

Study of Parton Energy Loss in Heavy Ion Collisions using Charged Particle Spectra Measured with CMS

by

Austin A. Baty

B.S., Texas A&M University (2013)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2019

© Massachusetts Institute of Technology 2019. All rights reserved.

Author
Department of Physics
May 1, 2019

Certified by.....
Gunther M. Roland
Professor of Physics
Thesis Supervisor

Accepted by.....
Nergis Mavalvala
Associate Department Head of Physics

Study of Parton Energy Loss in Heavy Ion Collisions using Charged Particle Spectra Measured with CMS

by

Austin A. Baty

Submitted to the Department of Physics
on May 1, 2019, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

The phenomenology of the strong nuclear force is still not well understood at low momentum transfers and requires experimental input to constrain. Collisions of heavy ions at the Large Hadron Collider provide a unique opportunity to explore this kinematic region because they create a novel form of matter: the quark-gluon plasma (QGP). Using the CMS detector, spectra of charged particles originating from proton-proton (pp), proton-lead (pPb), and lead-lead (PbPb) collisions at a center of mass energy per nucleon pair ($\sqrt{s_{NN}}$) of 5.02 TeV are examined as a function of transverse momentum and centrality. Nuclear modification factors and fragmentation functions are constructed from these spectra. By comparing to pp collision reference spectra, a puzzle concerning previous measurements in pPb collisions is clarified. A strong suppression of particle production observed in PbPb collisions is also quantified. Finally, collisions of xenon nuclei are also studied to constrain the path length dependence of parton energy loss. The strength of energy loss is found to increase with both $\sqrt{s_{NN}}$ and the average path length through the QGP. Comparisons to theoretical models and previous measurements indicate that the path length dependence is between linear and quadratic, as expected from a combination of collisional and radiative energy loss mechanisms.

Thesis Supervisor: Gunther M. Roland
Title: Professor of Physics

Acknowledgments

The MIT heavy ion group is a spectacularly unique group of people that I'm glad to have had the pleasure to work with. I would first like to thank my advisor, Gunther Roland, for getting me excited about heavy ion collisions in the first place when I was visiting MIT for the first time six years ago. He made me feel welcome in the group, provided valuable guidance to a fresh graduate student, and helped push me to where I am today. Thanks to Yen-Jie Lee for sharing his excitement of heavy ion physics and always making me feel like there is another new and exciting question just around the corner. I'm also thankful to Yen-Jie for a huge amount of advice about physics analysis that has accumulated over the years. Thanks to Bolek Wyslouch for remembering to look out for the MIT graduate students, even while they were at CERN. I want to thank Ivan Cali for his guidance on the various tracker-related projects I had once I got to CERN, and also for the very comfy chair. I also want to thank Camelia Mironov for the interesting conversations in the CERN office, sharing her excellent wisdom, and her tasty baked goods. Thanks also to Gian Michele Innocenti, Kristian Krajczar, and Yue-Shi Lai for their advice and friendship while I was at CERN. I also would be remiss to ignore the instrumental contributions Kristian and Yue-Shi made to the analyses discussed in this thesis.

That brings me to the most important part of this section: the other students. I'm indebted to Alex Barbieri, Dragos Velicanu, Doga Gulhan, and Mukund Varma for their patience and tutelage while I was a new student who didn't know how to use ROOT. Particular thanks goes to Doga, who helped me out when I first arrived at CERN and taught me all about the CMS tracking. Thanks to Chris McGinn for his friendship over the last six years, and all the crazy experiences we had in class, at conferences, working on analyses, and while staying at CERN. I'm also grateful for the companionship of Jing Wang, Ta-Wei Wang, Kaya Tatar, and Bi Ran, especially during the stressful heavy ion Runs every November. Finally I want to thank some of the newer students: Michael Peters, Zhaozhong Shi, and Anthony Badea for keeping things interesting at CERN even as I became more experienced.

I want to acknowledge the continued love and support I receive from all my family back home in Texas (and Idaho). Finally, thanks to Andrew Brown, Daniel Mann, Jonathan Warner, and Cole Spencer for continued friendship and entertainment as I travel all over the world.

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 23 |
| 1.1 | The Standard Model | 24 |
| 1.2 | QED and QCD | 25 |
| 1.2.1 | Hadron Collisions | 29 |
| 1.2.2 | Hadron Structure | 29 |
| 1.2.3 | Hard QCD Processes and Parton Showers | 30 |
| 1.2.4 | QCD Factorization | 32 |
| 1.3 | The Quark-Gluon Plasma | 33 |
| 1.3.1 | The Evolution of a Heavy Ion Collision | 37 |
| 1.3.2 | Jet Quenching | 38 |
| 1.3.3 | Hydrodynamic Flow | 41 |
| 1.3.4 | Thermal Photons | 43 |
| 1.3.5 | Strangeness Enhancement | 43 |
| 1.3.6 | Quarkonia Suppression | 44 |
| 1.3.7 | Cold Nuclear Matter Effects | 46 |
| 1.4 | Charged Particle Nuclear Modification Factors | 48 |
| 1.4.1 | Theoretical Models | 50 |
| 1.4.2 | Experimental Status | 56 |
| 1.4.3 | Questions Addressed in this Work | 64 |
| 2 | The CMS Detector | 67 |
| 2.1 | Pixel Tracker | 69 |
| 2.2 | Strip Tracker | 72 |

