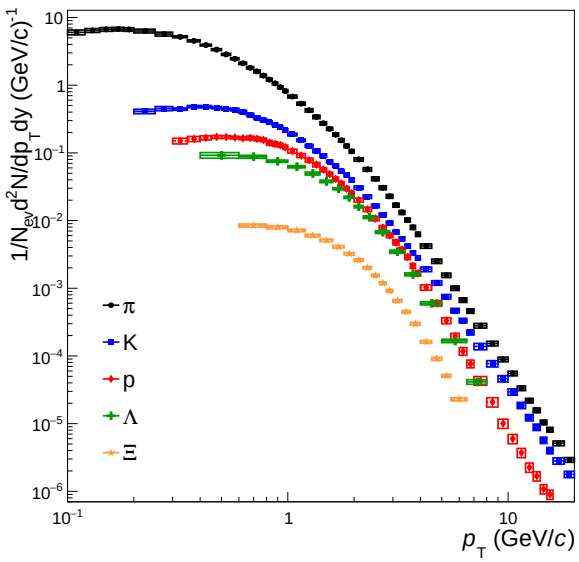

Figure 10.2 – Charge inclusive yields of various hadron species as a function of p_T, measured by the ALICE detector at a collision energy of $\sqrt{s} = 13\,\text{TeV}$, based on the results published in Ref. [128].

The strength of the underlying-event correlation function is expected to scale with the integrated spectra, which for ALICE data are shown in Fig. 10.2. This is the reason why the magnitude of the $\Xi - K$ correlation function is only about 15% of the

$\Xi - \pi$ correlation function. The physics should be similar though, which explains the similarity in shape, which is the case also for PYTHIA and to some extent also for EPOS. Also the $\Xi - \mathrm{p}$, $\Xi - \Lambda$ and $\Xi - \Xi$ correlation functions scale with the spectra, but one should bear in mind that the p_{T} and η ranges are somewhat reduced for these particle species, affecting the total yield. Yet, this explains why the statistical errors for $\Xi - \Xi$ correlations are so large compared to the others – about ten Λ baryons are produced for every Ξ baryon.

While there is a clear scaling in magnitude, there is also a clear division in shape between $\Xi-$meson and $\Xi-$baryon correlations. There does not appear to be any significant difference in shape between same-B $\Xi - \mathrm{p}$, $\Xi - \Lambda$, and $\Xi - \Xi$ correlations, proving that baryon number is more important than charge for describing these correlations. The near-side depletion indicates that a Ξ baryon is unlikely to be produced together with another baryon of the same baryon number, which means that they are more likely to appear in back-to-back jets than in the same jet, for instance. This is quite well described in PYTHIA, but less so in EPOS, which is not too surprising given that this requires a different mechanism. In PYTHIA, baryons are produced through double string breakings, cf. Fig. 3.24a, meaning that two baryons cannot be produced from the same string. This results in a natural separation in phase space, but studies of $\mathrm{p} - \mathrm{p}$ and $\Lambda - \Lambda$ correlations indicate that this cannot be the full explanation why the $\Xi - \mathrm{p}$ correlation in PYTHIA agrees with data. Same-baryon number $\mathrm{p} - \mathrm{p}$ and $\Lambda - \Lambda$ correlations have a similar shape as $\Xi - \mathrm{p}$ correlations in data (Fig. 9.15a), but in PYTHIA there is a clear near-side peak for these systems – which is much less pronounced for the results presented in this thesis – resulting in a very different shape overall, as shown in the lower panels of Fig. 10.1 [132]. Therefore, there must be another explanation, which is currently not understood.

The magnitudes of the $\Xi-\pi$, $\Xi-\mathrm{K}$, and $\Xi-\mathrm{p}$ underlying-event correlation functions are by far best described by the Monash tune of PYTHIA (except that the rope extension does slightly better for $\Xi - \mathrm{K}$ correlations), which among other things is tuned to accurately describe light-hadron spectra. The other models tend to overestimate these quantities in $\Xi - \pi$ and $\Xi - \mathrm{p}$ correlations. On the other hand, PYTHIA8 Monash severely underestimates the $\Xi - \Lambda$ and $\Xi - \Xi$ correlation functions, and here the other models do a better job – and in particular the rope model (EPOS still tends to overestimate these).

For PYTHIA, all these observations can be understood from studying the spectra, as seen in Fig. 10.3. The junction model adds an additional mechanism for producing baryons, cf. Fig. 3.24b, and since this mechanism is required also for the PYTHIA implementation of the ropes (cf. Section 3.8.3), this is present also in the rope model. Therefore, in order to produce the number of Λ and Ξ baryons observed in data – and the increase with multiplicity – the proton yields get overestimated, cf. Fig. 10.3b. It has turned out to be challenging for the PYTHIA developers to achieve strangeness enhancement without introducing an excess of protons, but since several mechanisms are involved, this would likely require a reduction of the baryon production mechanism present in the standard version of PYTHIA in order for this model to be successful.

For EPOS, the underlying-event correlations are overestimated for all particles, although this is not the case for most of the spectra. This means that the relative multiplicity of other particles to the number of triggers is overestimated in events containing

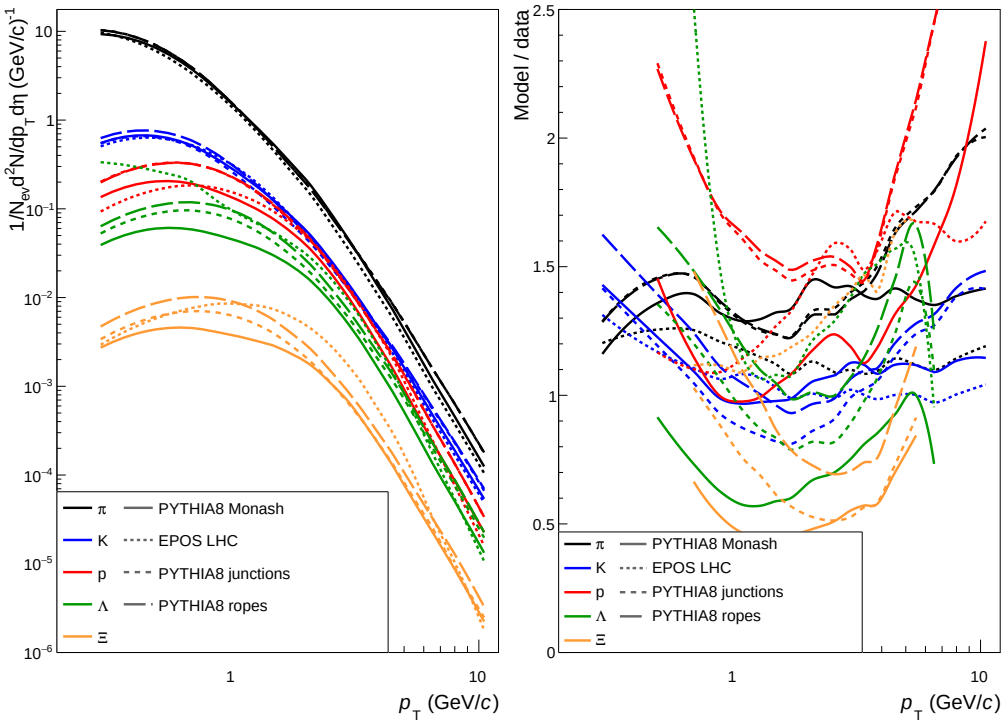

(a) **(b)**

Figure 10.3 – **(a)** p_{T} spectra of various particles simulated in pp collisions at $\sqrt{s} = 13\,\mathrm{TeV}$ by the event generators used in this thesis. **(b)** Ratios between models and data, based on the results shown in Fig. 10.2. Some interpolation errors may occur at the ends of the p_{T} intervals. In both plots, error bars have been excluded for clarity.

a Ξ baryon. A possible explanation for this is that Ξ baryons are mostly associated with events with very high multiplicity in EPOS LHC (cf. Fig 3.31), and the fraction of Ξ baryons in each such event likely is lower than in minimum-bias events data (this has not been investigated, but it seems reasonable given the increased baseline for the high-multiplicity results in data). For PYTHIA on the other hand, the number of Ξ baryons per event roughly scales with multiplicity (cf. Fig. 3.31).

Another peculiarity with EPOS is that the correlation function has some features which are not at all present in data – and in particular stronger correlations between the trigger and other particles with a large separation in Δy than with particles close in phase space. This would indicate that same-quantum number particle pairs are correlated transversely to the trigger (in phase space) rather than back-to-back, which is something that has not been observed in data. This is likely due to a limitation in the model description, since the micro-canonical ensemble only conserves quantum

numbers in slices in pseudorapidity, and not locally [90, 133], which very well may lead to an improper description of the event shape as well.

10.2 Correlations Originating from Ξ Baryon Interactions

Unlike same-quantum number particle pairs, oppositely charged Ξ−meson pairs and opposite-baryon number Ξ−baryon pairs share at least one quark-antiquark pair. In the case of pions, kaons, and protons, Ξ only has one quark of opposite quantum number, whereas Λ and Ξ baryons have two and three such quarks, respectively. Therefore, one can expect that the interactions between Ξ and the latter two are quite different from the other ones, which indeed seems to be the case, but as will discussed here there are strong indications that there are more properties that give rise to a distinction between the different particles and their interactions with the Ξ baryon. What is important is that the underlying-event correlations are expected to be the same for same-sign as for opposite-sign correlations, and therefore this part of the discussion focuses on the differences between the two, i.e. the lower panels of Figs. 9.13–9.17.

The near side of the correlation function gives information about particles that are locally correlated with the Ξ baryon, and therefore all quantum numbers originating from the initial state of the collision will be preserved in this region. This is particularly relevant for the near-side peak, also referred to as the *jet peak*, where an excess of particles is seen due to the trigger being more likely to be produced in the regions of the event with higher multiplicity, usually associated with jets. Particles in the jet peak are produced close in phase space relative to the Ξ baryon during the hadronisation phase, and the excess relative to the underlying event gives information about which quantum numbers are remaining close to the Ξ baryon in the later stages of the event. Correlations in this region can also be used to study the hadronisation mechanism.

The other extreme is the excess of particles observed on the away side, which gives information on which quantum numbers have decoupled from the trigger early on in the collision, which in turn gives insight into which processes are involved. This decoupling can be quantified both from the magnitude of the away side difference and on the width of the near-side peak, where the latter gives a measure of the decoupling rate. Any excess of particles on the away side not related to the Ξ production, should cancel out when averaging over many events. The flatness of the away-side excess indicates that the initial-state effects are distributed rather isotropically in the event. In some sense we are lucky about this, since it makes it possible to measure this more precisely by averaging over a large region in phase space, as is done in Fig. 9.18.

Finally, by integrating the balancing part of the correlation function over the full detector acceptance, one can get access to how much each particle species contributes to the global balance of the Ξ quantum numbers, as is done in Fig. 9.19. The corresponding balances for charge, strangeness, and baryon number[1] contained within the detector acceptance are summarised in Table 10.1 for both data and all MC models. In data as well as for EPOS LHC, and to some extent for the junction model, a much larger

[1]The charge balance is obtained from summing the contributions from $\pi+\mathrm{K}+\Xi-\mathrm{p}$, the strangeness balance from $\mathrm{K}+\Lambda+2\cdot\Xi$, and the baryon balance from $\mathrm{p}+\Lambda+\Xi$.

contribution of the charge is contained within the detector than for the other quantum numbers. This can to some extent be explained by additional contributions from $\Sigma^{\pm}$, Ξ^0, and Ω (cf. Appendix C), but the main reason is likely that these quantum numbers to a larger extent are balanced at large separations in rapidity, which is a recurrent observation in the following discussion. For PYTHIA8 (in particular the Monash tune), the lower charge balance than in data can be explained from an expected large contribution from Σ^- (dds, shares two quarks with Ξ^-), which is missing, yet much more of the detected strangeness and baryon number are balanced close to the trigger. For the junction extension, the description is closer to what is seen in data, but one should carefully investigate the contributions from other strange baryons before drawing too strong conclusions. By adding ropes, the description again gets closer to what is predicted by the Monash tune.

Table 10.1 – Sum of contributions to the charge, strangeness, and baryon number balance of the Ξ trigger, contained within $|\Delta y| < 1$ and originating from pions, kaons, protons, Λ, and Ξ baryons, for data and each MC model studied in this thesis. Data derived from Fig. 9.19. Only statistical errors are stated in the table.

Data/Model	Charge	Strangeness	Baryon number
ALICE data	$59.4 \pm 0.9\%$	$28.1 \pm 0.4\%$	$27.1 \pm 0.7\%$
PYTHIA8 Monash	$48.5 \pm 0.7\%$	$37.1 \pm 0.1\%$	$35.4 \pm 0.2\%$
EPOS LHC	$47.1 \pm 1.7\%$	$25.8 \pm 0.4\%$	$24.0 \pm 0.6\%$
PYTHIA8 Junctions	$62.8 \pm 1.4\%$	$32.4 \pm 0.3\%$	$27.4 \pm 0.5\%$
PYTHIA8 Ropes	$53.2 \pm 1.7\%$	$30.9 \pm 0.4\%$	$29.0 \pm 0.6\%$

To summarise, both strangeness and baryon number conservation are much more localised in PYTHIA than in data, although this seems to at least partly be remedied by adding junctions, and in this respect EPOS is more successful. This translates to an earlier decoupling of strangeness and baryon number in the system evolution – but not charge – in data than in PYTHIA.

10.2.1 $\Xi - \pi$ Correlations

Pions couple to the Ξ baryon by sharing a down quark, but they will also affect the correlation function through charge balance. Since all quantum numbers are preserved in the jet peak, any excess in charge induced by other correlations must be balanced. Balancing this with kaons would create an excess in strangeness, which in turn needs to be balanced by other strange particles. Since K^0_s is a superposition of K^0 and $\overline{\mathrm{K}}^0$ to nearly equal proportions [10, pp. 289-291], resulting in a net strangeness of zero, this would need to be done by another particle, such as Λ. Since no such excess has been observed in the $\Xi - \Lambda$ correlation function, one can conclude that kaons should not play a major role in the charge balance. Similarly, balancing charge with protons would violate baryon number conservation, which only leaves pions.

To extract the impact of the d quark to the $\Xi - \pi$ correlation function, one should look at the away side, where a small excess of $0.034 \pm 0.005\,(\mathrm{stat.}) \pm 0.002\,(\mathrm{syst.})$ pions at midrapidity per Ξ baryon is observed. The presence of the Ξ^- ($\overline{\Xi}^+$) does however result in an excess of antiprotons (protons) on the away side, resulting in an excess of

negative (positive) charge in the event. This needs to be balanced somehow, and since charge balance seems to be a fairly localised effect (see below), this has to cause at least part of the pion excess[2]. The away-side proton excess of $0.025 \pm 0.001\,(\text{stat.}) \pm 0.008\,(\text{syst.})$ per Ξ baryon is actually nearly consistent with the pion excess being fully due to charge balance of protons, meaning that most of the d quark balance is likely achieved by protons and not pions, unless the mechanism for the down quark coupling is very different from that of strange quarks, which is unlikely[3]. Therefore the excess in oppositely-charged $\Xi - \pi$ correlations is likely strongly dominated by charge balance.

While the study of $\Xi - \pi$ correlations give information of the charge distribution in Ξ triggered events, it is not equal to it. For a complete description, one would need to study the full correlation between the Ξ baryon and oppositely charged particles, which very well is achievable, since these are dominated by pions, kaons, and protons, but has not been measured locally since it is not the purpose of this study. Instead, the pions give a measure of the charge distribution not covered by other charged particles, particularly kaons. This explains why PYTHIA gives a much lower excess of pions than what is seen in data; this is simply an effect of the increased number of kaons. By summing the contributions from all charged particles which are experimentally accessible, all versions of PYTHIA get fairly close to the experimental value (exactly how close is difficult to say due to the impact of $\Sigma^{\pm}$), showing that PYTHIA quite well describes the total charge distribution. While the charge balance peak is the only feature that EPOS reproduces, summing all contributions does not give a value of the charge excess even close to what is seen in data. Instead, this is largely balanced by an excess on the away side, which is the reason why the global charge balance is not very different from the other MC models or data.

While not probing the full charge distribution, the narrowness of the $\Xi - \pi$ correlation peak shows that charge conservation is a quite localised phenomenon, mostly associated with correlations within the jet, although the peak is not quite as narrow as the contribution from pions alone. Given the success of the EPOS description of jet-hadron correlations (Fig. 3.29), it is quite expected that the jet peak should show up for inclusive charged particles. What this study shows is that only pions contribute to the jet correlations in EPOS, which is not exactly what is seen in data (but as I will return to, this picture may not be entirely wrong).

The charge imbalance seen on the away side should on the other hand arise from other processes, such as quark or baryon number decoupling in the initial phase, as discussed elsewhere in this chapter.

[2]The excess of strange quarks on the away side is coupled to the production of (what later becomes) the Ξ baryon. As such, this will create a dipole already in the initial stage and should not contribute to the proton charge balance.

[3]For Ξ^-, part of the down quark balance can however be achieved by the valence quarks from the colliding protons, which would reduce the balance for d quarks relative to s quarks. There is a significantly weaker correlation difference on the away side for $\Xi^- - \pi$ correlations than for $\overline{\Xi}^+ - \pi$ correlations (not shown in this thesis), but a corresponding separation is not observed for $\Xi - \text{p}$ correlations.

10.2.2 $\Xi - \mathrm{K}$ Correlations

With charge balancing ruled out as a source of the excess of oppositely charged kaons, this must be a probe of where strangeness is produced in the event. However, just as the pion excess results from charge balance, the kaon excess results from *strangeness balance* remaining after production of strange baryons. Despite a large excess on the near side, reaching slightly more than 50% of the underlying-event correlation at the peak (compared to about 15% for the pions), this is somewhat suppressed on the near side due to the impact of Λ and – to a lesser extent – Ξ baryons. The width of the $\Xi-\mathrm{K}$ peak is similar to the $\Xi-\Lambda$ peak however, indicating that $\Xi-\mathrm{K}$ correlations alone fairly well describe how strangeness is distributed in the event, at least qualitatively. The slow fall-off – particularly in Δy – and the large excess on the away side, clearly indicate that strangeness is diluted in the system, meaning that it largely is an initial-state effect.

While diluted, the clear near-side peak shows however that strangeness is not nearly isotropically distributed in the system, with the Ξ baryon being produced at the side with more strangeness. This is likely a combination of initial- and final-state effects, since from statistical arguments, the Ξ baryon should more likely be produced in a region with more strangeness produced from imbalance in the initial state, but also that the production of a Ξ will result in even more strange hadrons produced during hadronisation. With the amplitude of the near-side $\Xi-\mathrm{K}$ peak being about four times as large as the away-side magnitude (and the difference being even larger if including Λ), it is unlikely that the imbalance is due to the initial state alone, but also due to effects later in the system evolution. The narrow core of the jet peak present in the PYTHIA models is however missing in data, indicating that strangeness is not produced in the final stages either, but rather at some intermediate stage.

PYTHIA and EPOS are extreme cases of where the strangeness production is dominated by final- and initial-state effects, respectively. In the Monash tune of PYTHIA, the near-side $\Xi-\mathrm{K}$ peak is about 70% stronger than what is observed in data, with an even larger excess for total strangeness. Consequently, the peak is much narrower than what is seen in data, and the away-side excess is reduced to about 40% of the ALICE results. The fact that it is non-zero shows that $\mathrm{s\bar{s}}$ breakings may happen early enough in the collision for the hadrons to be emitted back-to-back, which is expected, but strangeness production is still strongly dominated by interactions within the jet. Adding junctions partly remedies the problem by reducing the excess in Λ baryons, as expected when adding mechanisms such as the one shown in Fig 3.24b, but although the away-side excess as well as correlations at large Δy, of oppositely charged kaons get close to what is seen in data, the strangeness in the jet peak is still too dominant. This model also has the problem of producing too little strangeness, which is the reason for adding rope interactions, but the little extra impact that has on the correlation function results in an even larger deviation from data – closer to what is seen for the Monash tune – which disfavours this mechanism for producing strangeness enhancement in pp collisions. Adding string shoving (cf. Section 3.8.3; this was not included in any of the tested models) would probably not resolve this, since it would shove strings apart prior to hadronisation. Since multistrange baryons are produced by strings that are grouped close together, this would rather increase the dilution of the $\Xi-\pi$ correlation function than the correlation to strange hadrons. It would be interesting to test nevertheless, I might well be wrong.