| | | |
|----------|--|------------|
| 2.3 | Electromagnetic Calorimeter | 77 |
| 2.4 | Hadronic Calorimeter | 78 |
| 2.5 | Trigger and Data Acquisition | 80 |
| 3 | Event Reconstruction | 83 |
| 3.1 | Centrality | 83 |
| 3.2 | Tracking | 85 |
| 3.2.1 | The CMS Tracking Strategy | 87 |
| 3.2.2 | Hit reconstruction | 89 |
| 3.2.3 | Track Seeding | 89 |
| 3.2.4 | Track Propagation and Fitting | 96 |
| 3.2.5 | Track Selection | 100 |
| 3.2.6 | Track Merging | 101 |
| 3.2.7 | Tracking Performance | 101 |
| 3.3 | Vertexing | 108 |
| 3.3.1 | Input Tracks | 108 |
| 3.3.2 | Track Clustering | 109 |
| 3.3.3 | Fitting | 111 |
| 3.3.4 | Resolution | 111 |
| 3.4 | Particle Flow and Calorimeter Matching | 113 |
| 3.4.1 | Input Candidates | 114 |
| 3.4.2 | Linking | 115 |
| 3.4.3 | Energy Calibration | 116 |
| 3.4.4 | Calorimeter Matching Performance | 118 |
| 3.5 | Jets | 118 |
| 4 | Data and Monte Carlo Samples | 123 |
| 4.1 | Monte Carlo Generators | 123 |
| 4.2 | Glauber Model | 127 |
| 4.3 | Data Samples | 133 |
| 4.4 | Event selection | 135 |

| | | |
|----------|--|------------|
| 4.5 | Triggers | 137 |
| 5 | pPb Fragmentation Function | 141 |
| 5.1 | Analysis Procedure | 141 |
| 5.1.1 | Jet Selection | 142 |
| 5.1.2 | Track Selection | 142 |
| 5.1.3 | Tracking Correction | 143 |
| 5.1.4 | Underlying Event Subtraction | 145 |
| 5.2 | Interpolated Reference | 147 |
| 5.3 | Systematic Uncertainties | 148 |
| 5.4 | Results | 151 |
| 6 | pPb and PbPb Nuclear Modification Factors | 161 |
| 6.1 | Analysis Procedure | 161 |
| 6.1.1 | Track Selection | 162 |
| 6.1.2 | Tracking Correction | 163 |
| 6.1.3 | Trigger Combination | 165 |
| 6.1.4 | Particle Composition Correction | 169 |
| 6.1.5 | Normalization | 173 |
| 6.2 | Systematic Uncertainties | 174 |
| 6.3 | Results | 179 |
| 6.3.1 | PbPb Nuclear Modification Factor | 179 |
| 6.3.2 | pPb Nuclear Modification Factor | 186 |
| 7 | XeXe Nuclear Modification Factor | 191 |
| 7.1 | Analysis Procedure | 192 |
| 7.2 | Extrapolation Procedure | 193 |
| 7.3 | Systematic Uncertainties | 195 |
| 7.4 | Results | 201 |
| 7.5 | Discussion | 215 |
| 8 | Conclusions | 219 |

List of Figures

| | | |
|------|---|----|
| 1-1 | A pictorial representation of the fundamental particles in the Standard Model of particle physics. | 24 |
| 1-2 | The allowed single-vertex Feynman diagrams in QED where a photon couples to charged particles (left), and QCD where quarks and gluons can couple to each other (right 3 panels). | 26 |
| 1-3 | The strength of the strong coupling α_s as a function of momentum transfer Q | 27 |
| 1-4 | The effective potential $V(r)$ between two interacting heavy quarks as a function of their separation. For large separation values, the potential increases roughly linearly. | 28 |
| 1-5 | An example of a parton shower, in which one energetic quark radiates a series of gluons that cascade into multiple partons. | 32 |
| 1-6 | The energy density and three times the pressure divided by the fourth power of the temperature, of QCD matter. A cross over transition can be seen around 170 MeV. | 35 |
| 1-7 | The phase diagram for nuclear matter. | 36 |
| 1-8 | An illustration of the temporal evolution of a heavy ion collision. | 37 |
| 1-9 | A 2.76 TeV PbPb event display from CMS showing a dijet event in which one jet has been heavily quenched. | 40 |
| 1-10 | A measurement by the ATLAS Collaboration of the p_T ratio of leading and subleading jet in 2.76 TeV PbPb and pp Collisions. The large difference between the two distributions is due to jet energy loss. | 41 |

| | | |
|------|---|----|
| 1-11 | The near side ridge observed in PbPb (left), pPb (middle), and pp collisions (right) by the CMS and ATLAS Collaboration. | 42 |
| 1-12 | A measurement by the ALICE Collaboration of the yields of photons coming from 2.76 TeV PbPb collisions. | 44 |
| 1-13 | The ALICE Collaboration's measurement of strange-hadron-to-pion ratios, showing the enhancement of strange baryon production as a function of collision size. | 45 |
| 1-14 | A measurement by the CMS Collaboration of the relative yields of the $\Upsilon(1S)$ and $\Upsilon(2S)$ mesons in PbPb and pp collisions. | 46 |
| 1-15 | The relative modification of the free proton PDF in the lead ion at $Q = 10$ GeV, as determined in nCTEQ15. The result is shown for gluons, and u , d , and s sea quarks. | 48 |
| 1-16 | An example of the $\sqrt{s_{NN}} = 200$ GeV measurements of the nuclear modification factor in AuAu and dAu collisions at RHIC. Figure taken from the BRAHMS Collaboration. | 58 |
| 1-17 | The nuclear modification factor of jets in pPb collisions, as measured with an extrapolated pp reference. A measurement by ATLAS is also shown. | 59 |
| 1-18 | Comparisons by the CMS Collaboration of the dijet pseudorapidity distribution in 5.02 TeV pPb and pp collisions. In the mid-rapidity region a strong anti-shadowing effect is seen for all p_T bins examined. | 60 |
| 1-19 | The nuclear modification factor for charged particles in pPb collisions, as measured with an interpolated pp reference. A model prediction and a result from ALICE are also shown. | 61 |
| 1-20 | A summary of some of the measurements of R_{AA} in the pre-RHIC era. | 62 |
| 1-21 | A compilation of the LHC charged particle R_{AA} measurements made at $\sqrt{s_{NN}} = 2.76$ TeV. | 63 |
| 1-22 | A summary of how the π^0 R_{AA} evolves as a function of center of mass energy, from SPS to LHC energies. | 64 |

| | | |
|-----|---|----|
| 2-1 | A cutaway view of CMS, showing the positions of the magnet and different subdetectors. | 68 |
| 2-2 | A schematic of the positions of the CMS tracker layers and disks in half of the tracker. The thick blue lines indicate sensors which give "stereo" hit information. The red lines at the bottom are the pixel detector. | 71 |
| 2-3 | A comparison of the original and phase one pixel detectors in CMS. The left panel shows a side view with the phase one detector on top and the original detector on bottom. An isometric view is shown in the right panel. | 72 |
| 2-4 | The simulated material budget of the CMS tracking system as a function of pseudorapidity. | 74 |
| 2-5 | An example of the readout from a set of 768 adjacent silicon strips. The raw signal is shown by the gray line, while the pedestal-subtracted data is given by the green line. The baseline (red) is distorted for the middle set of 128 strips, and must be reconstructed with a special algorithm. The final clusters after baseline subtraction are shown by the pink and blue lines. | 75 |
| 2-6 | A schematic comparing two different readout formats used in the CMS tracker. Each box outlined in black represents 1 byte. Grey fields indicate information from a strip. The 16-bit format wastes 6 bits of space per strip encoded, as indicated by the red boxes. A special 10-bit format for heavy ion data exchanges this for a small padding region (purple) before the footer of the data payload. | 76 |
| 2-7 | A schematic of the CMS trigger system. | 80 |
| 3-1 | A PbPb event collected by CMS in 2015, highlighting the objects resulting from reconstruction. The green lines are tracks originating from a common vertex, while the red and blue towers are calorimeter deposits that can be clustered into jets. | 84 |

| | | |
|-----|--|-----|
| 3-2 | The HF distribution in 2.76 TeV PbPb data, illustrating how centrality quantiles are determined. | 86 |
| 3-3 | An illustration of how hits in a toy 4-layer detector (top row) correspond to tracks (bottom row). Hits are shown by red circles. Tracks are black lines radiating from the center of the concentric rings, which are detector layers. For low detector occupancies (left column), it is fairly simple to find tracks, while the problem becomes very difficult as the occupancy grows (right column). | 88 |
| 3-4 | Pixel charge collection templates are shown for four different bias voltages. The calculation from PIXELAV is shown by the solid line, and a measurement with a test beam is given by the black dots. | 90 |
| 3-5 | An illustration of the CA algorithm in a toy 5-layer detector. (Left) Hits in all layers, shown by the blue dots, are connected to form cells, represented by the black lines. (Right) An iterative procedure gives the cells scores. Here black corresponds to a score of 0, red to 1, orange to 2, green to 3, and blue to 4. Seeds are then found by starting at the outside and moving inward. | 92 |
| 3-6 | A cartoon of a Kalman filter procedure. Hits are shown as red circles. The predicted trajectory and its uncertainty are shown by the blue arrow and the shaded triangular regions. A example of a track producing these hits is shown by the dashed green line. | 99 |
| 3-7 | The track p_T resolution for XeXe collisions for two different track selections. The ‘All Trk Selections’ points refer to the selections given for XeXe data in Table 5.1. | 102 |
| 3-8 | The DCA resolutions for 2015 PbPb tracks. | 103 |
| 3-9 | (Left) The tracking efficiency for 2015 5.02 TeV pp collisions as a function of p_T . (Right) The efficiency for inclusive 2015 PbPb collisions. Typical analysis selections have been applied. | 104 |

| | | |
|------|--|-----|
| 3-10 | The tracking efficiency for 5.44 TeV XeXe collisions. The efficiency is given for six different centrality selections. The colored bands show the statistical uncertainty. | 105 |
| 3-11 | (Left) The quantity $1/(1-f)$ for 2015 5.02 TeV pp collisions as a function of p_T . (Right) The same quantity for inclusive 2015 PbPb collisions. Typical analysis selections have been applied. | 106 |
| 3-12 | The value $(1 - f)$ for six different centrality selections in 5.44 TeV XeXe collisions. The results from three different types of simulation are shown. The fake rate becomes very large for central collisions. | 107 |
| 3-13 | The primary vertex spatial resolutions in x and z for pp collisions at 5.02 TeV as a function of the number of tracks in the event. | 112 |
| 3-14 | The primary vertex spatial resolutions in x and z as a function of centrality for PbPb and XeXe collisions. | 113 |
| 3-15 | A schematic of a set of two PF HCAL clusters. The dark gray squares are the two cluster seeds, surrounded by smaller cluster deposits. The red dots are the result of the cluster center calculation. The bottom of the green lines are where two tracks enter the HCAL and leave the deposit of energy. | 115 |
| 3-16 | The calibration coefficients used to fine tune the PF cluster energy. | 117 |
| 3-17 | (Right) The distribution of E_T/p_T as a function of centrality in PbPb PYTHIA+HYDJET. (Left) The relative efficiency from applying a calorimeter-matching required as a function of the track p_T in the same sample | 119 |
| 3-18 | The output of an anti- k_t jet clustering algorithm on a toy event. | 121 |