In EPOS LHC on the other hand, strangeness seems to purely be produced at the initial state, resulting in a nearly isotropic excess of kaons. This indicates that the Ξ baryon itself would need to be produced very close to the initial medium and not even close to the jet, which seems unlikely in a physics perspective to begin with, and from these measurements is proven to be wrong. This is rather a limitation of the model than anything else, however, since local conservation of quantum numbers is not implemented in the current version of EPOS (and is not yet implemented in EPOS3 either, but is planned for the future [133]). This does not disprove the core-corona model for the initial state – the missing core of the jet peak indicates that this may very well be the reason for the dilution of strangeness in the event – but the hadronisation description clearly needs a more careful treatment. Until this has been implemented, it is not possible to use these measurements to test whether the EPOS description of the initial state is compatible with data.

The integrated magnitude of the kaon excess on the away side is 0.105±0.003 (stat.)± 0.003 (syst.) kaons per Ξ baryon at small separations in rapidity. The much smaller excess of Λ baryons on the away side indicates that this covers a majority of the total strangeness increase in this region. While being a quite significant number, it is only a small fraction of the total strangeness imbalance of 2 induced by the Ξ baryon. Here one should be careful about comparing numbers though, since the flatness of the away-side excess indicates that it may spread far in rapidity, increasing the total yield significantly. Therefore, this number should rather be compared with the yield on the near side, but due to the strange hadrons that are not experimentally accessible, this is difficult to do for total strangeness.

One can also think of the near-side peak spreading even further in rapidity, which cannot be verified due to the quite narrow rapidity window of the ALICE detector. This would be a signature of elliptic flow and a possible indication of a hydrodynamic

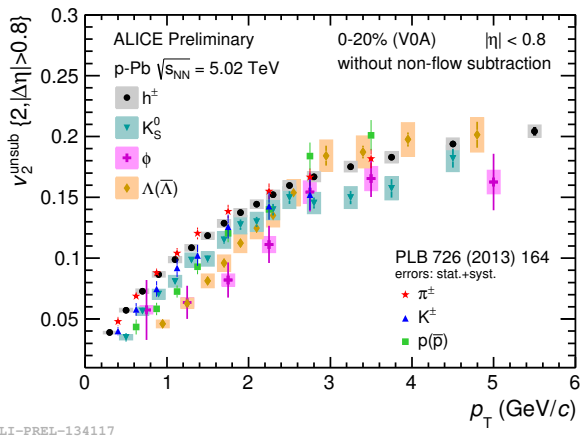

Figure 10.4 – Elliptic flow in p–Pb collisions measured with a rapidity gap of $|\Delta\eta| > 0.8$, but without any other non-flow subtraction for the preliminary results, measured for inclusive charged hadrons as well as a selection of identified particles. Figure taken from Ref. [134].

origin of the strangeness production in the initial state. A more robust measurement for studying this would be to measure flow of multistrange baryons, i.e. Ξ and/or Ω, in small systems, which has not been done yet. The closest measurement carried out so far is a measurement of flow of the Λ baryon and ϕ meson in p–Pb collisions, with only partial non-flow subtraction [134]. The preliminary results, shown in Fig. 10.4, show that the flow compared to inclusive charged hadrons is lower for these species at low p_{T}, which is expected due to mass ordering, but for Λ it is higher at high p_{T} (for ϕ it is consistent with results from other mesons). A similar increase is also observed for protons, and since the effect in Λ may be due to non-flow, it is not yet possible to draw any conclusions from this measurement.

10.2.3 Ξ – p Correlations

Just as for kaons, there is a large excess of protons on the away side, as well as a wide extension of the correlation function in Δy. While the magnitude is smaller than for strangeness, this still indicates that also baryon number is significantly diluted in the system, meaning that the net baryon content decouples already in the initial state. With a large excess also of Λ baryons on the away side, this decoupling is likely significant also for other strange baryons, as well as for neutrons. Given their similar properties (except for the charge), the number of protons and neutrons should be roughly equal, but depending on the mechanism for this decoupling, the number of neutrons may be larger since they contain two d quarks instead of one.

As already has been discussed, the protons seem to take up a larger fraction than the pions of the direct d quark associated with Ξ production at the initial state, although the pions indirectly add to some of the d quark balance due to charge balance. It is however unclear whether the shared $\mathrm{q\overline{q}}$ pair increases the baryon number decoupling, or if it is the baryon number decoupling itself that causes the protons to account for such a large fraction of the away-side d quark excess. One way to get some insight into this, and also would shed some light of the mechanism involved, would be to measure Ξ – n correlations, which unfortunately is impossible with the ALICE detector (and not so easy using any other detector either). A better way to access what fraction of the proton excess is due to baryon decoupling would be to measure Ω – p correlations, since they do not share the same quark flavours, which should be possible in Run 3.

In any case, if the Ξ – n correlations are similar to the Ξ – p correlations, the total baryon imbalance on the away side induced by the Ξ baryon is only slightly lower the strangeness imbalance (cf. Fig. 9.18), given that most of the net strangeness not contributing equally to baryon number comes from kaons (the contribution from Ξ baryons is quite weak, which is probably the case also for Ω given its much lower total yield). With the Ξ baryons having a strangeness-to-baryon number ratio of 2, this indicates that baryon number is even more diluted in the event than strangeness. I leave it to the theorists to figure out what could be causing this.

On the near side, on the other hand, the excess of protons is suppressed compared to e.g. Λ baryons, but the overall shape of it is quite similar to what is observed for kaons. This does not mean that the mechanisms are necessarily similar, though, since here protons (and neutrons) will act mostly as a baryon number balance. With the baryon number conservation appearing quite diluted in the event, it is not too

surprising that the near-side peak is quite small for the proton excess (although not nearly as suppressed as in EPOS). This is very much the case also for PYTHIA, where the Monash tune produces a relatively small excess, which can be explained by the large near-side $\Xi - \Lambda$ peak, which accounts for most of the baryon number balance. Therefore, the total $\Xi-$baryon per-trigger yield is much more enhanced on the near side in PYTHIA than what is seen in data. This is also indicated by previously published results of $\mathrm{p} - \mathrm{p}$ correlations in pp collisions, where the near-side peak of the symmetric correlation function is more prominent than in data, as shown in Fig. 10.5a [132]. The same has been observed also for $\Lambda - \Lambda$ and $\mathrm{p} - \Lambda$ correlations, so this is obviously a universal issue for baryon–baryon correlations.

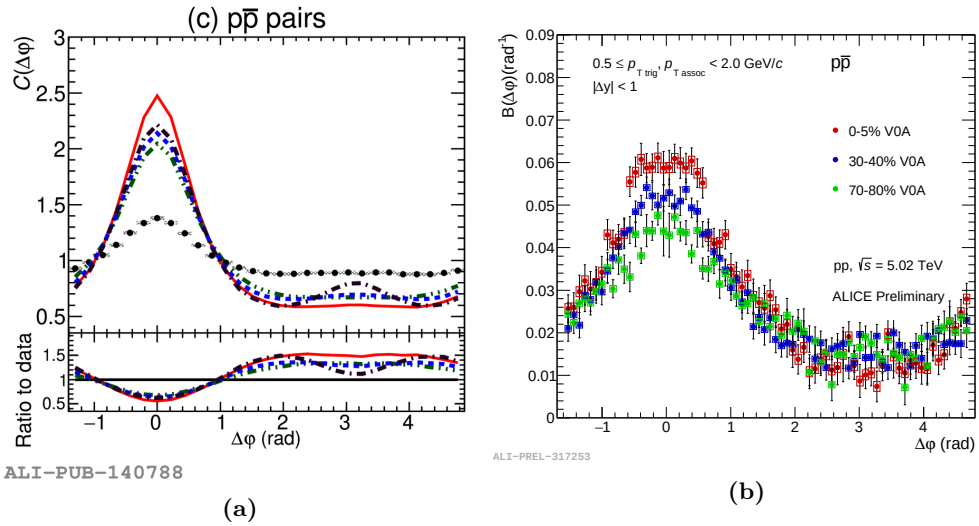

Figure 10.5 – **(a)** Symmetric correlation function (defined by Eq. (8.2)) between pairs of oppositely charged protons, along with model predictions, and **(b)** proton balance functions (differences between opposite-and same-sign particles) for a few multiplicity classes, measured in pp collisions at 7 TeV and 5.02 TeV, respectively. Figures taken from Refs. [132] and [80].

There has not yet been any publication of per-trigger normalised $\mathrm{p} - \mathrm{p}$ correlations by ALICE in small systems, but there has been a measurement of the proton balance function (i.e. the difference between opposite-sign and same-sign correlations, but it is mostly used for pairs of particles of the same kind or inclusive charged particles) in 5.02 TeV pp collisions, aiming at understanding the radial-flow like effects in small systems (cf. Sections 3.7.1 and 3.8.1) which is shown in Fig. 10.5b [80]. Although measured at a lower energy, the results are remarkably similar to the lower panel of Fig. 9.15a. The magnitude is greater (especially considering the narrower p_{T} interval), but this should not come as any surprise, since the proton trigger is not expected to be balanced by strange baryons to the same extent as the Ξ baryon. Moreover, the proton-to-neutron ratio should be greater due to a larger fraction of shared quarks, which could explain the increased magnitude on the away side. These results indicate that the baryon–baryon correlation function is not very different for different triggers.

For the junction and rope extensions, both the shape and magnitude of the near-side difference are very similar to what is observed in data. The reason for the similarity in shape seems to be due to a very weak core of the jet peak, which indicates that most of the correlations in these models are due to processes happening earlier in the evolution, just as in data. While the underlying mechanism might be different (in the PYTHIA models the reason is likely that the baryon balance in the jet peak is covered by other particles such as Λ), this indicates that the processes involved in this particular correlation happens at a similar stage in the evolution. Since strangeness (and total baryon number) seems to decouple only slightly later than this (due to the similar width but more enhanced near-side peak), it could be worth investigating. The similarity in magnitude might be a coincidence though, since the excess of Λ and Ξ baryons is lower in these models, requiring an extra balance of protons (although model dependent), and the total yield of protons is strongly overestimated by these models (cf. Fig. 10.3b).

10.2.4 $\Xi - \Lambda$ and $\Xi - \Xi$ Correlations

For $\Xi - \Lambda$ correlations, most of the excess of opposite-baryon number particles is on the near side, with a larger relative excess compared to the away side than for kaons or protons. The peak magnitude is about 50% larger than for protons, and about half of what is observed for kaons. Therefore, one can conclude that Λ baryons are favoured to be produced together with the Ξ baryon close to jet peak. This could be because here both the strangeness and baryon density are greater in the vicinity of the Ξ baryon, favouring that they hadronise close to each other in phase space, or that the Λ baryon is produced together with the Ξ baryon, as in PYTHIA. What is clearly not observed is a description like in EPOS, where the Λ baryons are produced nearly isotropically with the Ξ baryon in the event, but this may again just be due to quantum number conservation not being handled properly.

Now the question is whether the near-side peak is due to string breakings – like in PYTHIA – or some other mechanism is involved. A few of the results indicate that a pure string model is unlikely. Most importantly, the near-side peak is way too strong in PYTHIA as compared to data, which likely is a consequence of the baryon production mechanism in PYTHIA, cf. Fig 3.24a. With the $\Xi - \Lambda$ near-side excess being about 3.5 times higher in PYTHIA than in data, diquark breaking is unlikely to be the dominant mechanism for baryon production. The junction model reduces this excess to about two times what is seen in data, but one should bear in mind that the diquark breaking mechanism still exists, so this is likely part of the reason why the correlations are so strong. Therefore, baryon production through string junctions is a more plausible mechanism, but it has the drawback of creating more protons than what is observed in data. Nevertheless, it would be interesting to study what impact a pure junction mechanism has on the $\Xi - \Lambda$ correlation function. Adding ropes again moves the near-side yields away from data, which is another indication that this is not the mechanism responsible for strangeness enhancement.

Secondly, while junctions help in remedying the Λ excess, they do not significantly affect the near-side enhancement of opposite-sign kaons. Therefore, both the standard mechanism for baryon production and the junction one will result in a localised strangeness conservation picture, indicating that this is what one obtains from a pure string

model. Finally, this conclusion is consolidated by studying the width of the $\Xi - \Lambda$ peak; it is much wider than in any PYTHIA model, indicating that despite the low excess on the away side, Λ baryons are largely produced by indirect interactions, indicating that part of the near-side production is due to in-medium effects. Whether this is due to some hybrid mechanism or several different mechanisms is a subject for further research, although the multiplicity dependent study indicates that the mechanism(s) does not change with multiplicity.

For $\Xi - \Xi$ correlations, the correlation function is very similar to the one observed for $\Xi - \Lambda$ correlations, but weaker. On the other hand, the relative increase compared to the underlying event is much larger than for Λ, indicating that direct interactions with the trigger particles are important. This is the case also for all phenomenological models.

10.3 Multiplicity Dependent Results

The results from studying $\Xi-$hadron correlations when going from lower to higher multiplicity can be summarised by a few simple observations: the underlying-event correlation becomes stronger, while the part originating from the Ξ trigger remains remarkably similar, but with a slight enhancement in the peak region. These observations are common for all associated particles and there are not any statistically significant effects observed for any of the measurements.

The enhancement of the underlying event is a natural consequence of the increased collision multiplicity – higher-multiplicity events contain more tracks. While there is a bias towards events with a larger number of tracks at midrapidity (the multiplicity is measured at forward rapidities, and in pp correlations, the correlation is not as strong between the two regions as in heavy-ion collisions), this division is still apparent. Given the small correlation between pairs of Ξ baryons, the average number of Ξ baryons per event containing a Ξ is typically one, so the per-trigger yields are expected to scale with the number of tracks. This also tells us that the enhanced strangeness in high-multiplicity events is due to a higher probability that a Ξ baryon is formed in an event, and not that they are much more likely to form together.

That the magnitude of the balancing part of the correlation function is similar across multiplicity is also expected due to conservation of quantum numbers, so the interesting result is that the shape does not change. This tells us that the production mechanism is the same for high- and low-multiplicity events. The slight increase in peak amplitude at high multiplicity may be attributed to radial flow (cf. Section 3.7.1), since such an increase should be compensated by a depletion elsewhere in phase space, and hence the balance function should be narrower. There are indications of such a narrowing at least in $\Xi - \pi$ correlations, but in order to tell whether this is statistically significant, a careful quantification would be necessary, which there was unfortunately not enough time for here. Moreover, the narrowing is expected to be dominant at low p_{T}, so one should limit the study to low-p_{T} triggers to study this and ideally go lower in p_{T} than $1.2\,\mathrm{GeV}/c$, which is the low limit currently used[4].

[4]This limit was selected since there were large statistical fluctuations at lower p_{T}, which were distorting the results given the statistical sample used originally, but given the larger sample used in

In the case of a core-corona picture, this observation would indicate that the volume where strange quarks are produced is smaller at low multiplicity, but they would still be produced in the core part.

10.4 Conclusions

These results indicate that both strangeness and baryon number largely decouple during the initial phase, indicating that the initial state largely behaves like a medium, although what could be causing the baryon number decoupling is somewhat unclear. This is further indicated by radial-flow like effects in the scaling with multiplicity. This is consistent with a core-corona picture, but since EPOS does not conserve quantum numbers locally, it is not possible to test this model until these conservation laws are properly implemented in EPOS.

There are also strong final-state interactions resulting in a significant excess of opposite quantum numbers in the vicinity of the Ξ baryon, which in particular is seen as an excess of strange baryons. In regions of phase space where the baryon and strangeness density are lower, production of protons and kaons, i.e. non-strange baryons and strange mesons, is more favoured. The hadronisation mechanism could partly be due to a string breaking mechanism, and given the success of PYTHIA in describing many other observables, parts of this model may be close to reality. There are however indications of medium effects involved also here, which could either originate from the initial state or from some hybrid mechanism. Baryon production from diquark formation – the main mechanism in PYTHIA – and strangeness production from colour ropes are disfavoured to be among the dominant mechanisms for these processes.

The strangeness production mechanism does not seem to change with collision multiplicity.

The correlations between the Ξ baryon with same-quantum number particles are dominated by the underlying event and therefore probe different physics. These largely behave as described in PYTHIA.

To summarise, these measurements indicate that there is a possibility of a core-corona initial state and a hadronisation phase similar to what is observed in PYTHIA (with junctions), rather than a pure string based model. Therefore, such an approach should be considered some further phenomenological treatment.

10.5 Outlook

The following studies are proposed as an extension of this analysis:

- Measurements of p_{T} dependent $\Xi-$hadron correlations. This was originally planned for this thesis, but has not yet been done due to time constraints. This would test whether the correlations are different between hard and soft processes, and more conclusively test whether the correlation function is subject to radial flow.
- Ω-hadron correlation studies. The analysis presented here could be repeated with the Ω baryon as a trigger. Being a pure strange-quark state, this would reduce

the final results, this should now be possible to lift.

contamination from other processes, such as those induced by the down quark in Ξ. Therefore, $\Omega - \mathrm{p}$ correlations could be used to purely access the baryon correlation with the away side and $\Omega - \mathrm{K}$ correlations would more conclusively quantify how strangeness is distributed in the event. Especially interesting would be to compare $\Omega - \Lambda$ correlations with $\Xi - \Lambda$ correlations, since this would probe whether the enhancement of Λ baryons on the near side is mostly due to diquark sharing or if it is dominated by the high density of both strangeness and baryons. Such a study will be possible following in Run 3, following the detector upgrade, where a specific trigger for events with Ω baryons is planned (this will trigger on topology, so will also yield much more Ξ baryons) [106]. With the Ω baryon being about 20 times as rare as the Ξ baryon, only simpler studies such as minimum-bias $\Omega - \pi$ and $\Omega - \mathrm{K}$ correlations are currently accessible.

- More complete Ξ correlation studies in PYTHIA and its extensions. Since this is simulation, one can study interactions also with particles that cannot easily be detected in the experiment (including $\Sigma^{\pm}$). In this way one can study how the correlation between Ξ and total charge, baryon number, net strangeness etc. are related to each other, and how this changes by turning on and off different mechanisms such as colour reconnection (included in the Monash tune), junctions, and string shoving. For a more advanced study, one can tune various parameters to quantify the contribution from each mechanism to the correlation function.

- In-depth studies in PYTHIA on how the correlations between Ξ and other hadrons are produced. This means that one should dig into the various processes involved, which would support the discussion in this thesis. It would be particularly interesting to investigate what happens at the different stages of the evolution and how this enters into the correlation function, since this would give insights into the time scales of strangeness and baryon decoupling and charge balance etc. also in data.

- Extensions to larger collision systems such as p–Pb and Pb–Pb collisions (and intermediate systems possibly available in Run 3 and beyond). Such measurements would more robustly test how the mechanism for strangeness production depends on system size and collision multiplicity than what is possible from pp collisions alone.

- Measurements of $\mathrm{p} - \mathrm{p}$, $\mathrm{p} - \Lambda$, and $\Lambda - \Lambda$ correlations in data and various PYTHIA extensions. By comparing the results from such a study with the ones presented in this thesis, one can e.g. test if baryon number is distributed in a similar way when triggering on a proton or Λ baryon as when triggering on a Ξ baryons. In data, only symmetric correlations have been published for these observables so far, so for a direct comparison with this study, one needs to extend this to per-trigger yields.

- Theoretical development of models with a core-corona initial state and a hadronisation phase where all quantum numbers are locally conserved. This may be implemented in a future EPOS version, but given the success of PYTHIA, one can also think of corona based models partly based on PYTHIA. Such a model has

actually already been studied, where PYTHIA/Angantyr was used to generate particles from pp, p–Pb, and Pb–Pb collisions and a hydronamisation code was applied to generate a fluid state in regions with high parton density, whereas the PYTHIA string fragmentation was used for regions with lower density [135]. In this way, a core-corona state was produced, resulting in an increase in strange baryons with multiplicity. From the results presented in this thesis, it seems more likely that the core has a different origin (although this picture is a possible alternative if not believing that the CGC picture is valid in small systems) and the string picture is only valid in the hadronisation phase. Nevertheless, if an event generator based on this mechanism is released, it would be interesting to test it against the results presented in this thesis.

Part IV

Appendices

Appendix A

Mathematical Derivations

A.1 Derivation of Existence of a Local Inverse of the Distribution of the ADC Output

This mathematical derivation complements the work presented in Chapter 6. According to the Inverse Function Theorem, a function $\mathbf{g}(\mathbf{x})$ is locally invertible at a point $\mathbf{x_0}$ if the Jacobian

$$\mathcal{J} \equiv \det \frac{\mathrm{d}\mathbf{g}}{\mathrm{d}\mathbf{x}}$$

is non-zero at $\mathbf{x_0}$ [136, p. 3]. For a global inverse, a necessary (but not sufficient) condition is that $\mathcal{J}$ is non-zero over the entire domain of $\mathbf{g}$. In this analysis, the relevant function $\mathbf{g} : (\mathbb{R} \times \mathbb{R}_+) \mapsto (\mathbb{R} \times \mathbb{R}_+)$ is defined as (cf. Section 6.3.2)

$$\mathbf{g}(\mu,\sigma) = \begin{pmatrix} \mu_D(\mu,\sigma) \\ \sigma_D(\mu,\sigma) \end{pmatrix} = \begin{pmatrix} \sum_{n=-\infty}^{\infty} \frac{1}{2} n \left(\operatorname{erf}\left(\frac{n+1-\mu+\Delta_{(n+1)}}{\sigma\sqrt{2}} \right) - \operatorname{erf}\left(\frac{n-\mu+\Delta_{(n)}}{\sigma\sqrt{2}} \right) \right) \\ \sqrt{\sum_{n=-\infty}^{\infty} \frac{1}{2} (n-\mu_D(\mu,\sigma))^2 \left(\operatorname{erf}\left(\frac{n+1-\mu+\Delta_{(n+1)}}{\sigma\sqrt{2}} \right) - \operatorname{erf}\left(\frac{n-\mu+\Delta_{(n)}}{\sigma\sqrt{2}} \right) \right)} \end{pmatrix}.$$

For this function, after some simplifications one obtain

$$\frac{\partial \mathbf{g}}{\partial \mu} = \begin{pmatrix} \sum_{n=-\infty}^{\infty} \frac{1}{\sigma\sqrt{2\pi}} \exp\left(\frac{(n-\mu+\Delta_{(n)})^2}{2\sigma^2} \right) \\ \frac{1}{2\sigma_D(\mu,\sigma)\sigma\sqrt{2\pi}} \sum_{n=-\infty}^{\infty} \left(2n + \frac{1}{2} + \frac{3}{2} \sum_{m=-\infty}^{\infty} \operatorname{erf}\left(\frac{m-\mu+\Delta_{(m)}}{\sigma\sqrt{2}} \right) \right) \cdot \\ \exp\left(\frac{(n-\mu+\Delta_{(n)})^2}{2\sigma^2} \right) \end{pmatrix},$$

and

$$\frac{\partial \mathbf{g}}{\partial \sigma} = \begin{pmatrix} \sum_{n=-\infty}^{\infty} \frac{n-\mu+\Delta_{(n)}}{\sigma^2\sqrt{2\pi}} \exp\left(\frac{(n-\mu+\Delta_{(n)})^2}{2\sigma^2}\right) \\ \\ \frac{1}{2\sigma_D(\mu,\sigma)\sigma^2\sqrt{2\pi}} \sum_{n=-\infty}^{\infty}(n-\mu+\Delta_{(n)})\cdot \\ \left(2n+\frac{1}{2}+\frac{3}{2}\sum_{m=-\infty}^{\infty}\operatorname{erf}\left(\frac{m-\mu+\Delta_{(m)}}{\sigma\sqrt{2}}\right)\right)\exp\left(\frac{(n-\mu+\Delta_{(n)})^2}{2\sigma^2}\right) \end{pmatrix}.$$

Therefore,

$$\begin{aligned} \mathcal{J} = \begin{vmatrix} \frac{\partial \mu_D}{\partial \mu} & \frac{\partial \mu_D}{\partial \sigma} \\ \frac{\partial \sigma_D}{\partial \mu} & \frac{\partial \sigma_D}{\partial \sigma} \end{vmatrix} &= \frac{1}{4\pi\sigma_D\sigma^3}\left(\left(\sum_{n=-\infty}^{\infty} n\exp\left(\frac{(n-\mu+\Delta_{(n)})^2}{2\sigma^2}\right)\right)^2\right. \\ &\left. - \sum_{n=-\infty}^{\infty} n^2\exp\left(\frac{(n-\mu+\Delta_{(n)})^2}{2\sigma^2}\right)\sum_{m=-\infty}^{\infty}\exp\left(\frac{(m-\mu+\Delta_{(m)})^2}{2\sigma^2}\right)\right) \\ &\propto \mathbb{E}_Y[n^2] - (\mathbb{E}_Y[n])^2 = \mathbb{V}_Y[n], \end{aligned}$$

where Y is a distribution created from a discrete sampling of a normal distribution with mean μ and standard deviation σ. Thus,

$$\mathcal{J} = 0 \iff \mathbb{V}_Y[n] = 0 \iff \sigma = 0,$$

i.e. $\mathbf{g}$ is locally invertible whenever $\sigma > 0$.

A.2 Calculation of Sideband Coefficients Used for Λ in the Non-Linear Regime

This section describes how the scaling coefficients applied to the non-linear sideband regime in the Λ invariant-mass spectrum are calculated. The final coefficents are listed at the end of the section.

Let y_1 and y_2 be the integrals of the background in the lower and upper sidebands, respectively, and x_{L} and x_{NL} be the linear and non-linear components of the lower sideband, i.e. $y_1 = x_{\mathrm{L}} + x_{\mathrm{NL}}$. Then we can find coefficients a, b such that

$$\begin{cases} x_{\mathrm{L}} + & x_{\mathrm{NL}} = & y_1 \\ a\,x_{\mathrm{L}} + & b\,x_{\mathrm{NL}} = & y_2. \end{cases} \tag{A.1}$$

In order for the sideband subtraction to be correct, one needs to subtract the terms $(1+a)x_{\mathrm{L}} + f(1+b)x_{\mathrm{NL}}$ from the signal, where f is the scaling factor required for the non-linear part (the first term is just the usual linear sideband subtraction). We want to find scaling terms f_1 and f_2 for each sideband, such that

$$f_1 y_1 + f_2 y_2 = (1+a)x_{\mathrm{L}} + f(1+b)x_{\mathrm{NL}} \tag{A.2}$$

is satisfied.

By inverting Eq. (A.1), one obtains

$$\begin{cases} x_{\mathrm{L}} & = \dfrac{by_1 - y_2}{b-a} \\ x_{\mathrm{NL}} & = \dfrac{y_2 - ay_1}{b-a}. \end{cases}$$

Inserting this into Eq. (A.2), yields after some simplification

$$\begin{aligned} f_1 &= \frac{b - f\cdot a + (1-f)\cdot a\cdot b}{b-a} \\ f_2 &= \frac{f - 1 - a + f\cdot b}{b-a}. \end{aligned} \tag{A.3}$$

Now remains to evaluate a, b, and f, which is done by integrating various components of the fit function. By introducing $m = \Delta M_{\mathrm{inv}}$, the sidebands are defined as $-7\sigma < m - \mu \leq -4\sigma$ for the lower sideband and $4\sigma < m - \mu \leq 7\sigma$ for the upper sideband, and the signal region as $-3\sigma < m - \mu \leq 3\sigma$, where σ and μ are obtained from the fits shown in Fig. 8.7. The parameter a is the ratio of the linear component in the upper sideband to the one in the lower sideband, i.e.

$$a = \frac{\int_{\mu+4\sigma}^{mu+7\sigma} d_0 + d_1 m \mathrm{d}m}{\int_{\mu-7\sigma}^{\mu-4\sigma} d_0 + d_1 m \mathrm{d}m} = \frac{d_0 + d_1(\mu + 5.5\sigma)}{d_0 + d_1(\mu - 5.5\sigma)},$$

where d_0 and d_1 are the fit parameters in the linear component of the fit function. To evaluate b and f, let us first integrate the non-linear component. Its indefinite integral is

$$\begin{aligned} &I(m) = \\ &\int c_1 \cdot \exp\left(-4\left(\frac{m-c_2}{c_3} - 1\right)\right) \cdot \left(\frac{m-c_2}{c_3}\right)^4 \mathrm{d}m = -4c_1c_3 \exp\left(4\left(1 - \frac{m-c_2}{c_3}\right)\right) \\ &\cdot \left(\left(\frac{m-c_2}{c_3}\right)^4 + \left(\frac{m-c_2}{c_3}\right)^3 + \frac{3}{4}\left(\frac{m-c_2}{c_3}\right)^2 + \frac{3}{8}\left(\frac{m-c_2}{c_3}\right) + \frac{3}{32}\right) + C_1, \end{aligned}$$

where C_1 is an integration constant and the parameters $c_1 - c_3$ are obtained from the fit. Then the integrals in the various regions become

$$\begin{aligned} I_1 &= I(\mu - 4\sigma) - I(\mu - 7\sigma), \\ I_2 &= I(\mu + 7\sigma) - I(\mu + 4\sigma), \\ I_3 &= I(\mu + 3\sigma) - I(\mu - 3\sigma), \end{aligned}$$

where I_1 and I_2 are the integrals in the lower and upper sideband regions, respectively, and I_3 is the integral of the non-linear part of the background in the signal region. From these integrals,

$$b = \frac{I_2}{I_1}$$

and

$$f = \frac{I_3}{I_1 + I_2}.$$

The final scaling factors defined by Eq. (A.3) are summarised in Table A.1 for ALICE data and in Table A.2 for MC.

Table A.1 – Multiplication factors used for the lower and upper sidebands for Λ in the interval $0.4 < p_{\mathrm{T}} < 1.0\,\mathrm{GeV}/c$ for the ALICE data. Note that the lowest bin is not used for $\overline{\Lambda}$.

p_{T} interval	Λ		$\overline{\Lambda}$	
	lower sideband	upper sideband	lower sideband	upper sideband
$0.4 - 0.5\,\mathrm{GeV}/c$	1.80	0.56	–	–
$0.5 - 0.6\,\mathrm{GeV}/c$	1.12	1.14	0.67	1.57
$0.6 - 0.7\,\mathrm{GeV}/c$	0.32	1.87	0.37	1.82
$0.7 - 0.8\,\mathrm{GeV}/c$	0.32	1.88	0.33	1.89
$0.8 - 0.9\,\mathrm{GeV}/c$	0.27	2.07	0.27	2.09
$0.9 - 1.0\,\mathrm{GeV}/c$	0.01	2.61	0.06	2.51

Table A.2 – Multiplication factors used for the lower and upper sidebands for Λ in the interval $0.4 < p_{\mathrm{T}} < 1.0\,\mathrm{GeV}/c$ for the MC data set.

p_{T} interval	Λ		$\overline{\Lambda}$	
	lower sideband	upper sideband	lower sideband	upper sideband
$0.4 - 0.5\,\mathrm{GeV}/c$	1.73	0.64	–	–
$0.5 - 0.6\,\mathrm{GeV}/c$	1.21	0.84	1.20	0.78
$0.6 - 0.7\,\mathrm{GeV}/c$	0.68	1.39	0.63	1.44
$0.7 - 0.8\,\mathrm{GeV}/c$	0.20	2.15	0.20	2.13
$0.8 - 0.9\,\mathrm{GeV}/c$	0.14	2.32	0.16	2.24
$0.9 - 1.0\,\mathrm{GeV}/c$	0.01	2.59	0.01	2.58

Appendix B

List of Acronyms

A side	ATLAS side
ADC	Analogue-to Digital Converter
ALICE	A Large Ion Collider Experiment
ALTRO	ALICE TPC Read-Out
AOD	Analysis Oriented Data
ATLAS	A Toroidal Large ApparatuS
BGA	Ball Grid Array
BM@N	Baryonic Matter at Nuclotron
BR	Branching Ratio
C side	CMS side
CBM	Compressed Baryonic Matter
CGC	Color Glass Condensate
CERN	European Organization for Nuclear Research[1]
CKM matrix	Cabibbo-Kobayashi-Maskawa matrix
CMS	Compact Muon Solenoid
CP	Charge-Parity
CPV	Charged Particle Veto
CR	Colour Reconnection
CSA	Charge Sensitive Amplifier
CTF	Compressed Time Frame
CTP	Central Trigger Processor
DAQ	Data AQuisition
DC	Direct Current
DCA	Distance of Closest Approach
DCal	Di-jet Calorimeter
DPG	Data Preparation Group
DIPSY	Dipoles in Impact-Parameter Space and rapidity (Y)
EKRT model	Eskola-Kajantie-Ruuskanen-Tuominen model

[1] Originally an abbreviation for *Conceil Européen pour la Recherche Nucléaire*, the research council that was formed when it was decided to build an international nuclear physics laboratory in Europe.

EMCal	Electromagnetic Calorimeter
EPOS	Energy-conserving quantum mechanical multiple scattering approach, based on Partons (parton ladders), Off-shell remnants, and Splitting of parton ladders
ENC	Equivalent Noise Charge
ESD	Event Summary Data
FAIR	Facility for Antiproton and Ion Research
FPGA	Field-Programmable Gate Array
GDC	Global Data Concentrator
GEANT	GEometry ANd Tracking
GEM	Gas Electron Multiplier
HLT	High-Level Trigger
HM	High Multiplicity
HMPID	High-Momentum Particle IDentification detector
HRG	Hadron Resonance Gas
IROC	Inner Read-Out Chamber
IS	Initial State
ITS	Inner Tracking System
JTAG	Joint Test Action Group
LDC	Local Data Concentrator
LEP	Large Electron-Positron collider
LHC	Large Hadron Collider
LHCb	LHC beauty
MB	Minimum Bias
MC	Monte Carlo
MCH	Muon tracking CHamber
MC-KLN	MC implementation of the Kharzeev-Levin-Nardi model
MOS	Metal Oxide Semiconductor
MPI	MultiParton Interaction
MRPC	Multi-gap Resistive-Plate Chamber
MTR	Muon TRigger
MWPC	Multi-Wire Proportional Chamber
NCCA	Negatively polarised Circuit Card Assembly
NICA	Nuclotron-based Ion Collider fAcility
NLO	Next-to-Leading-Order
O^2	Offline-Online system
OROC	Outer Read-Out Chamber
PA	Pointing Angle
PASA	PreAmplifier/Shaping Amplifier
PHENIX	Pioneering High Energy Nuclear Interaction eXperiment
PCCA	Positively polarised Circuit Card Assembly
PID	Particle IDentification
PV	Primary Vertex
QCD	Quantum ChromoDynamics
QED	Quantum ElectroDynamics

QGP	Quark Gluon Plasma
RAM	Random-Access Memory
RHIC	Relativistic Heavy-Ion Collider
SDD	Silicon Drift Detector
SoCKit	System-on-Chip development Kit
SPD	Silicon Pixel Detector
sPHENIX	super PHENIX
SPS	Super Proton Synchrotron
SSD	Silicon Strip Detector
STAR	Solenoidal Tracker At RHIC
SV	Secondary Vertex
T0	Time 0 detector
TOF	Time-Of-Flight
TPC	Time Projection Chamber
TRD	Transition Radiation Detector
T_RENTo	Reduced Thickness Event-by-event Nuclear Topology model
TV	Tertiary Vertex
UV	Ultra-Violet
VZERO/V0	Vertex 0 detector
ZDC	Zero-Degree Calorimeter

Appendix C

List of Common Hadrons

Table C.1 summarises some of the basic properties of the hadrons involved in the analyses presented in this thesis (excluding antiparticles), along with a few other light hadrons, which either are used in related analysis, or have similar properties, but are not analysed due to detector limitations. Only charged particles give tracks in the detector. In order to measure the momenta of neutral particles, one requires a calorimeter. ALICE only has electromagnetic calorimeters, which enable photon detection, but such measurements are challenging due to the large background involved. Neutrons, for instance, require a hadronic calorimeter, and are thus not possible to detect by the ALICE detector. Alternatively, particles which decay into one charged and one neutral particle, it is possible to detect the momentum kink of the charged track, but this requires a higher tracking resolution in the vicinity of the beam pipe than what is currently available, but may be possible after the ITS upgrade (this is currently being evaluated). For the particles that can be more easily identified, they can either be detected directly, through identification of the decay topology (V^0 or cascade, cf. Section 8.5.2), or – in the case of resonances – through their invariant mass alone.

The proper lifetime listed in the table is multiplied by the speed of light c, since this is close to the decay length of a particle with high energy in its rest frame. In the lab frame, the particles may travel a significantly larger distance due to the Lorentz boost. Individual particles may travel up to a few decay lengths due to the statistical nature of the decay. The weakly decaying particles listed all have a decay length $> 1\,\mathrm{cm}$.

Table C.1 – Summary of some of the most important properties of a selection of light hadrons. Data retreived from Ref. [11].

Particle	name	quark content	main decay modes	detection method	proper lifetime ($c\tau$)	mass ($\times c^2$)
π^+	charged pion	$\mathrm{u\bar{d}}$	$\mu^+\nu_\mu$ ($\sim 100\%$)	direct tracking	7.80 m	140 MeV
π^0	neutral pion	$\frac{1}{\sqrt{2}}(\mathrm{u\bar{u}} - \mathrm{d\bar{d}})$	2γ (99%)	EM cal. + inv. mass	25.5 nm	135 MeV

Particle	name	quark content	main decay modes	detection method	proper lifetime ($c\tau$)	mass ($\times c^2$)
K^+	charged kaon	$u\bar{s}$	$\mu^+\nu_\mu$ (64%) $\pi^+\pi^0$ (21%)	direct tracking	3.71 m	494 MeV
K^0_s	neutral kaon, short	$\frac{1}{\sqrt{2}}(d\bar{s}+\bar{d}s)^a$	$\pi^+\pi^-$ (69%) $2\pi^0$ (31%)	V^0 topology	2.68 cm	498 MeV
K^0_L	neutral kaon, long	$\frac{1}{\sqrt{2}}(d\bar{s}-\bar{d}s)^a$	$\pi^\pm e^\mp \nu_e$ (41%) $\pi^\pm \mu^\mp \nu_\mu$ (27%) $3\pi^0$ (20%) $\pi^+\pi^-\pi^0$ (13%)	hadronic calorimetry	15.3 m	498 MeV
ϕ^0	ϕ meson	$s\bar{s}$	K^+K^- (49%) $K^0_s K^0_L$ (34%) $\pi^+\pi^-\pi^0$ (15%)	invariant mass	46.3 fm	1019 MeV
p	proton	uud	stable	direct tracking	N/A	938 MeV
n	neutron	udd	$pe^-\overline{\nu_e}$ (100%)	hadronic calorimetry	$2.6 \cdot 10^8$ km	940 MeV
Λ^0	Λ baryon	uds	$p\pi^-$ (64%) $n\pi^0$ (36%)	V^0 topology	7.89 cm	1116 MeV
Σ^+	charged Σ baryon	uus	$p\pi^0$ (52%) $n\pi^+$ (48%)	calorimeter + inv. mass / track kink	2.40 cm	1189 MeV
Σ^0	neutral Σ baryon	uds	$\Lambda^0\gamma$ ($\sim 100\%$)	EM cal. + inv. mass	22.2 pm	1193 MeV
Σ^-	charged Σ baryon	dds	$n\pi^-$ ($\sim 100\%$)	had. cal. + inv. mass / track kink	4.43 cm	1197 MeV
Ξ^0	neutral Ξ baryon	uss	$\Lambda^0\pi^0$ ($\sim 100\%$)	EM cal. + inv. mass	8.71 cm	1315 MeV
Ξ^-	charged Ξ baryon	dss	$\Lambda^0\pi^-$ ($\sim 100\%$)	cascade topology	4.91 cm	1322 MeV
Ω^-	Ω baryon	sss	$\Lambda^0 K^-$ (68%) $\Xi^0\pi^-$ (24%) $\Xi^-\pi^0$ (8.6%)	cascade topology	2.46 cm	1672 MeV

[a] Neglecting CP violation.

Appendix D

Complementary Figures

D.1 Complementary Results From the MCH Noise Measurement

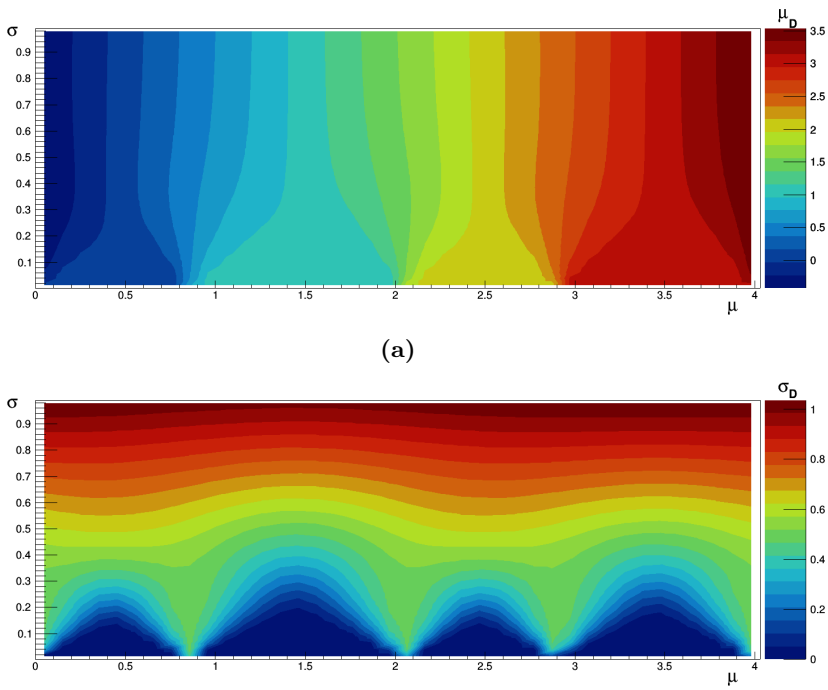

(a)

(b)

Figure D.1 – Discrete mean μ_D **(a)** and σ_D **(b)** as a function of analogue mean μ ($0 \leq \mu < 4$) and standard deviation σ, for $\Delta_1 = -0.15$ and $\Delta_2 = 0.05$.