| | | |
|-----|--|-----|
| 4-1 | Results from Glauber MC simulations, showing the initial state geometry for different collision systems. Yellow and orange circles represent spectator nucleons from each nucleus, while pink circles are wounded nucleons. The collision systems are PbPb (Top left), XeXe (Top right), pPb (Bottom left), and PbPb (Bottom right). Produced with TGLAUBERM v3.2. | 130 |
| 4-2 | Turn-on curves for 2.76 TeV pp jet triggers. | 138 |
| 4-3 | Turn-on curves for 5.02 TeV pPb jet triggers. | 139 |
| 4-4 | Turn-on curves for 7 TeV pp jet triggers. | 139 |
| 4-5 | Turn-on curves for 5.02 TeV PbPb jet and track triggers, as well as the peripheral centrality gating. | 140 |
| 5-1 | The tracking nonclosure for 2.76 TeV pp (left), 5.02 TeV pPb (middle), and 7 TeV pp (right) collisions. | 145 |
| 5-2 | The tracking nonclosure for two different event selections in 7 TeV pp collisions. The selections were chosen to study the effect of pileup on the tracking efficiency. | 146 |
| 5-3 | The fraction of gluon-initiated jets as a function of either generator-level or reconstruction-level jet p_T , as calculated by PYTHIA. | 149 |
| 5-4 | A graphical summary of the systematic uncertainty contributions for the 5.02 TeV pPb (left) and pp reference (right) fragmentation functions. (Middle) the uncertainty on the ratio is also shown. These values are for the $140 < p_T^{jet} < 200$ bin. | 152 |
| 5-5 | The fragmentation functions for 2.76 TeV pp (top row) and 7 TeV pp (bottom row) collisions. The selected jet p_T increase from left to right in each of the panels. | 157 |
| 5-6 | The fragmentation functions for 5.02 TeV pPb collisions (top row) and an interpolated 5.02 TeV pp reference (bottom row). The selected jet p_T increase from left to right in each of the panels. | 158 |

| | | |
|-----|---|-----|
| 5-7 | The ratio of the pPb and pp reference fragmentation functions as a function of track p_T . The selected jet p_T increase from left to right in each of the panels. | 159 |
| 5-8 | The ratio of the pPb and pp reference fragmentation functions as a function of ξ . The selected jet p_T increase from left to right in each of the panels. | 160 |
| 6-1 | (Left) The distribution of $\chi^2/((N_{layers})(d.o.f))$ in 30–100% PbPb collisions for data and MC. (Right) The required amount of scaling of the MC distribution to achieve adequate agreement with data. | 163 |
| 6-2 | The tracking efficiency for PbPb collisions as a function of track p_T and collision centrality. The values are calculated with PYTHIA+HYDJET. | 164 |
| 6-3 | Distributions of the leading jets for pp collisions after each trigger has been weighted by an effective prescale. Refer to the main text for more details. | 167 |
| 6-4 | Ratios of the distributions of the leading jets for pp collisions after each trigger has been weighted by an effective prescale. | 168 |
| 6-5 | (Top) The uncorrected total track invariant yields for pp collisions, as calculated with jet (left) and track (right) triggers. (Bottom) The relative contribution of each trigger to the spectra shown in the top row. | 169 |
| 6-6 | (Left) Ratios of the counts from various jet triggers after weighting each one by its ‘relative prescale’ in 0–30% PbPb collisions. (Right) The resulting uncorrected track spectrum in the 0–5% centrality range. | 170 |
| 6-7 | (Left) Ratios of the counts from various track triggers after weighting each one by its ‘relative prescale’ in 0–30% PbPb collisions. (Right) The resulting uncorrected track spectrum in the 0–5% centrality range. | 171 |
| 6-8 | (Left) The tracking efficiency as a function of p_T for the most common particles produced in hadronic collisions. (Right) The tracking efficiency for various strange baryons. | 171 |

| | | |
|------|---|-----|
| 6-9 | The relative fraction of charged particles coming from various particle species, as a function of p_T , for three different MC generators. The particles are, from the top left panel to the bottom right panel, π^\pm , K^\pm , p , \bar{p} , Σ^+ , Σ^- , Ξ^- and Ω^- | 177 |
| 6-10 | The particle species residual correction for 5.02 TeV PbPb collisions (left) and pp collisions (right). | 178 |
| 6-11 | A comparison of the PbPb spectra in the 10–30% and 50–100% centrality bins, after varying the track selections to looser criteria. | 178 |
| 6-12 | The charged particle spectra for 5.02 TeV PbPb and pp collisions. The pp data has been scaled by a constant factor of 70 mb so it can be directly compared to the PbPb measurement. | 180 |
| 6-13 | The nuclear modification factor in 5.02 TeV PbPb collisions, for six different centrality selections. Previous results at 2.76 TeV are shown by the red, blue, and green data points. | 182 |
| 6-14 | Comparison of the nuclear modification factor to six theoretical predictions in the 0–10% and 30–50% centrality bins. | 183 |
| 6-15 | A summary of R_{AA} measurements and predictions for central collisions. | 184 |
| 6-16 | A selection of heavy flavor measurements from CMS, compared to the measurement of the nuclear modification of inclusive charged particles described here. The charged particle measurement is a crucial reference for interpreting these results. | 186 |
| 6-17 | The updated result of R_{pA} using a measured 5.02 TeV pp reference, shown in by the points surrounded by green boxes. The inclusive R_{AA} for PbPb collisions is shown by the points with yellow boxes. | 187 |
| 7-1 | The extrapolation factor needed to calculate a pp reference at 5.44 TeV for comparison with XeXe data. | 195 |

| | | |
|-----|--|-----|
| 7-2 | A comparison of the d_0 DCA significance for XeXe data (black points) and MC (black line). The red and blue lines show the fraction of read and fake tracks making up the MC, respectively. The unscaled data is shown in the left panel, while MC that has had the fake fraction scale up to match data in a side band region is shown on the right panel. | 197 |
| 7-3 | The relative difference in the XeXe particle spectra after varying the track selections to tighter and looser cuts. The left panel shows 0–10% events, while the right panel shows 50–70% events. Refer to the text for explanations about selections D and E. | 198 |
| 7-4 | A graphical summary of the systematic uncertainties related to measurements of the 0–5% XeXe spectra (top), 0–5% XeXe R_{AA} (middle), and 0–5% R_{Pb}^{Xe} (bottom). | 200 |
| 7-5 | The charged particle spectra for 5.44 TeV XeXe and pp collisions. The pp data has been scaled by a constant factor of 70 mb so it can be directly compared to the XeXe measurement. | 201 |
| 7-6 | The nuclear modification factor in 5.44 TeV XeXe collisions, shown by the black points with pink boxes for six different centrality selections. The nuclear modification factor in 5.02 TeV PbPb collisions is shown by the blue points. | 203 |
| 7-7 | The N_{coll} scaled differential charged particle yields as a function of N_{coll} in two different MC generators for two different p_T selections. Two sets of lines are shown, one that uses the generator-level N_{coll} values (solid lines), and a second that uses the $\langle N_{coll} \rangle$ of a centrality bin calibrated in a fashion similar to data. The two curves differ for low values of N_{coll} , which correspond to peripheral events. | 205 |
| 7-8 | An estimation of the potential bias to R_{AA} caused by the effects of the centrality calibration procedure. Central events are largely unbiased, but peripheral events can be strongly affected. The magnitude of the bias depends on the MC model used in the calculation. | 207 |

| | | |
|------|---|-----|
| 7-9 | A measurement of R_{Pb}^{Xe} for five centrality selections. The blue line shows the unmodified expectation from the difference in center of mass energies between the numerator and denominator. | 209 |
| 7-10 | The values of R_{AA} and R_{AA}^* for PbPb and XeXe collisions as a function of N_{part} for two p_T selections. The left panel shows charged particles at the local minimum in R_{AA} , while the right panel shows particles that are in the middle of the rising high p_T trend. | 210 |
| 7-11 | Measurements of R_{Pb}^{Xe} comparing different centrality bins for XeXe and PbPb collisions. The relevant centralities were selected to have similar N_{part} values. The left panel shows the result for central XeXe collisions, while the right panel shows the measurement in peripheral collisions. | 211 |
| 7-12 | A comparison of R_{AA}^* for both 0–10% (left) and 30–50% (right) XeXe collisions with various theoretical models. | 212 |
| 7-13 | A comparison of four theoretical models to measurements of R_{Pb}^{Xe} . The 5–10% centrality range is shown on the left, while the 30–50% range is shown on the right. | 214 |
| 7-14 | The average energy loss of a fast parton, as calculated in a phenomenological study using the data from this measurement. | 216 |
| 7-15 | An example of the tools being developed by the JETSCAPE Collaboration to compare charged hadron spectra to theoretical models that have similar parameters. Here three energy loss models having the same hydrodynamic background and hadronization code are compared to 2.76 TeV CMS results. | 218 |

List of Tables

| | | |
|-----|---|-----|
| 1.1 | Previous measurements of the nuclear modification factor for light particle species. The first grouping consists of pA and dA collisions, while the second grouping is AA collisions. | 66 |
| 3.1 | The track seeding iterations for the pp and pPb data used in this work. Here ‘BS’ and ‘PV’ refer to tracking regions centered on the beam spot and primary vertex, respectively. Quantities given in units of σ refer to a number of standard deviations of the primary vertex or beam spot uncertainty in the z direction. | 94 |
| 3.2 | The track seeding iterations for the PbPb and XeXe data used in this work. Here ‘BS’ and ‘PV’ refer to tracking regions centered on the beam spot and primary vertex, respectively. Quantities given in units of σ refer to a number of standard deviations of the primary vertex or beam spot uncertainty in the z direction. $\Delta\theta$ and $\Delta\phi$ are provided when Cellular Automaton seeding is employed. | 95 |
| 3.3 | The track selections used during primary vertex reconstruction. | 109 |
| 4.1 | The Glauber model settings used for modeling the heavy nuclei in pPb, PbPb, and XeXe collisions. | 129 |
| 4.2 | The values of $\langle N_{\text{part}} \rangle$, $\langle N_{\text{coll}} \rangle$, T_{AA} , and their uncertainties, for $\sqrt{s_{NN}} = 5.44$ TeV XeXe collisions and 5.02 TeV PbPb collisions in the centrality ranges used. The relevant values for inclusive pPb at 5.02 TeV collisions are also given. | 132 |
| 4.3 | The data samples used in this thesis. | 133 |

| | | |
|-----|---|-----|
| 5.1 | The analysis level track selections used for various data sets. | 154 |
| 5.2 | Summary of the systematic uncertainties for the 5.02 TeV pPb fragmentation functions. The p_T values are given in units of GeV. | 155 |
| 5.3 | Summary of the systematic uncertainties for the ratio between pPb and pp fragmentation functions. The p_T values are given in units of GeV. | 156 |
| 6.1 | Event classes for the measurement of charged particle spectra in pp collisions. | 166 |
| 6.2 | Event classes for the measurement of charged particle spectra in PbPb collisions. | 170 |
| 6.3 | Systematic uncertainties associated with the measurement of the charged-particle spectra and R_{AA} using $\sqrt{s_{NN}} = 5.02$ TeV pp and PbPb collision data. The ranges quoted cover both the p_T and the centrality dependence of the uncertainties. The combined uncertainty in R_{AA} does not include the integrated luminosity and the T_{AA} uncertainties. | 174 |
| 7.1 | The systematic uncertainties related to the XeXe measurements reported here. The values quoted cover the centrality and p_T dependence of each uncertainty. They are separated into normalization uncertainties and all other systematic uncertainties. | 196 |

Chapter 1

Introduction

Over the last century, we have made great strides in distilling the fundamental laws of the universe down to their very essence. Over time many classical theories have given way to quantum mechanics and, eventually, quantum field theories. We are now able to describe most of the interactions in the universe using a relatively small set of quantum fields that interact according to rules that are well-defined, albeit with seemingly stochastic strengths. Yet we know there are still gaps in our understanding, and we are still pushing to understand how matter and energy behave in extreme conditions.

Historically we have studied the structure of atomic nuclei, how they are held together and how they decay. Such studies can be seen as a vital bridge that seeks to understand how the fundamental laws discovered by high energy physicists give rise to the complex objects that sit at the center of every atom in the universe. But this bridge still has some gaps in it. It is difficult to perform calculations with the theory of the most relevant force — the strong nuclear force — when small momentum transfers between the fundamental constituents of a proton (or neutron) are involved. It may be possible to fill in these gaps, or at least better understand them, by using the particle physics strategy of subjecting nuclear matter to extreme conditions.