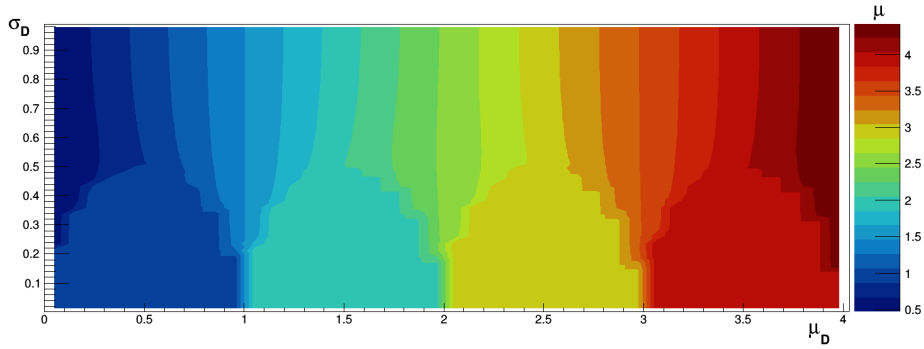

(a)

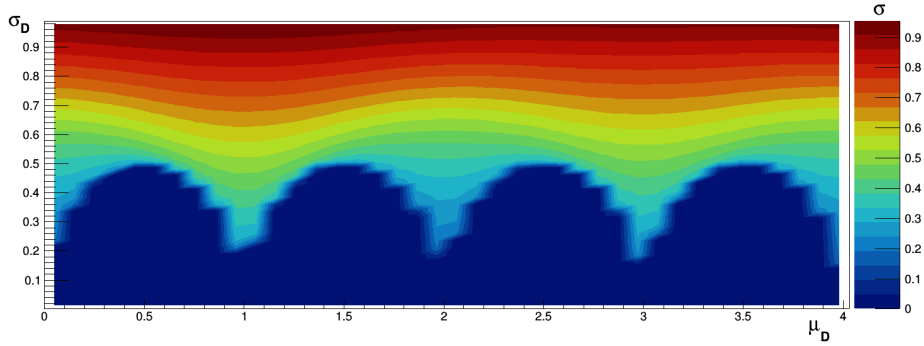

(b)

Figure D.2 – Inverse to Fig. D.1, i.e. analogue mean μ **(a)** and σ **(b)** as a function of discrete mean μ_D ($0 \leq \mu_D < 4$) and standard deviation σ_D, for $\Delta_1 = -0.15$ and $\Delta_2 = 0.05$.

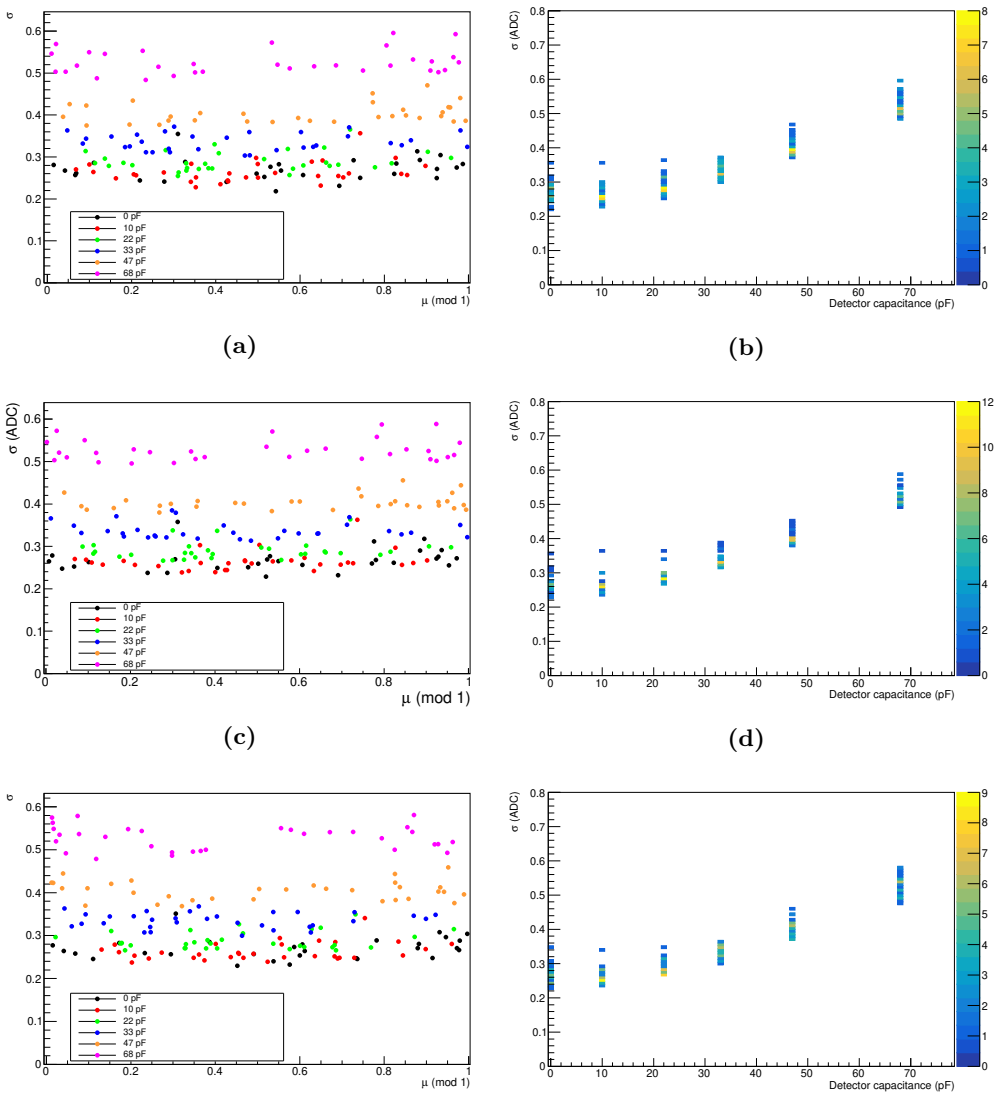

Figure D.3 – Left: Extracted noise σ versus truncated mean μ_0 for different input capacitances. Right: Spread in σ as a function of detector capacitance. **(a)**-**(b)** Hypothesis 2, **(c)**-**(d)** Hypothesis 3, and **(e)**-**(f)** Hypothesis 4, defined in Section 6.5.1.

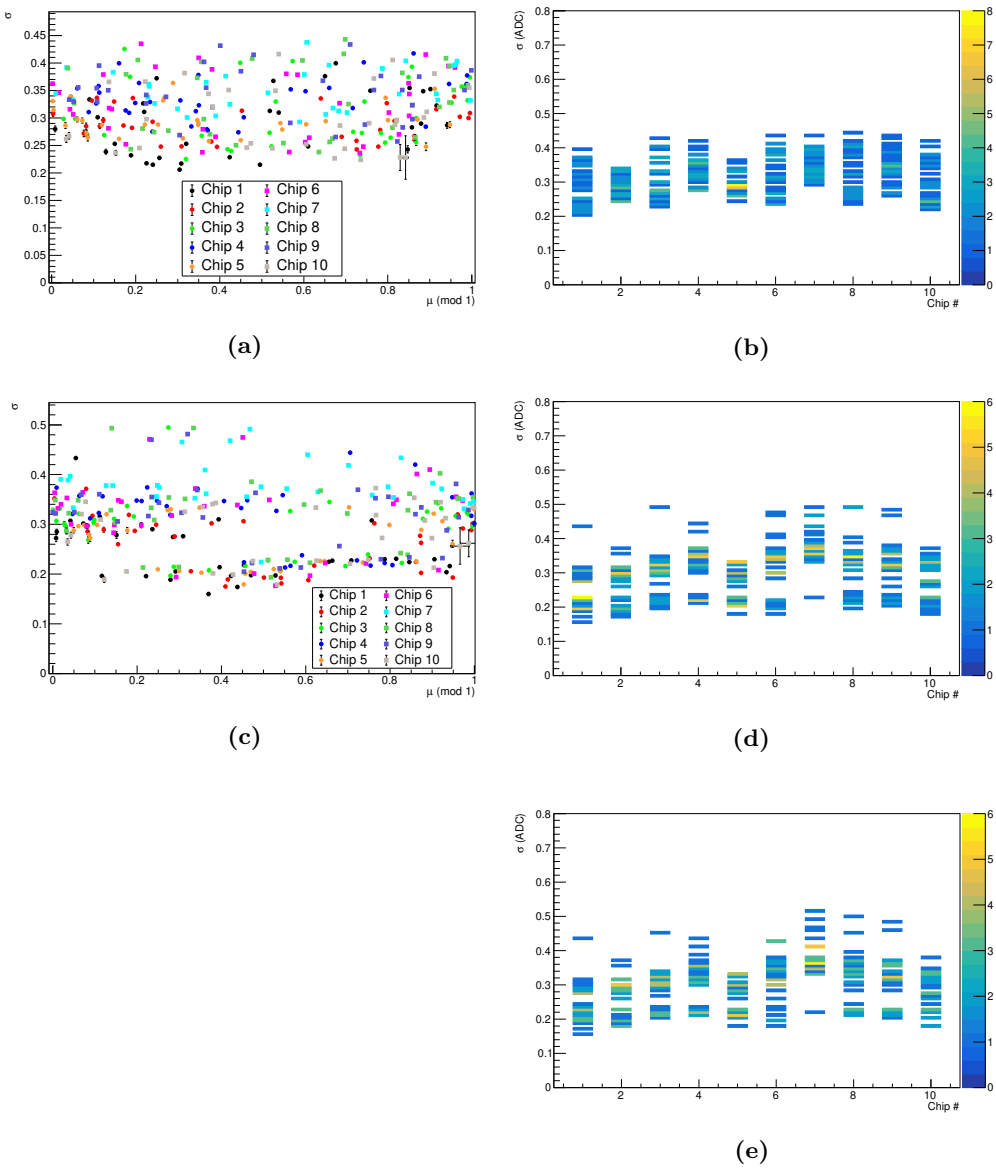

Figure D.4 – Left: Extracted noise σ versus truncated mean μ_0 for 10 different chips from the automatic testing, using a gain of 4 mV/fC. Right: Spread in σ divided into different chips. **(a)-(b)** Hypothesis 2, **(c)-(d)** Hypothesis 3, and **(e)** Hypothesis 5, defined in Section 6.5.1.

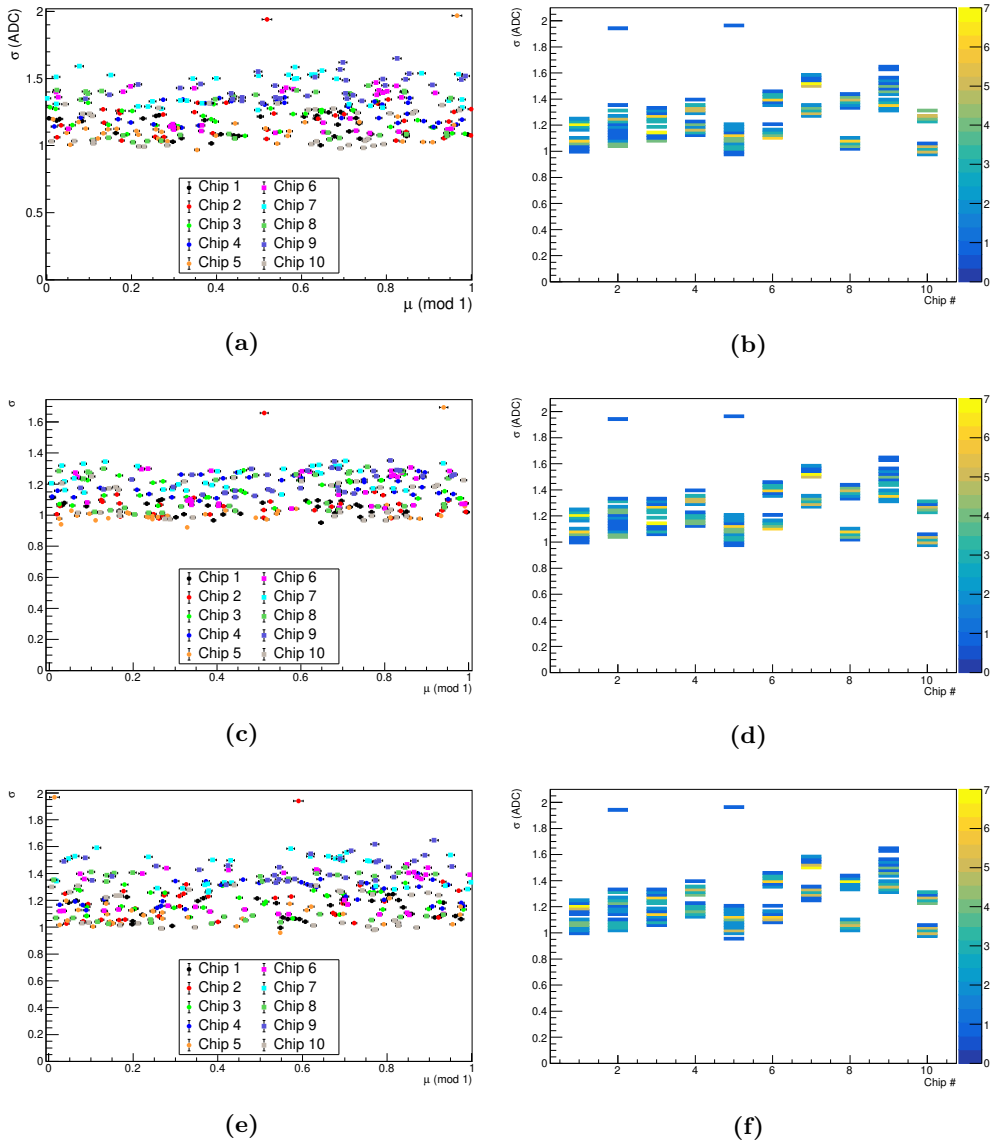

Figure D.5 – Left: Extracted noise σ versus truncated mean μ_0 for 10 different chips from the automatic testing, using a gain of 20 mV/fC. Right: Spread in σ divided into different chips. **(a)**-**(b)** Hypothesis 1, **(c)**-**(d)** Hypothesis 2, and **(e)**-**(f)** Hypothesis 3, defined in Section 6.5.1.

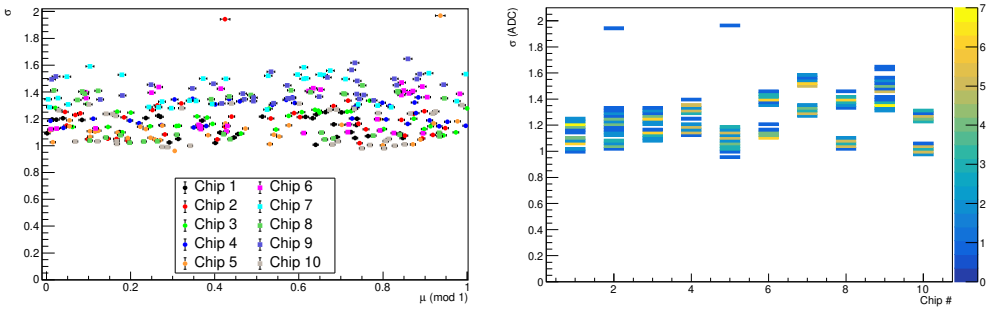

(a) **(b)**

Figure D.6 – **(a)** Extracted noise σ versus truncated mean μ_0 for Hypothesis 4 defined in Section 6.5.1, for 10 different chips from the automatic testing, using a gain of 20 mV/fC. **(b)** Spread in σ divided into the different chips (for the same hypothesis).

D.2 Supporting Figures to the Ξ−Hadron Analysis Description

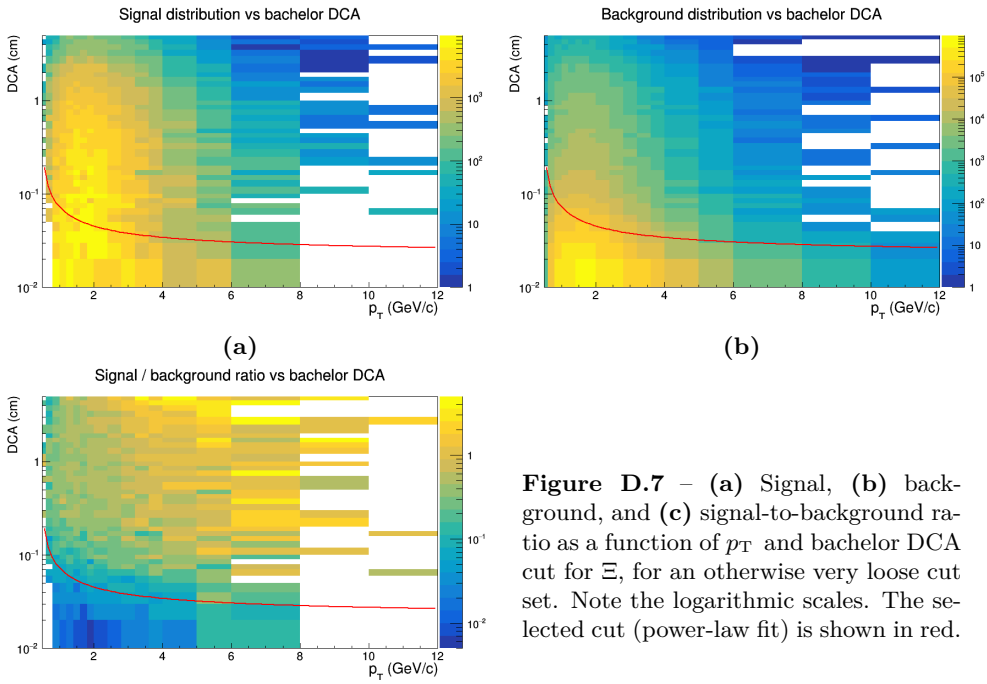

(a) **(b)**

(c)

Figure D.7 – **(a)** Signal, **(b)** background, and **(c)** signal-to-background ratio as a function of p_{T} and bachelor DCA cut for Ξ, for an otherwise very loose cut set. Note the logarithmic scales. The selected cut (power-law fit) is shown in red.

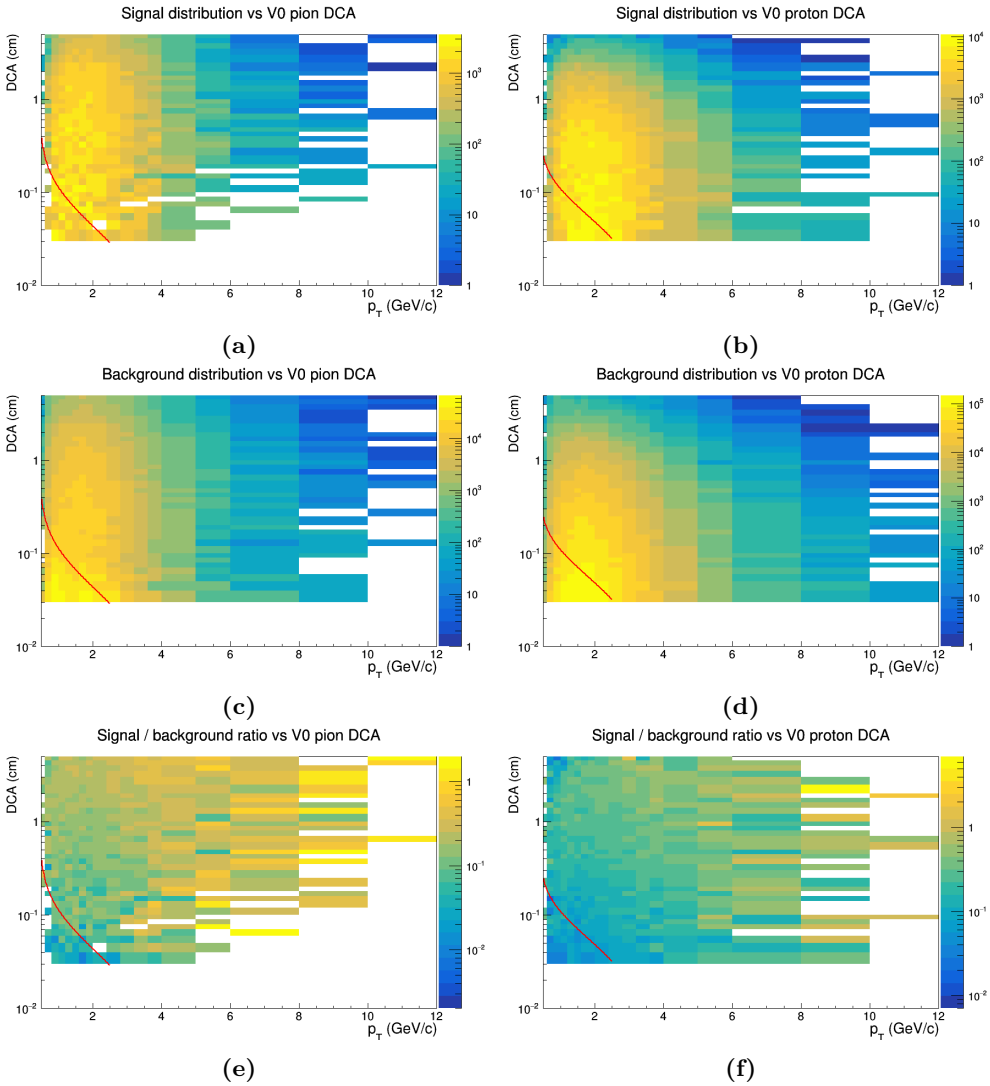

Figure D.8 – **(a)-(b)** Signal, **(c)-(d)** background, and **(e)-(f)** signal-to-background ratio as a function of p_{T} and **(left)** V^0 pion and **(right)** V^0 proton daughter DCA cut for Ξ, for an otherwise very loose cut set. The selected cuts (power-law fit) are shown in red.