The field of relativistic heavy ion physics uses collisions of atomic nuclei to attempt to understand the behavior of the strong nuclear force at low momentum transfers. As will be discussed in Section 1.3, a unique form of matter — the quark-gluon plasma




Figure 1-1: A pictorial representation of the fundamental particles in the Standard Model of particle physics. Image courtesy of CERN [1].

(QGP) — is created in these collisions. This system may be the smallest system ever studied to exhibit complex emergent behavior, despite being governed by a set of relatively simple laws that can be written down on a single sheet of paper. It is also the hottest man-made form of matter in the universe. This thesis explores a small portion of the rich dynamics found in this state of matter.

1.1 The Standard Model

The Standard Model of particle physics is perhaps the most successful theoretical model in history. It has been able to describe the behavior of three of the four known fundamental forces — electromagnetism, the weak nuclear force, and the strong nuclear force — with astounding accuracy. Remarkably, it accomplishes this by describing the interaction of a relatively small number of fundamental particles. A pictorial representation of the various components of the Standard Model is shown in Fig. 1-1.

Five vector bosons — particles having a spin of 1 — are shown in a group on the right side of the figure. These particles give rise to the fundamental forces described

by the Standard Model. The photon, γ , mediates the electromagnetic force between particles having electric charge. The W^\pm and Z^0 bosons mediate the weak nuclear force between particles based on the values of quantum numbers known as weak isospin and weak hypercharge. The gluon, g , gives rise to the strong nuclear force for particles having color charge.

The twelve fermions — particles having a spin of one-half — are shown on the left hand side of the same figure. These are split into two groups, each of six particles. The first group are known as quarks. These particles have both electric charge, weak isospin and hypercharge, and color charge. They therefore interact with all three forces. The up, down, and strange quarks have relatively light masses and are fairly common in regular matter, while the heavier charm, bottom, and top quarks are much rarer. The second group of six particles are leptons. These particles consist of the electrically charged electron, muon, and tau. An electrically neutral neutrino is paired with each of these particles, and interacts with it via the weak nuclear force. Leptons have no color charge and therefore do not interact with the strong nuclear force. All fermions have a corresponding antimatter partner that has opposite electric charge.

The final component of the Standard Model is the recently-discovered Higgs boson. This boson interacts with all the other particles, except the photon and gluon, to give them mass [2, 3]. Although all components of the Standard Model are needed to describe the wide range phenomena in our universe, this work focuses primarily on quarks and gluons. Furthermore, heavy quarks are sufficiently rare that most of the relevant dynamics for this thesis can be described by just the up, down, and strange quarks as well as the gluon.

1.2 QED and QCD

The modern theory of electromagnetism, Quantum Electrodynamics (QED), has been known for around sixty years. The theory describes the fundamental interaction between electrically charged fermions and photons. A Feynman diagram for this funda-




Figure 1-2: The allowed single-vertex Feynman diagrams in QED where a photon couples to charged particles (left), and QCD where quarks and gluons can couple to each other (right 3 panels).

mental interaction can be seen in the left panel of Fig. 1-2. The wavy line represents the photon, while the straight line represents a fermion. Using perturbation theory, any QED scattering process can be expanded as a sum of scattering amplitudes involving a different number of these fundamental interactions, which show up as vertices in a Feynman diagram. Each vertex weights the contribution of the diagram to the total amplitude by an extra factor of the fine structure constant $\alpha \approx 1/137$. This means that diagrams involving many vertices have diminishing importance and the calculated result will quickly converge to the true value after only including a few terms in the sum. Thus, quantities can be calculated within any desired precision as long as sufficient computational power is available. Because of this, QED is responsible for some of the most accurate predictions made by the Standard Model, such as the anomalous magnetic moment of the electron.

Another key pillar of the Standard Model is the theory of the Quantum Chromodynamics (QCD), which describes how the force-carrier of the strong nuclear force, the gluon, interacts with quarks. In many respects, QCD is similar to QED. However, this theory also comes with a twist: gluons can also directly interact with other gluons. This is very different from QED where the force-carrier photons can not directly interact with other photons. The right three panels in Fig. 1-2 show the allowed fundamental interactions. Curly lines represent gluons. Unlike QED, interaction vertices of three or four gluons are possible. The consequences of this gluon-gluon interaction are far-reaching and profound. In 1973 it was shown that a model with this property, along with three types of color charge, would have a coupling strength α_s ,




Figure 1-3: The strength of the strong coupling α_s as a function of momentum transfer Q . Taken from Ref. [6].

that rapidly increases for low values of momentum transfer Q^2 , but is fairly small for large values of Q^2 [4, 5]. A plot of the current understanding of this evolution is shown in Fig. 1-3. The so-called *running* of α_s implies that perturbative methods are still valid computational tools in QCD, as long as the characteristic momentum transfer of the interaction process is large enough such that $\alpha_s < 1$. This property is known as *asymptotic freedom*, and the theory in this kinematic region is known as perturbative QCD. On the other hand, at low Q^2 the coupling constant becomes large and perturbative methods break down. Calculations are notoriously difficult to perform at these values of Q^2 , and experimental input is required to constrain the dynamics of the theory.

For low Q^2 the gluon field between two quarks interacts with itself strongly and tends to form a collimated color flux tube. As the distance between the two quarks increases, the force between them stays constant rather than decreasing. This means a large amount of energy is stored in the gluon field between the two quarks if they are




Figure 1-4: The effective potential $V(r)$ between two interacting quarks as a function of their separation. For large separation values, the potential increases roughly linearly. Figure from Ref. [7].

significantly separated in space. An illustration can be seen in Fig. 1-4, which shows the effective potential between two interacting heavy quarks as a function of their separation. In the region of large separation, the potential increases approximately linearly. At some point, it becomes energetically favorable for a quark anti-quark pair to be created and break the color flux tube. This mechanism prevents the presence of an isolated color-charged object in nature; all color-charged particles must be nearby each other to prevent the formation of large gluon fields. This phenomenon is known as *color confinement* and causes quarks and gluons to coalesce into composite particles that must be color neutral or color singlets (having no net color charge). Such composite particles are known as hadrons. Protons, kaons, and pions are all examples of hadrons.

A majority of the physics processes probed at machines such as the Large Hadron Collider (LHC) are governed by QCD. Hadron-hadron collisions can result in multiple scattering processes that produce many partons. However, the cocktail of partons produced by a collision must ultimately result in a collection of color neutral hadrons

because of confinement. The procedure by which these partons become grouped together into a set of final particles is known as *hadronization*. Because it is a low- Q^2 QCD process, the exact dynamics of this hadronization process are not well understood and still require experimental input.

1.2.1 Hadron Collisions

High-energy hadron collisions are one of our best tools for creating nuclear matter in extreme conditions. To better understand these collisions, both the initial structure of the colliding particles, as well as the actual scattering process itself, must be described. In the parton model, the colliding hadrons are composite particles created out of many interacting quarks and gluons, which are collectively known as partons. In general many complex scattering processes at a variety of Q^2 values can occur in a hadron collision, but occasionally two partons, one from each of the incoming hadrons, undergo a high- Q^2 interaction. In the following sections a framework for calculating the cross sections of these hard scatterings is discussed. It involves probabilistic descriptions of the initial kinematics of the two scattering partons, as well as the fate of the outgoing partons. These two descriptions are then combined with other perturbative QCD calculations using an assumption known as *QCD factorization*.

1.2.2 Hadron Structure

In a simplified picture where only the hardest partonic interaction is considered, a hadron collision can be viewed as a collision of two partons that live inside each hadron. For high energy collisions, this means the longitudinal momentum structure of the partons inside the hadrons is of great importance for the resulting dynamics of the system. To describe this structure, functions known as parton distribution functions (PDFs) are defined. These functions are denoted as $f_i(x, Q^2)$. There are separate PDFs for each type of parton that makes up the structure of the hadron, and these are indicated by the index i . The PDF itself depends on the fraction of the total hadron momentum (x) the parton has, as well as a characteristic energy

scale Q^2 . Because all protons are identical, the proton PDF is universal. It does not vary based on what type of scattering process is considered. These PDFs are non-perturbative objects because the average momentum exchange between partons in a hadron is relatively low. Therefore, they can not yet be calculated with high precision from first principles and must be constrained with experimental data. Precision deep inelastic scattering measurements have strongly constrained these functions [8, 9, 10] across a wide range of x and Q^2 .

1.2.3 Hard QCD Processes and Parton Showers

Once a particle collision occurs, a parton can be scattered with a large initial energy. This scattering can involve a very large transfer of momentum, and therefore may be described by perturbative QCD. This parton can then split into multiple partons, initiating what is known as a *parton shower*. A Feynman diagram of a parton shower is shown in Fig. 1-5 as an example. Each new parton may be able to radiate its own particles, causing a cascading effect. Using perturbative QCD, the probability P of a parton i splitting into two more partons j and k is given by

$$dP = \frac{\alpha_s}{4\pi} \frac{d^2 k_T}{k_T^2} P_{i \rightarrow jk}(z) dz. \quad (1.1)$$

In this equation k_T is the transverse momentum with respect to the original parton's direction, and z is the fraction of the total momentum carried by parton j (the parton k carries a momentum fraction $(1 - z)$). The quantities $P_{i \rightarrow jk}(z)$ are known as the Altarelli-Parisi splitting functions [11] and depend on which types of partons are considered:

$$P_{q \rightarrow qg}(z) = \frac{4}{3} \frac{1+z^2}{1-z} \quad (1.2)$$

$$P_{g \rightarrow gg}(z) = 3 \frac{(1-z)(1-z)^2}{z(1-z)} \quad (1.3)$$

$$P_{g \rightarrow qq}(z) = \frac{z^2 + (1-z)^2}{2}. \quad (1.4)$$

The subscripts q and g stand for quarks and gluons, respectively. It should be noted that the probability of splitting diverges as k_T approaches zero, e.g. for collinear parton emissions. Also, the splitting functions involving a gluon in the final state diverge if the gluon is soft, such that $(1 - z)$ approaches 0. This is known as an infrared divergence.

Quantum mechanically, when a soft gluon is emitted its initial wave function has a significant spatial overlap with the parent parton. During this time, the two partons are said to be coherent and emit further radiation as if they were a single object. The two partons decohere and behave like separate objects once their transverse separation is greater than $\lambda_{\perp} = 1/k_T = 1/(\omega\theta)$, where ω is the gluon's frequency and θ is the emission angle. The time it takes a gluon to decohere is known as the gluon formation time:

$$t_{form} = \frac{2k_z}{k_T^2} = \frac{2}{k_z\theta^2}. \quad (1.5)$$

Here k_z is the longitudinal momentum component of the gluon. The values of t_{form} are larger for small emission angle and small transverse momenta. This leads to an angular ordering (or k_T ordering) of the parton shower, where emissions at large angles are temporally resolved first.

As the shower evolves, the characteristic energy scale of each particle in the shower decreases towards the non-perturbative region. Because of confinement, the partons in the shower must then recombine to form a collection of color neutral hadrons that can be detected. This hadronization procedure is also described by non-perturbative QCD. Because of the collinear divergence of the splitting probability, the resulting hadrons are usually collimated into a spray known as a jet. Thus, most of the initially scattered parton's energy and momentum end up being carried by the hadrons making up the jet. For this reason, the kinematics of a jet are used as an approximate experimental proxy for those of the initially scattered parton.