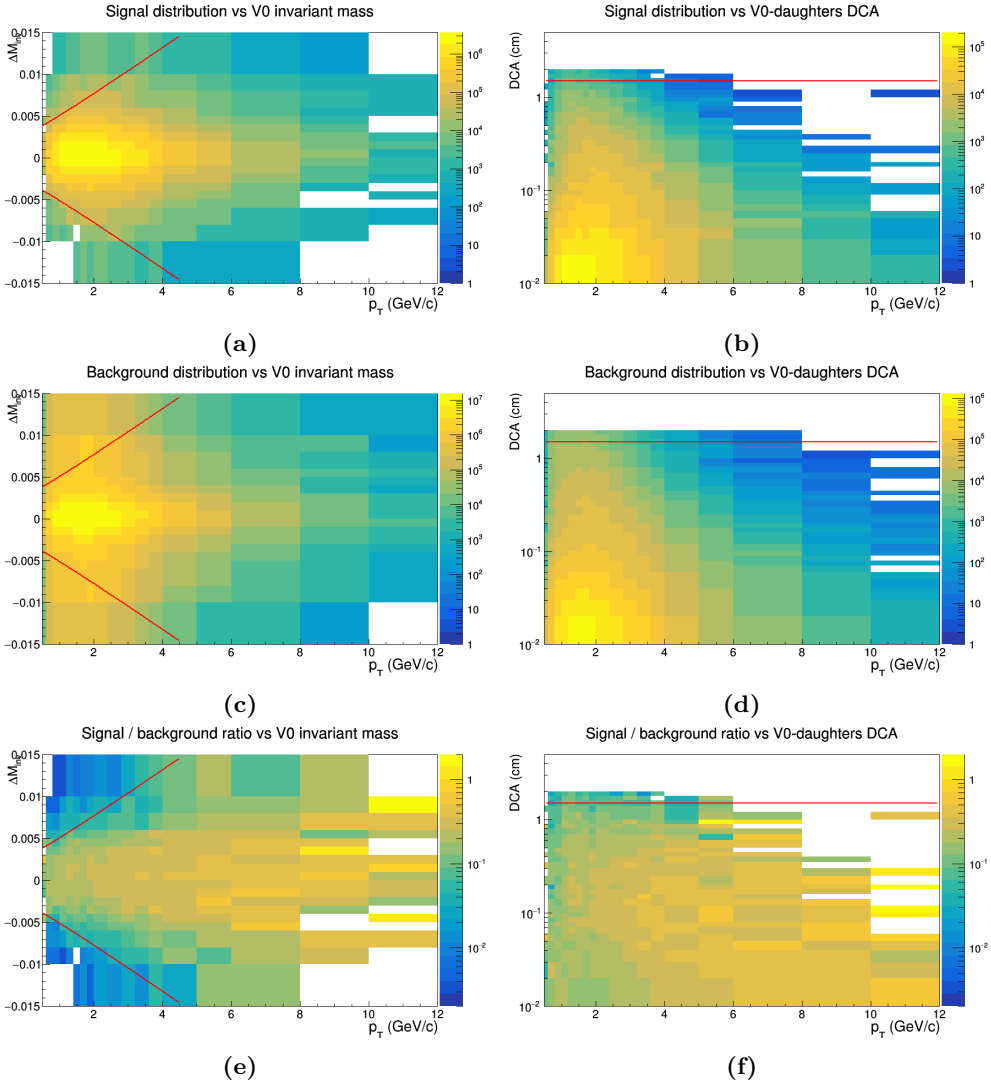

Figure D.9 – **(a)-(b)** Signal, **(c)-(d)** background, and **(e)-(f)** signal-to-background ratio as a function of p_{T} and **(left)** V^0 invariant mass and **(right)** V^0 $\mathrm{DCA_{d-d}}$ cut for Ξ, for an otherwise very loose cut set. The selected cuts (power-law fit) are shown in red.

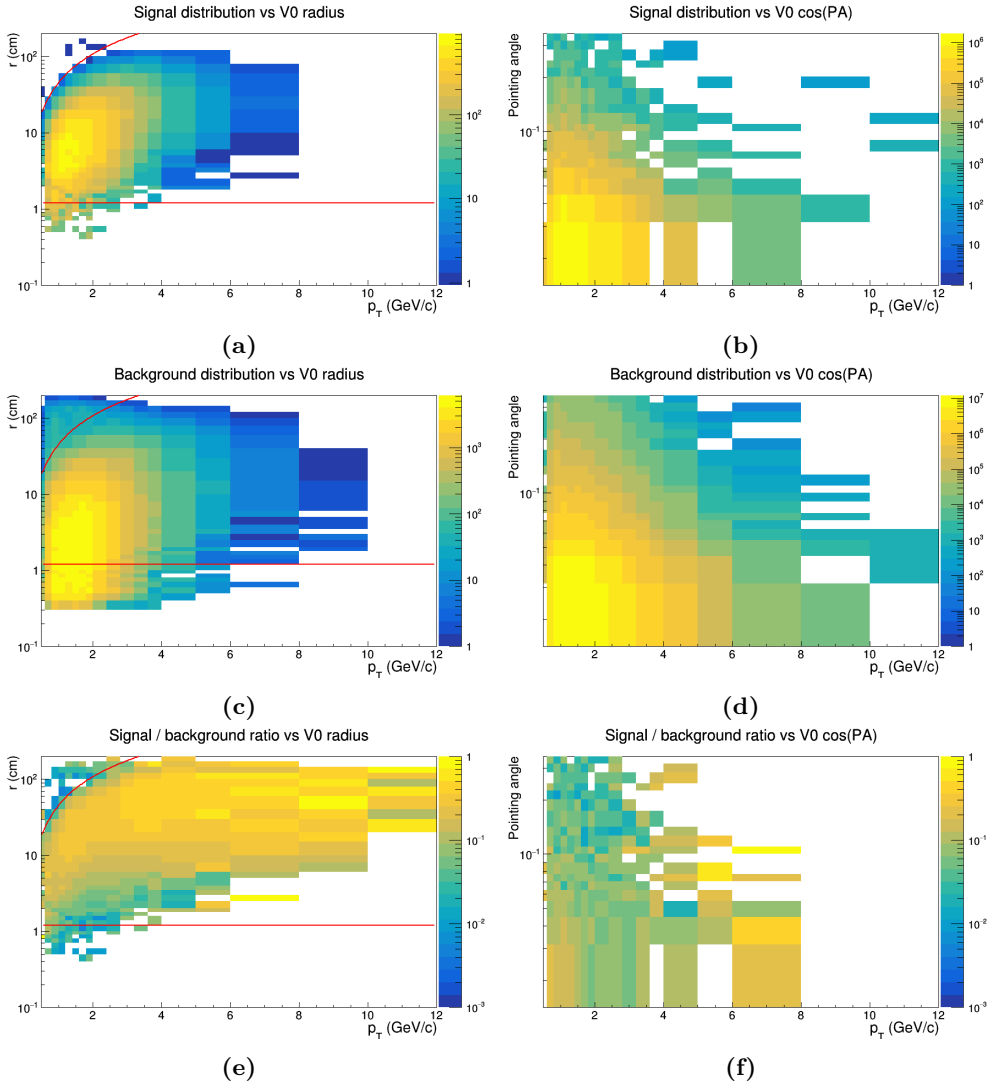

Figure D.10 – **(a)-(b)** Signal, **(c)-(d)** background, and **(e)-(f)** signal-to-background ratio as a function of p_{T} and **(left)** V^0 radius and **(right)** V^0 cos(PA) cut for Ξ, for an otherwise very loose cut set. The selected V^0 radius cuts (power-law fits) are shown in red; based on this no cut was applied in V^0 cos(PA).

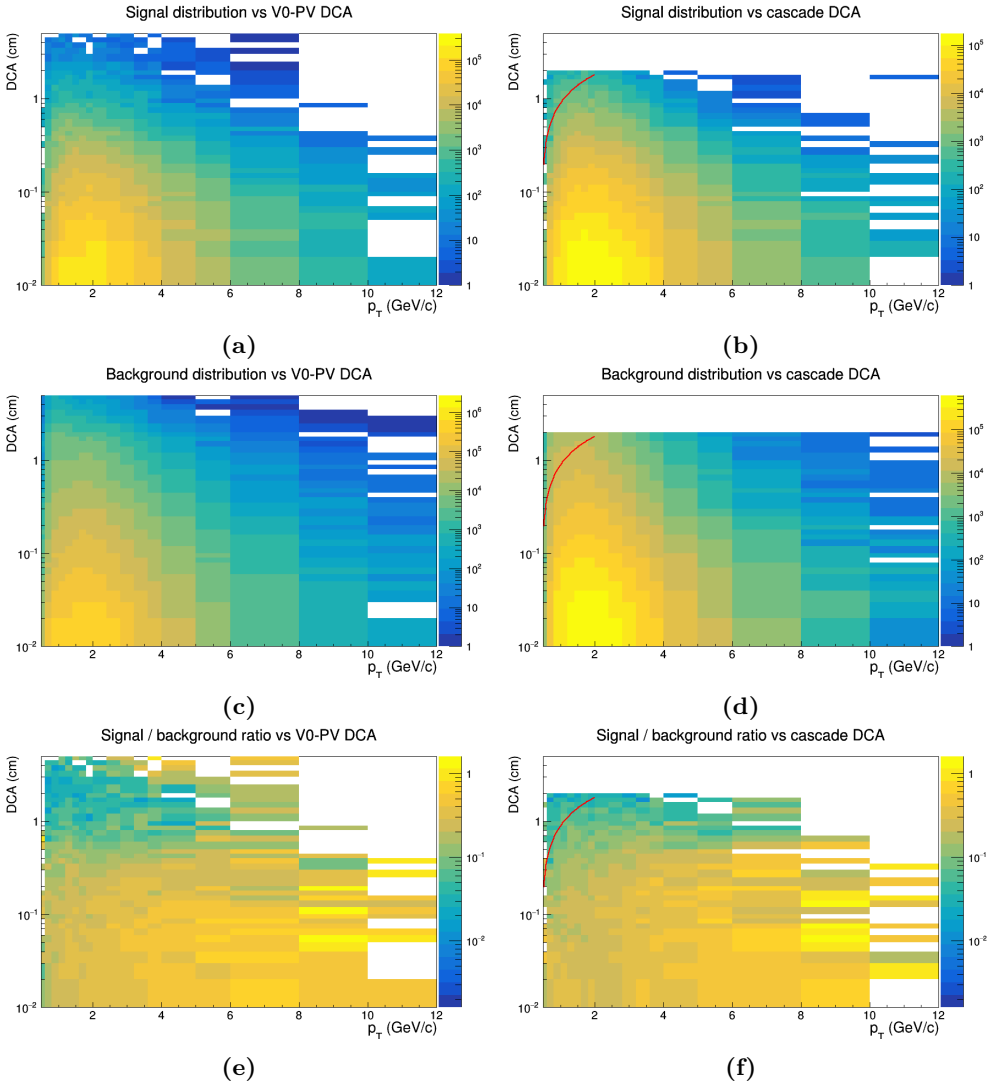

Figure D.11 – **(a)-(b)** Signal, **(c)-(d)** background, and **(e)-(f)** signal-to-background ratio as a function of p_{T} and **(left)** V^0-primary vertex DCA and **(right)** cascade-primary vertex cut for Ξ, for an otherwise very loose cut set. The selected $\mathrm{DCA}_{\mathrm{casc-PV}}$ cut (power-law fit) is shown in red; based on this no cut was applied in V^0-primary vertex DCA.

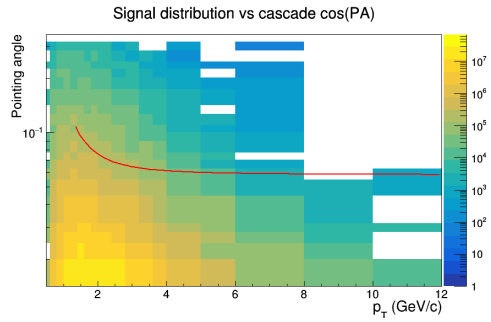

(a)

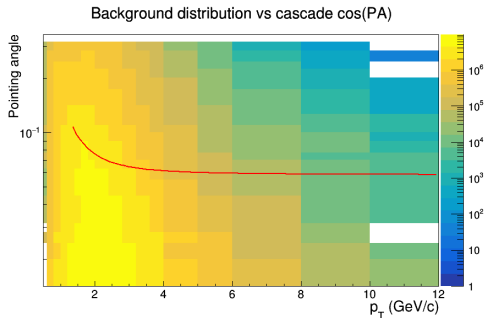

(b)

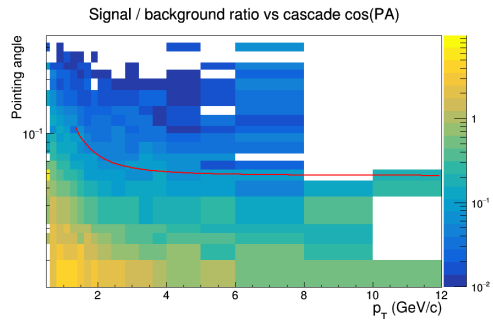

(c)

Figure D.12 – **(a)** Signal, **(b)** background, and **(c)** signal-to-background ratio as a function of p_{T} and cascade $\cos(PA)$ cut for Ξ, for an otherwise very loose cut set. Note the logarithmic scales. The selected cut (power-law fit) is shown in red.

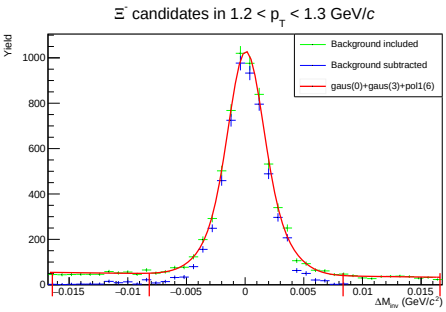

(a)

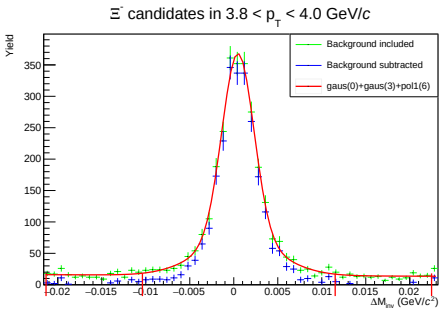

(b)

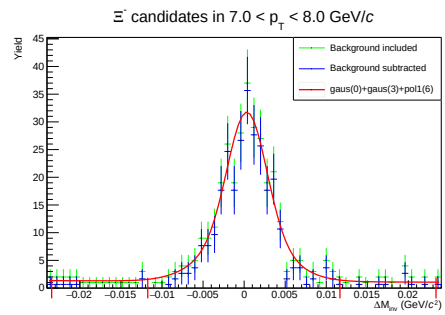

(c)

Figure D.13 – Invariant-mass distribution of Ξ^- candidates in the transverse-momentum regions **(a)** $1.2 < p_{\mathrm{T}} \leq 1.3\,\mathrm{GeV}/c$, **(b)** $3.8 < p_{\mathrm{T}} \leq 4.0\,\mathrm{GeV}/c$, and **(c)** $7.0 < p_{\mathrm{T}} \leq 8.0\,\mathrm{GeV}/c$, including a double-Gaussian fit + one-dimensional polynomial background. The red ticks mark the 3σ limits of the wider Gaussian component, which are used as limits for the signal region. The background is estimated from the sideband region, located between 3 and 6σ away from the mean.

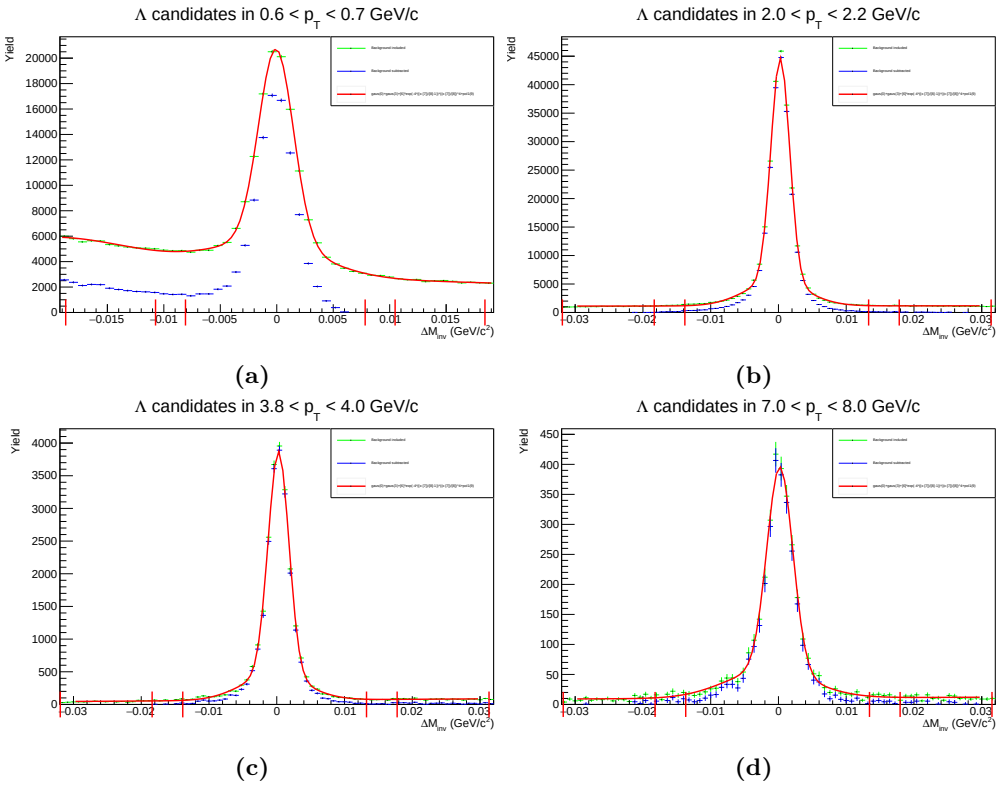

Figure D.14 – Invariant-mass distribution of Λ candidates in the transverse-momentum regions **(a)** $0.6 < p_{\mathrm{T}} \leq 0.7\,\mathrm{GeV}/c$, **(b)** $2.0 < p_{\mathrm{T}} \leq 2.2\,\mathrm{GeV}/c$, **(b)** $3.8 < p_{\mathrm{T}} \leq 4.0\,\mathrm{GeV}/c$, and **(c)** $7.0 < p_{\mathrm{T}} \leq 8.0\,\mathrm{GeV}/c$, including a double-Gaussian fit + one-dimensional polynomial background. In the lowest-p_{T} interval an additional component defined by Eq. (8.10) is included. The red ticks mark the 3σ limits of the wider Gaussian component, which are used as limits for the signal region. The background is estimated from the sideband region, located between 4 and 7σ away from the mean.

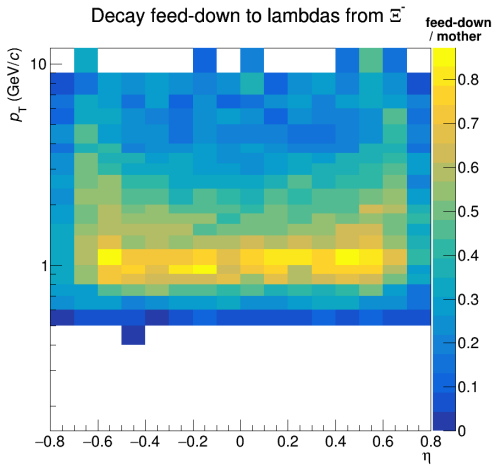

(a)

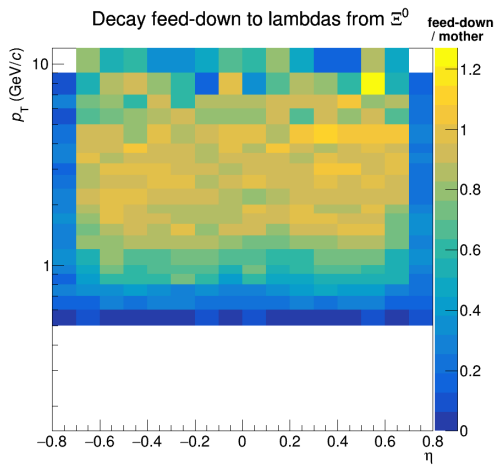

(b)

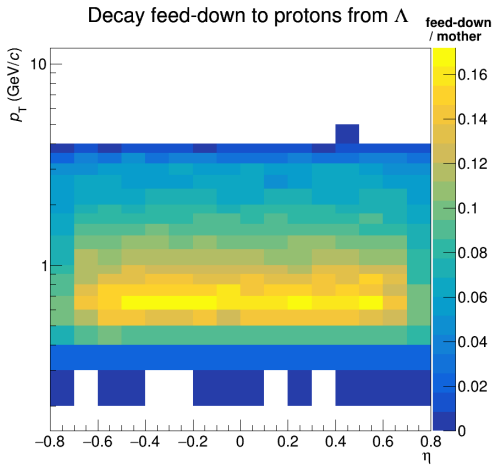

(c)

Figure D.15 – Efficiency-corrected probabilities of detecting **(a)** a Λ baryon in the range $0.6 < p_{\mathrm{T}} < 12\,\mathrm{GeV}/c$ and $|\eta| < 0.72$ originating from $\Xi^{\pm}$ decay, **(b)** a Λ baryon in the same kinematic region originating from Ξ^{0} decay, and **(c)** a proton in the range $0.4 < p_{\mathrm{T}} < 3\,\mathrm{GeV}/c$ and $|\eta| < 0.8$ originating from Λ decay, as a function of pseudorapidity and transverse momentum of the mother particle.

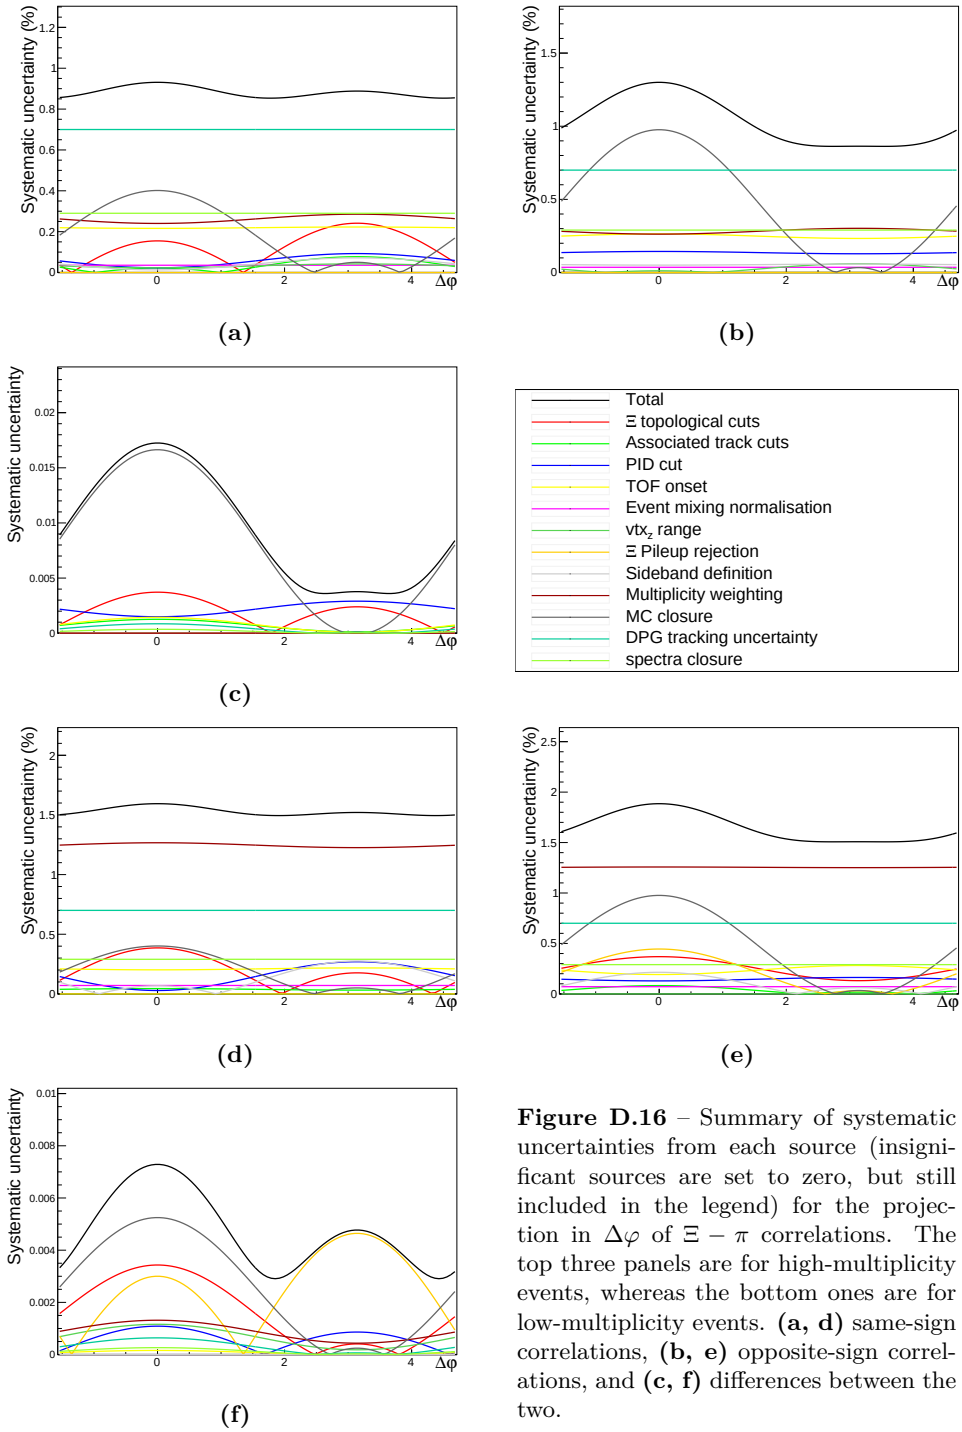

Figure D.16 – Summary of systematic uncertainties from each source (insignificant sources are set to zero, but still included in the legend) for the projection in $\Delta\varphi$ of $\Xi - \pi$ correlations. The top three panels are for high-multiplicity events, whereas the bottom ones are for low-multiplicity events. **(a, d)** same-sign correlations, **(b, e)** opposite-sign correlations, and **(c, f)** differences between the two.

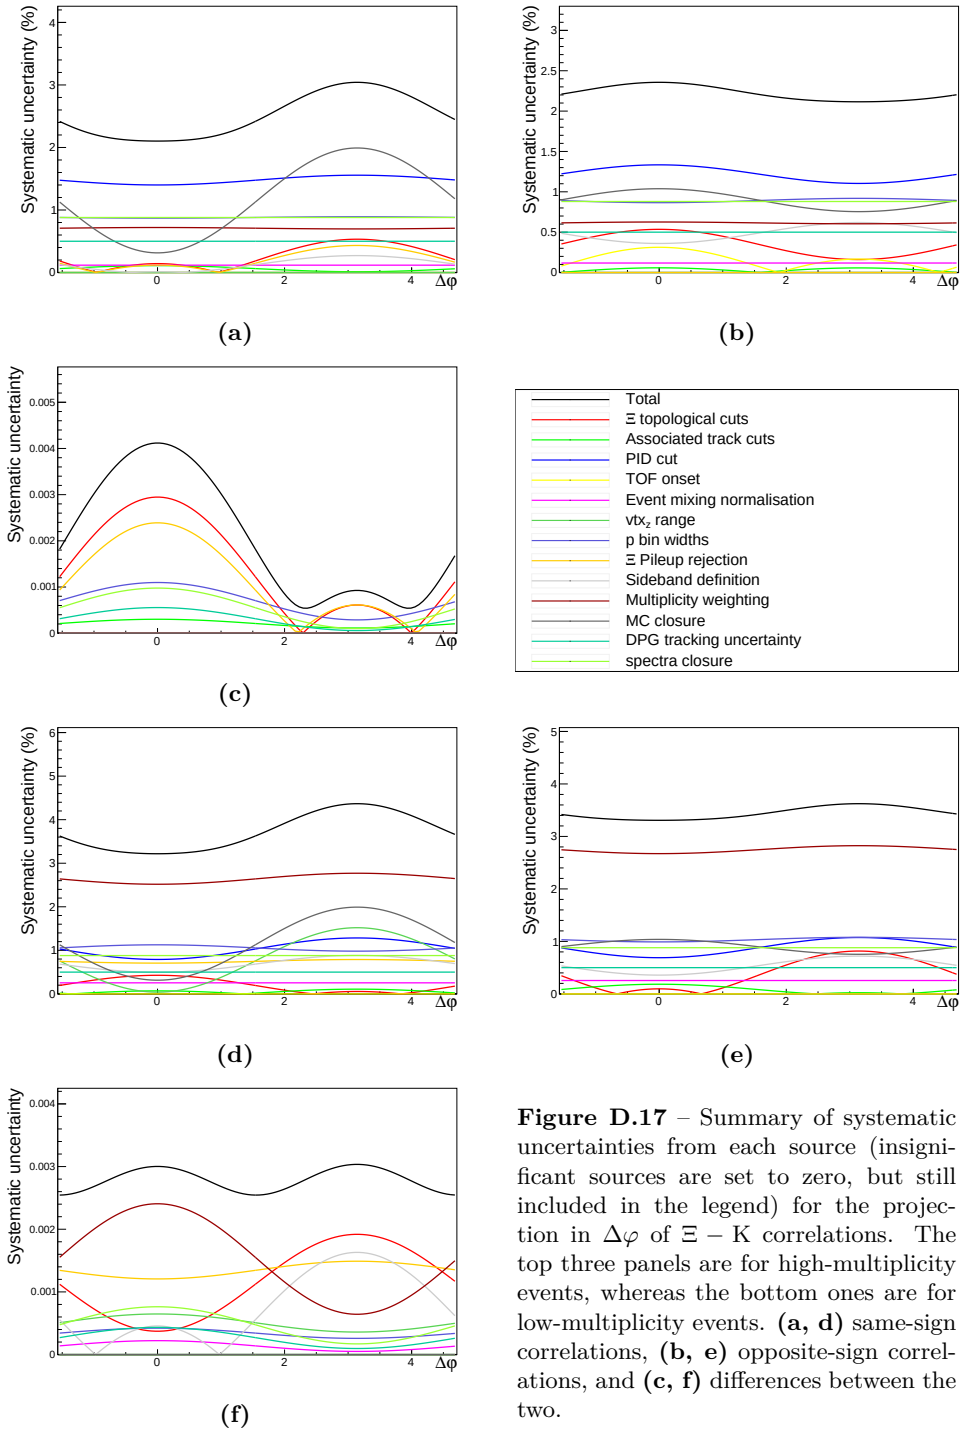

Figure D.17 – Summary of systematic uncertainties from each source (insignificant sources are set to zero, but still included in the legend) for the projection in $\Delta\varphi$ of $\Xi - \mathrm{K}$ correlations. The top three panels are for high-multiplicity events, whereas the bottom ones are for low-multiplicity events. **(a, d)** same-sign correlations, **(b, e)** opposite-sign correlations, and **(c, f)** differences between the two.

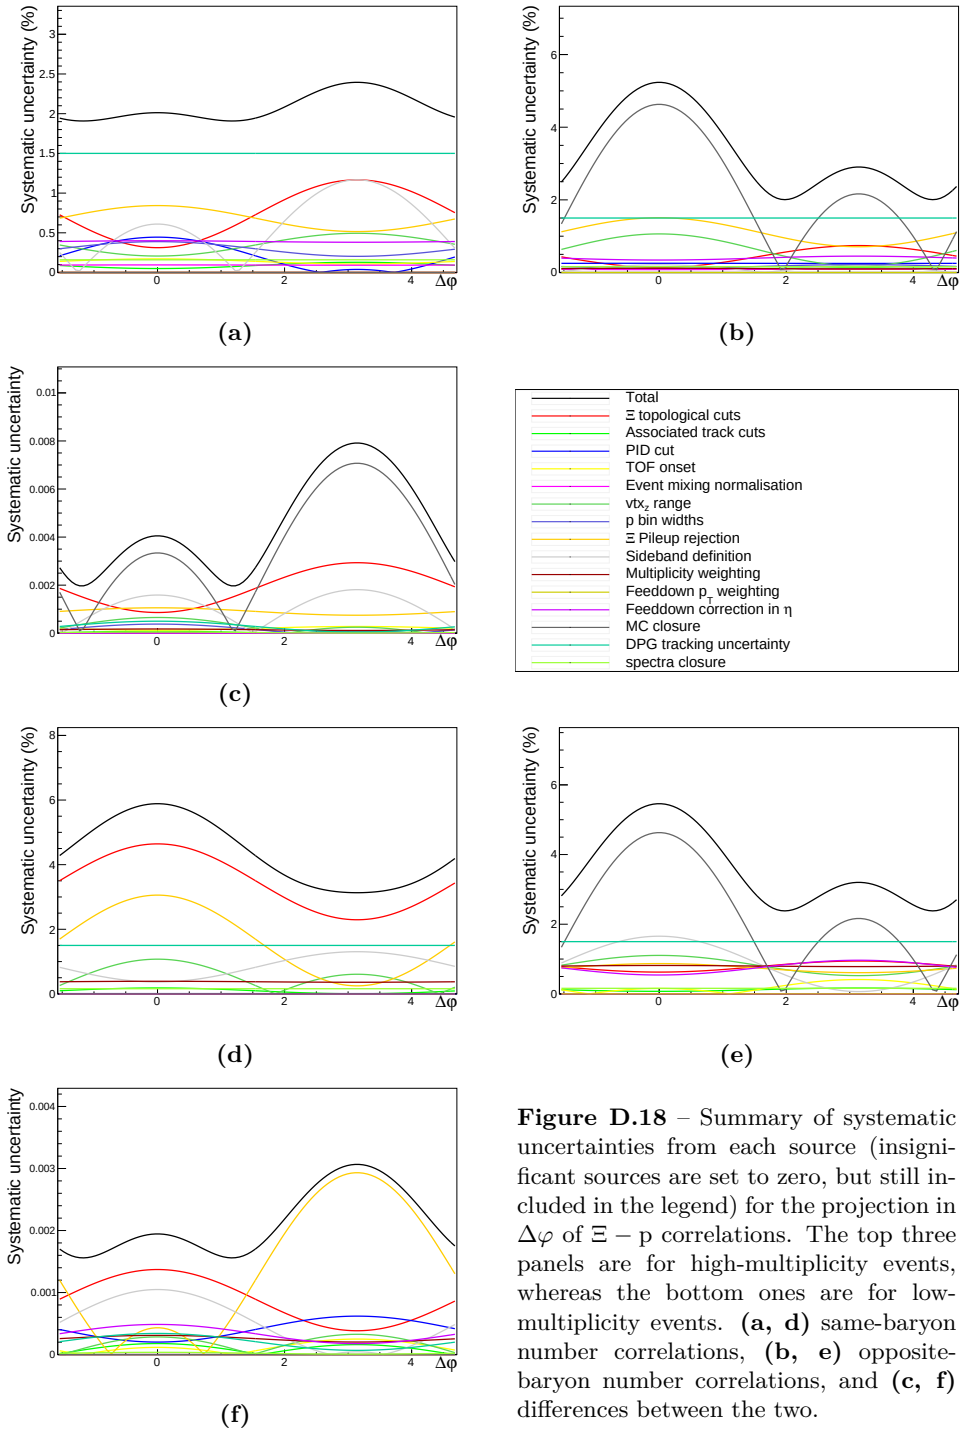

Figure D.18 – Summary of systematic uncertainties from each source (insignificant sources are set to zero, but still included in the legend) for the projection in $\Delta\varphi$ of $\Xi - \mathrm{p}$ correlations. The top three panels are for high-multiplicity events, whereas the bottom ones are for low-multiplicity events. **(a, d)** same-baryon number correlations, **(b, e)** opposite-baryon number correlations, and **(c, f)** differences between the two.

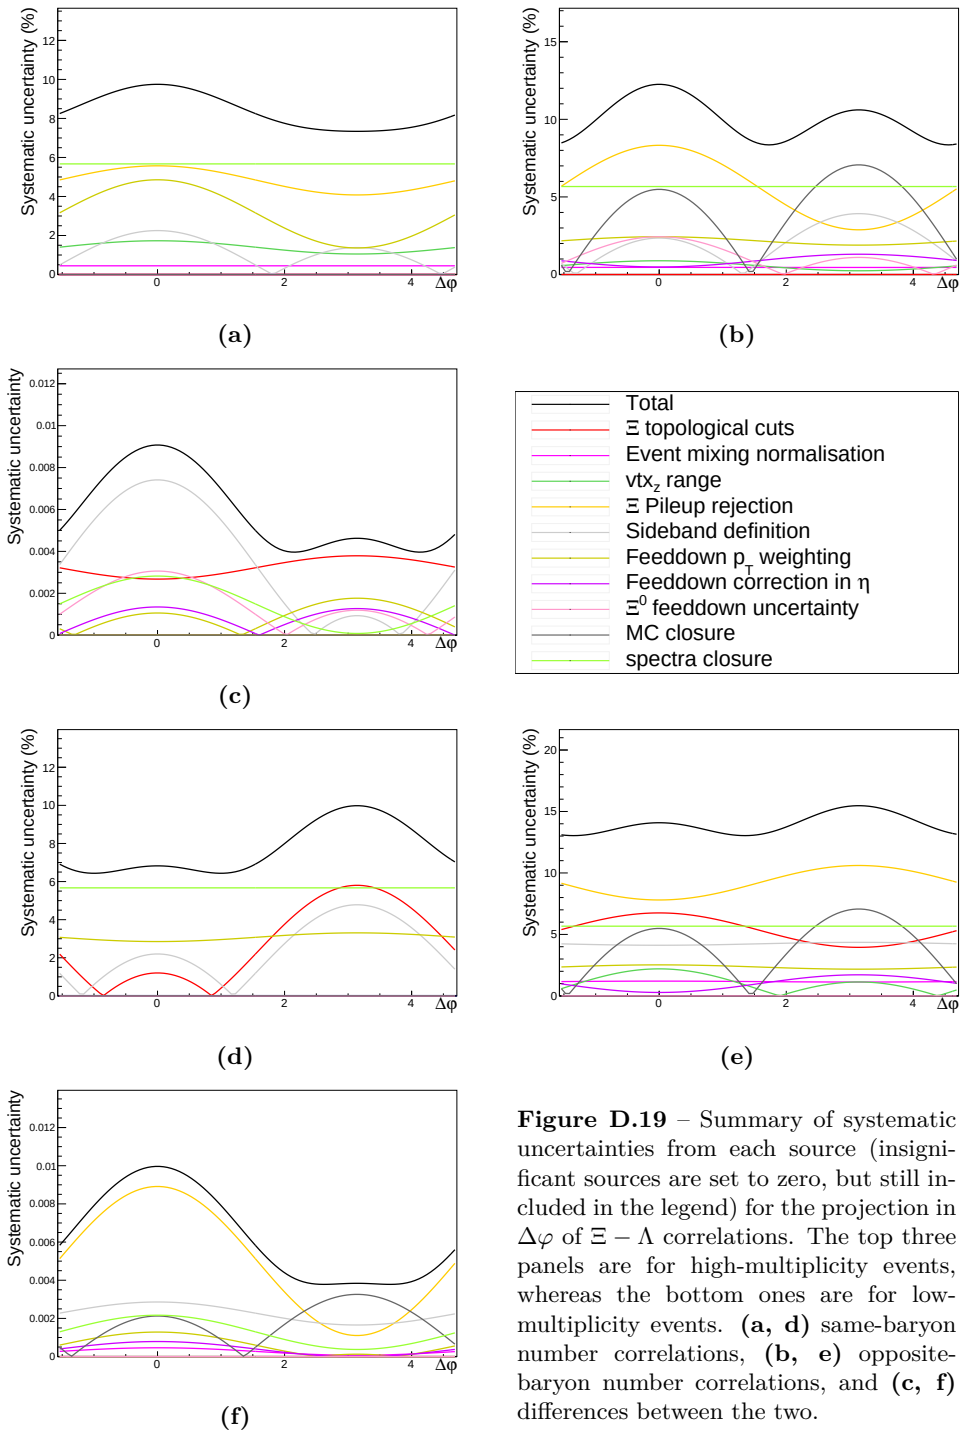

Figure D.19 – Summary of systematic uncertainties from each source (insignificant sources are set to zero, but still included in the legend) for the projection in $\Delta\varphi$ of $\Xi - \Lambda$ correlations. The top three panels are for high-multiplicity events, whereas the bottom ones are for low-multiplicity events. **(a, d)** same-baryon number correlations, **(b, e)** opposite-baryon number correlations, and **(c, f)** differences between the two.

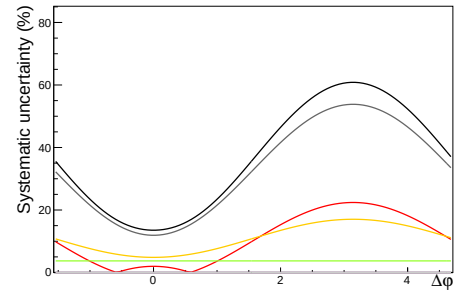

(a)

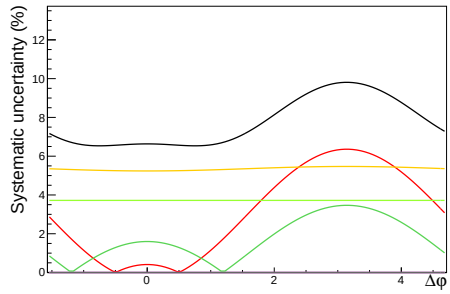

(b)

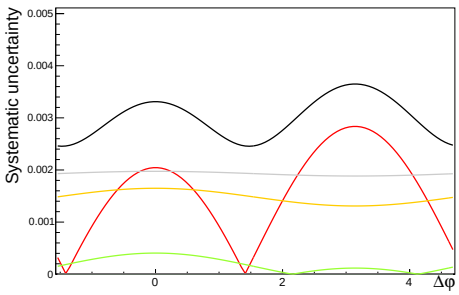

(c)

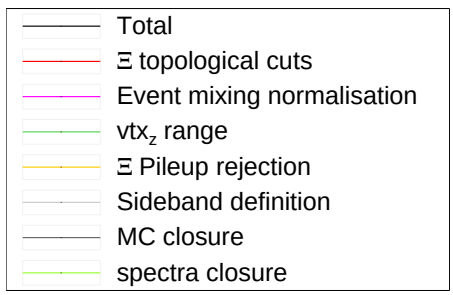

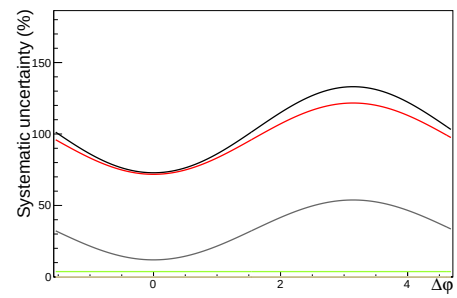

(d)

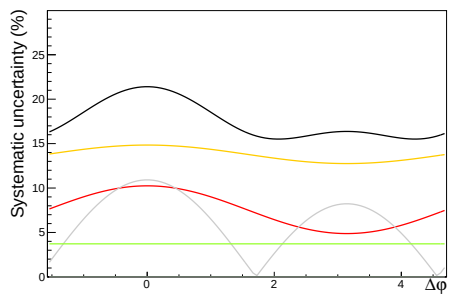

(e)

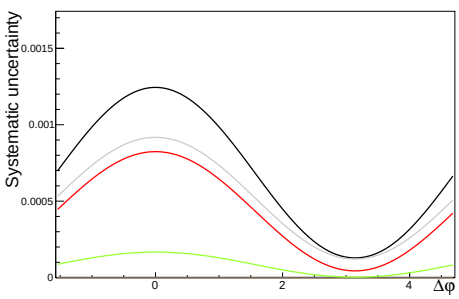

(f)

Figure D.20 – Summary of systematic uncertainties from each source (insignificant sources are set to zero, but still included in the legend) for the projection in $\Delta\varphi$ of $\Xi - \Xi$ correlations. The top three panels are for high-multiplicity events, whereas the bottom ones are for low-multiplicity events. **(a, d)** same-baryon number correlations, **(b, e)** opposite-baryon number correlations, and **(c, f)** differences between the two.

D.3 Complementary Ξ-Hadron Correlation Results

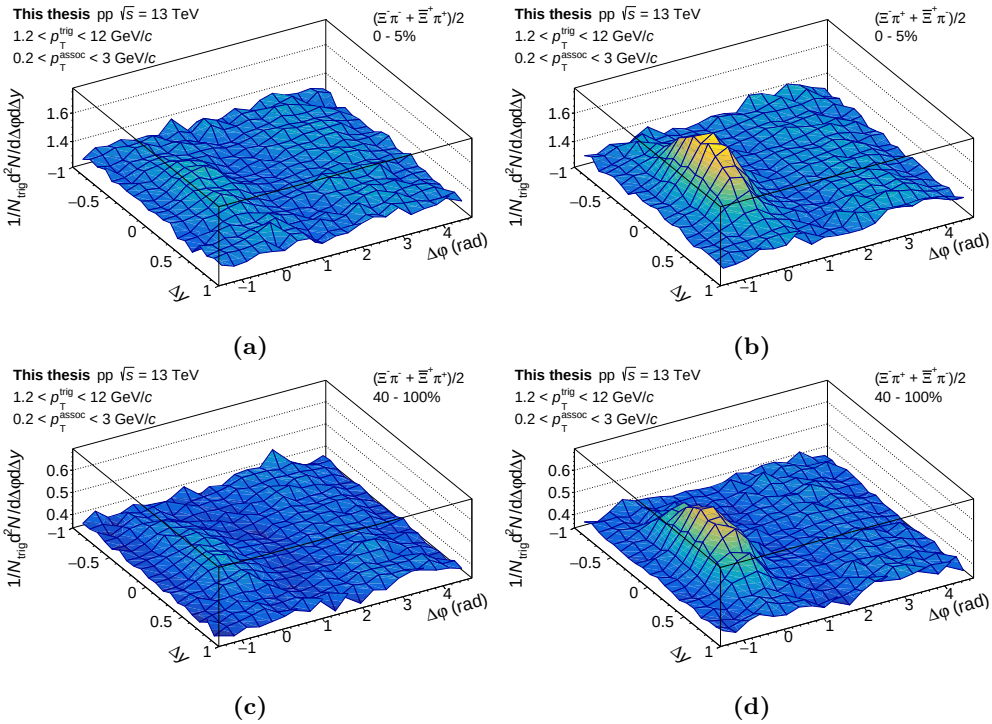

Figure D.21 – Correlations in $(\Delta y, \Delta\varphi)$ space between Ξ baryons and pions for **(a)** same-sign pairs at high multiplicity, **(b)** opposite-sign pairs at high multiplicity, **(c)** same-sign pairs at low multiplicity, and **(d)** opposite-sign pairs at low multiplicity

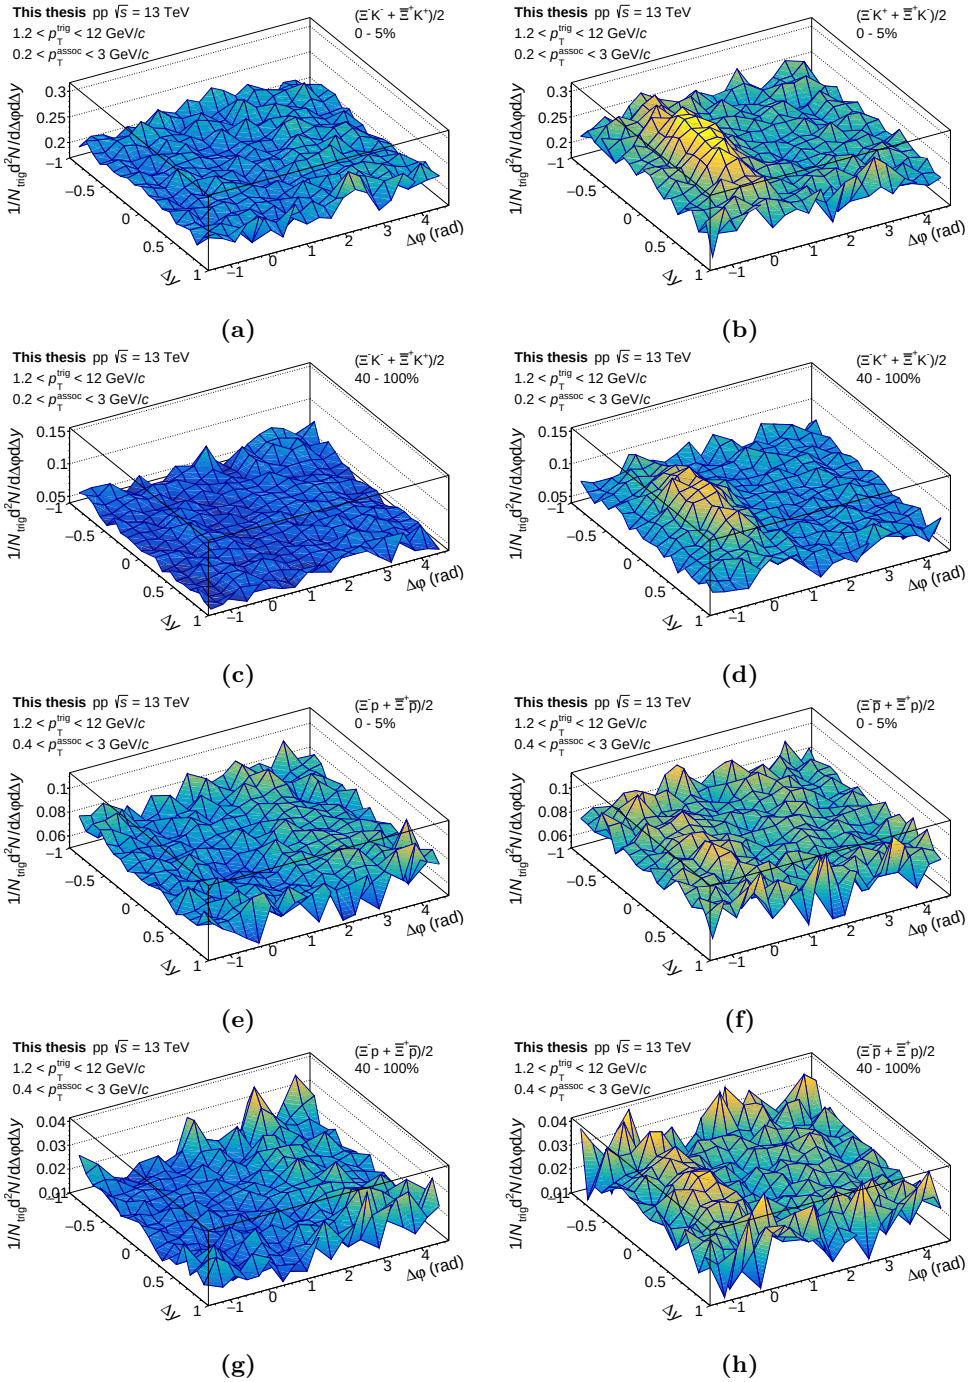

Figure D.22 – Correlations in $(\Delta y, \Delta\varphi)$ space between Ξ baryons and **(a, c)** same-sign kaons, **(b, d)** opposite-sign kaons, **(e, g)** same-baryon number protons, and **(f, h)** opposite-baryon number protons, with high-multiplicity events in the upper panels and low-multiplicity events in the lower ones.

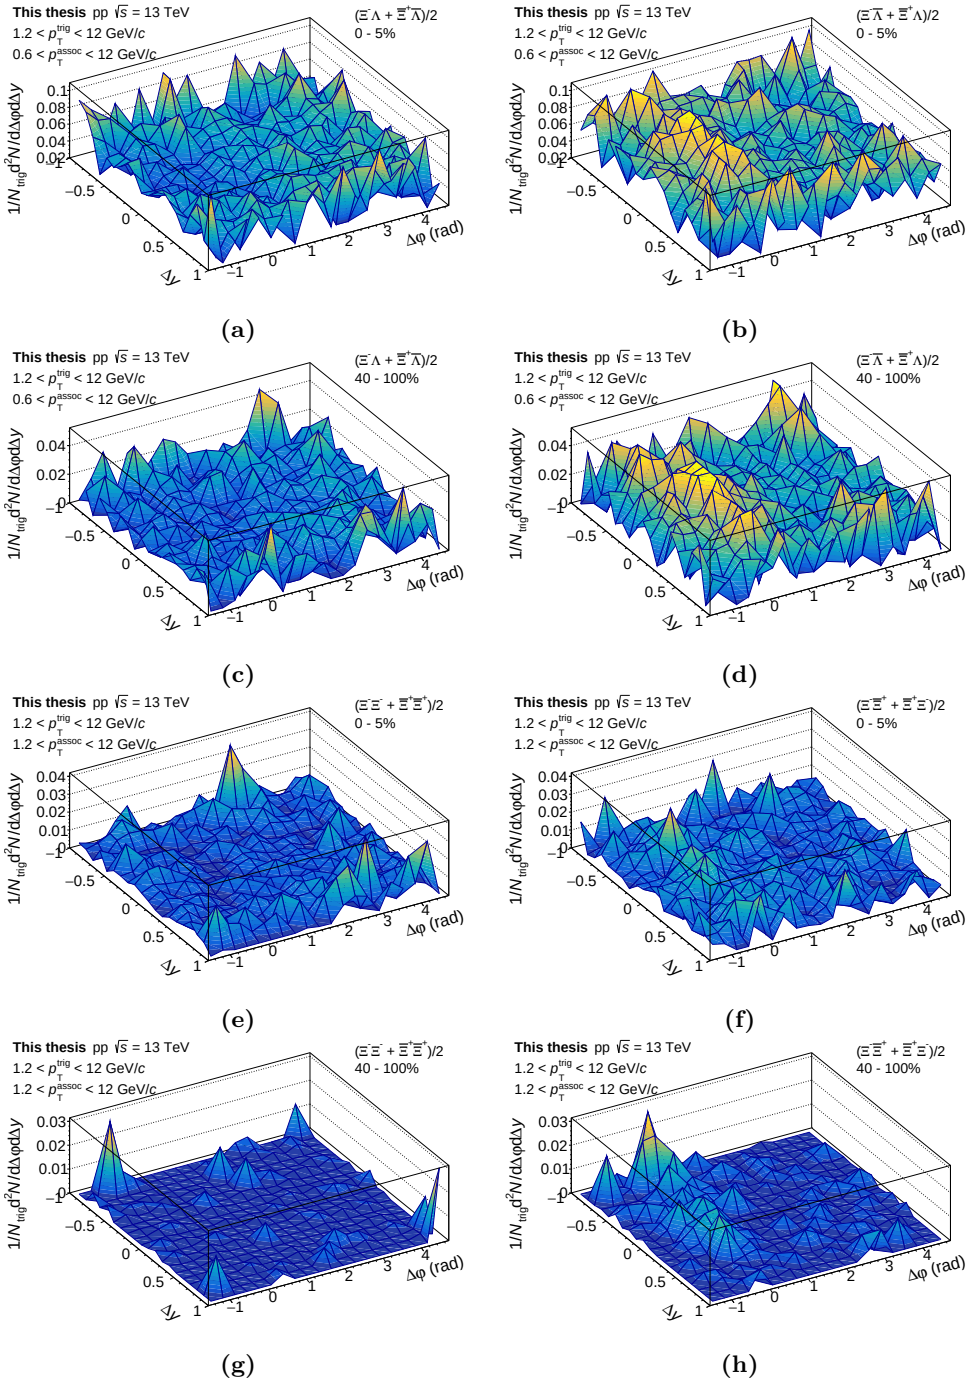

Figure D.23 – Correlations in $(\Delta y, \Delta\varphi)$ space between Ξ baryons and **(a, c)** same-B Λ baryons, **(b, d)** opposite-B Λ baryons, **(e, g)** same-B Ξ baryons, and **(f, h)** opposite-B Ξ baryons, with high-multiplicity events in the upper panels and low-multiplicity events in the lower ones.

Bibliography

[1] M. Brice. *Aerial View of the CERN taken in 2008.* CERN Document Server, CERN-MI-0807031, 2008.

[2] A. Andronic and S.G. Weber. *ALICE event display of a Pb-Pb collision at 2.76A TeV.* CERN Document Server, ALICE-PHO-GEN-2015-004, 2015.

[3] MissMJ, Cush. *Standard Model of Elementary Particles.* `https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg`. Accessed on 10 October 2020.

[4] LHCb Collaboration. Observation of $J/\psi p$ Resonances Consistent with Pentaquark States in $\Lambda_b^0 \to J/\psi K^- p$ Decays. *Phus. Rev. Lett.* **115**, 072001, 2015.

[5] S. Weinberg. A Model of Leptons. *Phys. Rev. Lett.* **19**, 1264-1266, 1967.

[6] P.W. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phus. Rev. Lett.* **13**, 508-509, 1964.

[7] ATLAS Collaboration. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett.* **B 716**, 1-29, 2012.

[8] CMS Collaboration. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett.* **B 716**, 30-61, 2012.

[9] D. Kaiser. Physics and Feynmans Diagrams. *American Scientist* **93**:2, 156-165, 2005.

[10] B.R. Martin and G. Shaw. *Particle Physics.* Third edition, John Wiley & Sons, 2008.

[11] Particle Data Group. Review of Particle Physics. *Chin. Phys.* **C** 40, 100001, 2016.

[12] V.E. Barnes *et al.* Observation of a Hyperon with Strangeness Minus Three. *Phys. Rev. Let.* **12**, 204-206, 1964.

[13] G. Kane. *Modern Elementary Particle Physics: The Fundamental Particles and Forces?.* Updated edition, Addison-Wesley, 1993.

[14] K.G. Wilson. Confinement of quarks. *Phys. Rev. D* **10**, 2445-2459, 1974.

[15] N. Cabibbo and G. Parisi. Exponential hadronic spectrum and quark liberation. *Phys. Lett.* **B 59**, 67-69, 1975.

[16] E.V. Shuryak. Quark-gluon plasma and hadronic production of leptons, photons and psions. *Phys. Lett.* **B 78**, 150-153, 1978.

[17] A. Bazavov *et al.* Equation of state and QCD transition at finite temperature. *Phys. Rev. D* **80**, 014504, 2009.

[18] V. Vovchenko, M.I. Gorenstein, and H. Stoecker. van der Waals Interactions in Hadron Resonance Gas: From Nuclear Matter to Lattice QCD. *Phys. Rev. Lett.* **118**, 182301, 2017.

[19] The DOE/NSF Nuclear Science Advisory Committee. The Frontiers of Nuclear Science, A Long Range Plan. *arXiv*: 0809.3137, 2008.

[20] M. Alford, K. Rajagopal, and F. Wilczek. QCD at finite baryon density: nucleon droplets and color superconductivity. *Phys. Lett.* **B 422**, 247-256, 1998.

[21] D. d'Enterria. Quark-gluon matter. *J. Phys.* **G 34**, S53-S82, 2007.

[22] J. Maldacena. The large N limit of superconformal field theories and supergravity. *Adv. Theor. Math. Phys.* **2**, 231-252, 1998.

[23] P.K. Kovtun, D.T. Son, and A.O. Starinets. Viscosity in Strongly Interacting Quantum Field Theories from Black Hole Physics. *Phys. Rev. Lett.* **94**, 111601, 2005.

[24] M. Müller, J. Schmalian, and L. Fritz. Graphene: A Nearly Perfect Fluid. *Phys. Rev. Lett.* **103**, 025301, 2009.

[25] ALICE Collaboration. Two-pion Bose–Einstein correlations in central Pb–Pb collisions at $\sqrt{s_{\rm NN}} = 2.76$ TeV. *Phys. Lett.* **B 696**, 328-337, 2011.

[26] P. Braun-Munzinger and B. Dönigus. Loosely-bound objects produced in nuclear collisions at the LHC. *Nucl. Phys.* **A 987**, 144-201, 2019.

[27] W. Florkowski. *Phenomenology of Ultra-Relativistic Heavy-Ion Collisions*. World Scientific, 2010.

[28] H1 and ZEUS Collaborations. Combined measurement and QCD analysis of the inclusive $e^{\pm}p$ scattering cross sections at HERA. *J. High Energ. Phys.* **2010**, 109, 2010.

[29] T. Lappi and L. McLerran. Some features of the glasma. *Nucl. Phys.* **A772**, 200-212, 2006.

[30] A. Kurkela and G.D. Moore. Thermalization in weakly coupled nonabelian plasmas. *J. High Energ. Phys.* **2011**, 44, 2011.

[31] C. Loizides, J. Kamin, and D. d'Enterria. Improved Monte Carlo Glauber predictions at present and future nuclear colliders. *Phys. Rev. C* **97**, 054910, 2018.

[32] A. Toia. Participants and spectators at the heavy-ion fireball. *CERN Courier* **53**:4, 31-33, 2013.

[33] ALICE Collaboration. Centrality determination of Pb–Pb collisions at $\sqrt{s_{\rm NN}} =$ 2.76 TeV with ALICE. *Phys. Rev.* **C 88**, 044909, 2013.

[34] M.L. Miller, K. Reygers, S.J. Sanders, and P. Steinberg. Glauber Modeling in High-Energy Nuclear Collisions. *Annu. Rev. Nucl. Part. Sci.* **57**, 205-243, 2007.

[35] C. Loizides. Glauber modeling of high-energy nuclear collisions at the subnucleon level. *Phys. Rev. C* **94**, 024914, 2016.

[36] ALICE Collaboration. Direct photon production in Pb–Pb collisions at $\sqrt{s_{\rm NN}} =$ 2.76 TeV. *Phys. Lett.* **B 754**, 235-248, 2016.

[37] Guinness World Records. *Highest artificial temperature.* `https://www.guinnessworldrecords.com/world-records/highest-man-made-temperature`. Accessed on 15 October 2020.

[38] ALICE Collaboration. Direct photon production at low transverse momentum in proton-proton collisions at $\sqrt{s} = 2.76$ and 8 TeV. *Phys. Rev. C* **99**, 024912, 2019.

[39] CMS Collaboration. *CMS collision events: from lead ion collisions.* CERN Document Server, CMS-PHO-EVENTS-2010-003, 2010.

[40] ALICE Collaboration. Centrality dependence of the nuclear modification factor of charged pions, kaons, and protons in Pb-Pb collisions at $\sqrt{s_{\rm NN}} = 2.76$ TeV. *Phys. Rev. C* **93**, 034913, 2016.

[41] S. Peigne and A.V. Smilga. Energy losses in relativistic plasmas: QCD versus QED. *Phys. Usp.* **52**, 659-685.

[42] K.M. Burke *et al.* Extracting the jet transport coefficient from jet quenching in high-energy heavy-ion collisions. *Phys. Rev. C* **90**, 014909, 2014.

[43] P.M. Jacobs. Search for jet quenching effects in high multiplicity pp collisions at $\sqrt{s} = 13$ TeV. *arXiv*: 2001.09517, 2020.

[44] CMS Collaboration. Measurement of nuclear modification factors of Υ(1S), Υ(2S), and Υ(3S) mesons in PbPb collisions at $\sqrt{s_{\rm NN}} = 5.02$ TeV. *Phys. Lett.* **B 790**, 270-293, 2019.

[45] T. Matsui and H. Satz. J/ψ Suppression by Quark-Gluon Plasma Formation. *Phys. Lett.* **B178**, 416-422, 1986.

[46] ALICE Collaboration. J/ψ Suppression at Forward Rapidity in Pb-Pb Collisions at $\sqrt{s_{\rm NN}} = 2.76$ TeV. *Phys. Rev. Lett.* **109**, 072301, 2012.

[47] ALICE Collaboration. Centrality, rapidity and transverse momentum dependence of J/ψ suppression in PbPb collisions at $\sqrt{s_{\rm NN}} = 2.76$ TeV. *Phys. Lett.* **B 734**, 314-327, 2014.

[48] ALICE Collaboration. Centrality dependence of ψ(2S) suppression in p-Pb collisions at $\sqrt{s_{\rm NN}} = 5.02\,\text{TeV}$. *J. High Energ. Phys.* **2016**, 50, 2016.

[49] J. Rafelski and B. Müller. Strangeness Production in the Quark-Gluon Plasma. *Phys. Rev. Lett.* **48**, 1066-1069, 1982.

[50] A. Andronic, P. Braun-Munzinger, and J. Stachel. Thermal hadron production in relativistic nuclear collisions. *Acta Phys. Polon. B* **40**, 1005-1012, 2009.

[51] ALICE Collaboration. Enhanced production of multi-strange hadrons in high-multiplicity proton-proton collisions. *Nature Phys.* **13**, 535-539, 2017.

[52] A. Khuntia. Latest results on the production of hadronic resonances inALICE at the LHC. *arXiv*: 2001.09023, 2020.

[53] S.A. Voloshin. Two particle rapidity, transverse momentum, and azimuthal correlationsin relativistic nuclear collisions and transverse radial expansion. *Nucl. Phys.* **A 749**, 287c-290c, 2005.

[54] ALICE Collaboration. Multiplicity and transverse momentum evolution of charge-dependent correlations in pp, p–Pb, and Pb–Pb collisions at the LHC. *Eur. Phys. J. C* **76**, 86, 2016.

[55] B. Alver, M. Baker, C. Loizides, and P. Steinberg. The PHOBOS Glauber Monte Carlo. *arXiv*: 0805.4411, 2008.

[56] H. Holopainen, H. Niemi, and K.J. Eskola. Event-by-event hydrodynamics and elliptic flow from fluctuating initial states. *Phys. Rev. C* **83**, 034901, 2011.

[57] A. Bilandzic, R. Snellings, and S. Voloshin. Flow analysis with cumulants: Direct calculations. *Phys. Rev. C* **83**, 044913, 2011.

[58] STAR Collaboration. Elliptic flow from two- and four-particle correlations in Au+Au collisions at $\sqrt{s_{\rm NN}} = 130\,\text{GeV}$. *Phys. Rev. C* **66**, 034904, 2002.

[59] N. Borghini, P.M. Dinh, and J.-Y. Ollitrault. Flow analysis from multiparticle azimuthal correlations. *Phys. Rev. C* **64**, 054901, 2001.

[60] A. Bilandzic *et al.* Generic framework for anisotropic flow analyses with multiparticle azimuthal correlations. *Phys. Rev. C* **89**, 064904, 2014.

[61] S.A Voloshin, A.M. Poskanzer, and R. Snellings. Collective phenomena in non-central nuclear collisions. *arXiv*: 0809.2949, 2008.

[62] ALICE Collaboration. Anisotropic Flow of Charged Particles in Pb-Pb Collisions at $\sqrt{s_{\rm NN}} = 5.02$ TeV. *Phy. Rev. Lett.* **116**, 132302, 2016.

[63] H. Niemi, K.J. Eskola, R. Paatelainen, and K. Tuominen. Predictions for 5.023 TeV Pb+Pb collisions at the LHC. *Phys. Rev. C* **93**, 014912, 2016.

[64] J. Noronha-Hostler, M. Luzum, and J.-Y. Ollitrault. Hydrodynamic predictions for 5.02 TeV Pb-Pb collisions. *Phys. Rev. C* **93**, 034912, 2016.

[65] ALICE Collaboration. Anisotropic flow of identified particles in Pb-Pb collisions at $\sqrt{s_{\rm NN}} = 5.02$ TeV. *J. High Energ. Phys.* **2018**, 6, 2018.

[66] R.J. Fries, B. Müller, C. Nonaka, and S.A. Bass. Hadronization in Heavy-Ion Collisions: Recombination and Fragmentation of Partons. *Phys. Rev. Lett.* **90**, 202303, 2003.

[67] R.D. Weller and P. Romatschke. One fluid to rule them all: Viscous hydrodynamic description of event-by-event central p+p, p+Pb and Pb+Pb collisions at *Phys. Lett.* **B 774**, 351-356, 2017.

[68] Y. Zhou, W. Zhao, K. Murase, and H. Song. One fluid might not rule them all. *arXiv*: 2005.02684, 2020.

[69] T. Pierog *et al.* EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider. *Phys. Rev. C* **92**, 034906, 2015.

[70] K. Werner, B. Guiot, Iu. Karpenko, and T. Pierog. Analyzing radial flow features in $p-$Pb and $p-p$ collisions at several TeV by studying identified-particle production with the event generator EPOS3. *Phys. Rev. C* **89**, 064903, 2014.

[71] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand. Parton Fragmentation and String Dynamics. *Phys. Rept.* **97**, 31-145, 1983.

[72] M. Bähr *et al.* Herwig++ physics and manual. *Eur. Phys. J. C* **58**, 639-707, 2008.

[73] R. Engel, J. Ranft, and S. Roesler. Hard diffraction in hadron-hadron interactions and in photoproduction. *Phys. Rev. D* **52**, 1459-1468, 1995.

[74] T. Sjöstrand *et al.* An introduction to PYTHIA 8.2. *Comput. Phys. Commun.* **191**, 159-177, 2015.

[75] C. Bierlich, G. Gustafson, L. Lönnblad, and H. Shah. The Angantyr model for heavy-ion collisions in PYTHIA8. *J. High Energ. Phys.* **2018**, 134, 2018.

[76] ALICE Collaboration. Multiplicity dependence of inclusive J/psi production at midrapidity in pp collisions at $\sqrt{s} = 13$ TeV. *Phys. Lett.* **B 810**, 135758, 2020.

[77] S.G. Weber. Measurement of $J\psi$ production as a function of event multiplicity in pp collisions at $\sqrt{s} = 13$ TeV with ALICE. *Nucl. Phys.* **A 967**, 333-336, 2017.

[78] C. Bierlich, G. Gustafson, L. Lönnblad, and L. Tarasov. Effects of overlapping strings in *pp* collisions. *J. High Energ. Phys.* **2015**, 148, 2015.

[79] C. Flensburg, G. Gustafson, and L. Lönnblad. Inclusive and exclusive observables from dipoles in high energy collisions. *J. High Energ. Phys.* **2011**, 103, 2011.

[80] Z. Khabanova. *Collective effects in pp collisions with the balance function of the identified particles.* Poster presented at the *XXVIIIth International Conference on Ultrarelativistic Nucleus-Nucleus Collisions* (Quark Matter 2019).

[81] C. Bierlich, S. Chakraborty, G. Gustafson, and L. Lönnblad. Setting the string shoving picture in a new frame. *arXiv*: 2010.07595, 2020.

[82] V. Vislavičius. *Identified Hadron Production as a Function of Event Multiplicity and Transverse Spherocity in pp Collisions at $\sqrt{s} = 7$ and 13 TeV with the ALICE Detector*. Lund University, Faculty of Science, Department of Physics, 2018.

[83] T. Sjöstrand. *Particle Physics Phenomenology 7. Hadronization*. Material used in the course "Particle Physics Phenomenology", Lund University, 2018.

[84] P. Skands, S. Carazza, and J. Rojo. Tuning PYTHIA 8.1: the Monash 2013 tune. *Eur. Phys. J. C* **74**, 3024, 2014.

[85] T. Sjöstrand. *Particle Physics Phenomenology 6. Multiparton interactions and MB/UE*. Material used in the course "Particle Physics Phenomenology", Lund University, 2018.

[86] T. Sjöstrand. Colour reconnection and its effects on precise measurements at the LHC. *arXiv*: 1310.8073, 2013.

[87] J.R. Christiansen and P.Z. Skands. String Formation Beyond Leading Colour. *arXiv*: 1505.01681, 2015.

[88] J. Adolfsson *et al.* QCD Challenges from pp to A-A collisions. *arXiv*: 2003.10997, 2020.

[89] C. Bierlich, G. Gustafson, and L. Lönnblad. A shoving model for collectivity in hadronic collisions. *arXiv*: 1612.05132, 2016.

[90] K. Werner. Core-Corona Separation in Ultrarelativistic Heavy Ion Collisions. *Phys. Rev. Lett.* **98**, 152301, 2007.

[91] ALICE Collaboration. Centrality Dependence of the Charged-Particle Multiplicity Density at Midrapidity in Pb-Pb Collisions at $\sqrt{s_{\mathrm{NN}}} = 5.02\,\mathrm{TeV}$. *Phys. Rev. Lett.* **116**, 222302, 2016.

[92] J. Wenninger. *Operation and Configuration of the LHC in Run 2*. CERN Document Server, CERN-ACC-NOTE-2019-0007, 2019.

[93] A. Tauro. *ALICE Schematics*. CERN Document Server, ALICE-PHO-SKE-2017-001, 2017.

[94] ALICE Collaboration. The ALICE experiment at the CERN LHC. *JINST* **3** S08002, 2008.

[95] W.R. Leo. *Techniques for Nuclear and Particle Physics Experiments - A How-to Approach*. Second edition, Springer-Verlag, 1994.

[96] R. Brenner. *Laboratory Instruction on Ionizing Detectors "Silicon Sensor and Read Out Electronics"*. Material used in the course "Research training course in Detector Technology for particle physics", Helsinki, 2018.

[97] P. Gasik. Upgrade of the ALICE central barrel tracking detectors: ITS and TPC. *Nucl. Phys.* **A 982**, 943-946, 2019.

[98] ALICE Collaboration. Technical Design Report for the Upgrade of the ALICE Inner Tracking System. *J. Phys. G: Nucl. Part. Phys.* **41**, 087002, 2014.

[99] Dougsim. *Proportional counter avalanches.* `https://commons.wikimedia.org/wiki/File:Proportional_counter_avalanches.jpg`. Accessed on 1 June 2020.

[100] J. Alme *et al.* The ALICE TPC, a large 3-dimensional tracking device with fast readout for ultra-high multiplicity events. *NIM* **A 622**, 316-367, 2010.

[101] ALICE Collaboration. *Technical Design Report for the Upgrade of the ALICE Time Projection Chamber.* CERN Document Server, CERN-LHCC-2013-020 / ALICE-TDR-016, 2013.

[102] F.V. Böhmer *et al.* Simulation of space-charge effects in an ungated GEM-based TPC. *NIM* **A 719**, 101-108, 2013.

[103] ALICE Collaboration. *The ALICE Time of Flight Detector.* `http://aliceinfo.cern.ch/Public/en/Chapter2/Chap2_TOF.html`. Accessed on 9 June 2020.

[104] A. Maire. *Track reconstruction principle in ALICE for LHC run I and run II.* CERN Document Server, ALICE-PHO-SKE-2011-001, 2011.

[105] Martino Gagliardi (ALICE Trigger Coordination). Personal communication.

[106] ALICE Collaboration. *Future high-energy pp programme with ALICE.* CERN Document Server, CERN-LHCC-2020-018 / ALICE-PUBLIC-2020-005, 2020.

[107] R.E. Bosch, A.J. de Parga, B. Mota, and L. Musa. The ALTRO chip: a 16-channel A/D converter and digital processor for gas detectors. *IEEE Trans. Nucl. Sci.* **50**, 6, 2460-2469, 2003.

[108] S.H.I. Barboza *et al.* SAMPA chip: a new ASIC for the ALICE TPC and MCH upgrades. *JINST* **11** C02088, 2016.

[109] Oak Ridge National Laboratory, U.S. Dept. of Energy. *Major upgrades of particle detectors and electronics prepare CERN experiment to stream a data tsunami.* `https://www.ornl.gov/news/major-upgrades-particle-detectors-and-electronics-prepare-cern-experiment-stream-data-tsunami`. Accessed on 15 June 2020.

[110] ALICE Collaboration. *Upgrade of the ALICE Readout & Trigger System.* CERN Document Server, CERN-LHCC-2013-019 / ALICE-TDR-015, 2013.

[111] Terasic. *SoCKit - the Development Kit for New SoC Device.* `https://www.terasic.com.tw/cgi-bin/page/archive.pl?Language=English&CategoryNo=205&No=816&PartNo=1`. Accessed on 17 June 2020.

[112] J. Adolfsson *et al.* SAMPA Chip: the New 32 Channels ASIC for the ALICE TPC and MCH Upgrades. *JINST* **12** C04008, 2017.

[113] P. Carlson & S. Johansson. *Modern elektronisk mätteknik*. Liber, 1997.

[114] L.-C. Böiers. *Mathematical Methods of Optimization*. Studentlitteratur, 2010.

[115] R.P. Brent. *Algorithms for Minimization Without Derivatives*. Prentice-Hall, 1973.

[116] J. Adolfsson. Measurements of Anisotropic Flow in Xe–Xe Collisions at $\sqrt{s_{\mathrm{NN}}} = 5.44\,\mathrm{TeV}$ Using the ALICE Detector. *MDPI Proc.* **10**, 41, 2019.

[117] ALICE Collaboration. Anisotropic flow in Xe–Xe collisions at $\sqrt{s_{\mathrm{NN}}} = 5.44\,\mathrm{TeV}$. *Phys. Lett.* **B 784**, 82-95, 2018.

[118] M. Gyulassy and X.-N. Wang. HIJING 1.0: a Monte Carlo program for parton and particle production in high-energy hadronic and nuclear collisions. *Comput. Phys. Commun.* **83**, 307-331, 1994.

[119] G. Giacalone, J. Noronha-Hostler, M. Luzum, and J.-Y. Ollitrault. Hydrodynamic predictions for 5.44 TeV Xe+Xe collisions. *Phys. Rev. C* **97**, 034904, 2018.

[120] J.S. Moreland, J.E. Bernhard, and S.A. Bass. Alternative ansatz to wounded nucleon and binary collision scaling in high-energy nuclear collisions. *Phys. Rev. C* **92**, 011901, 2015.

[121] K. Hagino, N.W. Lwin, and M. Yamagami. Deformation parameter for diffuse density. *Phys. Rev. C* **74**, 017310, 2006.

[122] P. Möller, A.J. Sierk, T. Ichikawa, and H. Sagawa. Nuclear ground-state masses and deformations: FRDM(2012). *At. Data Nucl. Data Tables* **109-110**, 1-204, 2016.

[123] K.J. Eskola, H. Niemi, R. Paatelainen, and K. Tuominen. Predictions for multiplicities and flow harmonics in 5.44 TeV Xe+Xe collisionsat the CERN Large Hadron Collider. *Phys. Rev. C* **97**, 034911, 2018.

[124] ALICE Collaboration. Centrality and pseudorapidity dependence of the charged-particle multiplicity density in Xe–Xe collisions at $\sqrt{s_{\mathrm{NN}}} = 5.44\,\mathrm{TeV}$. *Phys. Lett.* **B 790**, 35-48, 2019.

[125] H.-J. Drescher, Y. Nara. Eccentricity fluctuations from the color glass condensate in ultrarelativistic heavy ion collisions. *Phys. Rev. C* **76**, 041903, 2007.

[126] ATLAS Collaboration. Measurement of long-range multiparticle azimuthal correlations with the subevent cumulant method in *pp* and *p*+Pb collisions with the ATLAS detector at the CERN Large Hadron Collider. *Phys. Rev. C* **97**, 024904, 2018.

[127] CERN Application Software Group. *GEANT Detector Description and Simulation Tool*. CERN Program Library Long Writeup, 1993.

[128] ALICE Collaboration. Production of light-flavor hadrons in pp collisions at $\sqrt{s} = 7\,\mathrm{TeV}$ and $\sqrt{s} = 13\,\mathrm{TeV}$. *arXiv*: 2005.11120, 2020.

[129] J. Thaeder. *General ALICE Cross-Section without L3 Magnet.* ALICE General Figures. `https://alice-figure.web.cern.ch/system/files/figures/General/jthaeder/2012-Aug-02-General_ALICE_Cross-Section_without_Magnet.jpg`. Accessed on 8 September 2020.

[130] T. Richert. Λ/K^0_s *Associated with a Jet in Central Pb–Pb Collisions at* $\sqrt{s_{NN}} =$ 2.76 *TeV Measured with the ALICE Detector.* Lund University, Faculty of Science, Department of Physics, 2016.

[131] M. Angelsmark. Ξ *Production in pp Collisions at Different Multiplicity and Spherocity.* Lund University Libraries, LUP Student Papers, 8983483, 2019.

[132] ALICE Collaboration. Insight into particle production mechanisms via angular correlations of identified particles in pp collisions at $\sqrt{s} = 7$ TeV. *Eur. Phys. J. C* **77**, 569, 2017.

[133] Klaus Werner. Personal communication.

[134] V. Pacík. *Anisotropic azimuthal correlations of identified hadrons in p-Pb collisions at* $\sqrt{s_{\mathrm{NN}}} = 5.02\,\mathrm{TeV}$. Poster presented at the *17th International Conference on Strangeness in Quark Matter*, 2017.

[135] Y. Kanakubo, M. Okai, Y. Tachibana, and T. Hirano. Enhancement of strange baryons in high-multiplicity protonproton and protonnucleus collisions. *Prog. Theor. Exp. Phys.* **2018**, 121D01, 2018.

[136] M. Renardy & R. Rogers. *An Introduction to Partial Differential Equations.* Springer-Verlag, 2004.