Figure 1-5: An example of a parton shower, in which one energetic quark radiates a series of gluons that cascade into multiple partons.

1.2.4 QCD Factorization

For a QCD scattering that is sufficiently hard, i.e. having a momentum transfer large enough to enter into the perturbative region, the assumption of QCD factorization allows the separation of non-perturbative and perturbative contributions to the total scattering cross section. This assumption can be stated mathematically as [3, 12]

$$\sigma^{pp \rightarrow \text{hadrons}} = f_i(x_1, Q^2) \otimes f_j(x_2, Q^2) \otimes \sigma^{ij \rightarrow kl}(x_1, x_2, Q^2) \otimes D_k(z, Q^2) \otimes D_l(z, Q^2). \quad (1.6)$$

This equation involves an implicit summation over all possible partons denoted as i , j , k , and l . Here $\sigma^{pp \rightarrow \text{hadrons}}$ is the cross section for a scattering process between two protons producing some generic state of outgoing hadrons. The $f_i(x, Q^2)$ functions are the PDFs of the incoming hadrons, describing the probability of a parton i having a momentum fraction x at a given scattering energy scale Q^2 . The term $\sigma^{ij \rightarrow kl}(x_1, x_2, Q^2)$ is a scattering cross section that can be calculated fairly precisely using perturbative methods, given that the relevant value of $\alpha_s(Q^2)$ for the scattering is small enough. The scattering described is between a state having partons i and j

and resulting in a system of partons k and l .

The final terms, $D_k(z, Q^2)$, are known as *fragmentation functions* and describe how the parton, k , originating from the hard scattering showers and hadronizes. This function depends on the fraction of the total parton momentum, z , each of the final particles have. Because it involves hadronization, this is a non-perturbative process. It is usually assumed that this quantity is also universal, i.e. that a parton at a given Q^2 should fragment similarly regardless of how it was produced. The fragmentation functions of partons have been constrained using measurements of jets at the LEP e^+e^- collider [13].

The power of QCD factorization should now be apparent. Non-perturbative features of the scattering — the PDF and the fragmentation function — are sequestered into universal objects that can be experimentally measured across a wide kinematic range of x and Q^2 . A procedure known as DGLAP evolution can then be applied to extrapolate these discrete measurements into a continuous space that describes the proton structure and parton fragmentation at any desired value of Q^2 [11, 14, 15]. The remaining parts of the equation can be calculated directly using perturbative QCD techniques. The use of QCD factorization has been remarkably successful at describing jet and hadron production at hadron colliders, particularly in the region of high transverse momentum (p_T).

1.3 The Quark-Gluon Plasma

If a gas of atoms is heated to a sufficient temperature, the average kinetic energy of the atomic electrons can exceed the atomic binding energy caused by electromagnetic attraction. When this happens, the gas becomes ionized into a plasma. A similar phenomenon was originally thought to happen in matter interacting via QCD. It was believed that if a gas of hadrons was heated to a high enough temperature, the average strength of partonic interactions would decrease enough to allow the hadrons to melt into a gas of weakly interacting quarks and gluons. For this reason, this predicted state of matter was called a *quark-gluon plasma* (QGP).

To achieve the temperatures needed to produce a QGP, a large amount of energy needs to be compressed into a small volume. Collisions of heavy atomic nuclei at very high energies were identified as a way to achieve such energy densities. The nuclei are electromagnetically charged and can therefore be accelerated using conventional accelerator technologies. They also allow the simultaneous collision of hundreds of nucleons within a volume of around 1 fm^3 ; something that cannot currently be achieved with proton-proton collisions. Significant fixed-target heavy ion experimental programs were conducted at Brookhaven National Lab (BNL) Alternating Gradient Synchrotron (AGS) and the European Organization for Nuclear Research (CERN) Super Proton Synchrotron (SPS) between 1986 and around 2000. These experiments explored collisions of a variety of nuclear species up to center-of-mass energies per nucleon pair ($\sqrt{s_{NN}}$) of around 10 GeV (20 GeV) at the AGS (SPS). Many of these measurements found an enhancement in the fraction of strange hadrons produced in heavy ion collisions as compared to smaller collision systems. This signal was one of the predicted signs of the production of a QGP. In 2000 another signal — the suppression of J/Ψ meson yields — was observed at SPS. This prompted the announcement of the creation of ‘a new state of matter, in which quarks and gluons are “deconfined”’ [16]. Although this claim is debatable, subsequent studies at Relativistic Heavy Ion Collider (RHIC) confirmed the creation of a QGP [17]. However, the RHIC experiments also discovered that the QGP did not seem to behave like a weakly-interacting collection of partons. Instead it was observed that the initial geometric configuration of the nuclei in the collision strongly affected the spatial and momentum distributions of the final state particles. This indicated that a QGP behaves more like a fluid evolving according to the laws of hydrodynamics rather than a gas. Thus, although the quarks and gluons are liberated from hadrons, they still interact strongly.

QCD matter can be described by its temperature T and its baryon chemical potential (μ_b). The baryon chemical potential is the energy needed to increase the system’s baryon number by one unit (by adding more quarks), and scales with the system’s net quark density. Lattice QCD calculations [18] at $\mu_b = 0$ predict that




Figure 1-6: The energy density and three times the pressure divided by the fourth power of the temperature, of QCD matter. A cross over transition can be seen around 170 MeV. Figure from Ref. [18].

a phase transition between a hadron gas and a quark-gluon plasma is expected in QCD matter at a temperature of around 170 MeV. This phase transition is a smooth crossover transition without any discontinuities associated with it. Figure 1-6 shows a lattice QCD calculation of this transition by plotting the energy density and pressure of QCD matter as a function of the temperature. As the temperature continues to increase, the behavior of the QGP approaches that of a free gas (shown by the short black line in the top right corner of the figure) very slowly. At temperatures over 500 MeV, the QGP is expected to still interact relatively strongly.

A summary of the current understanding of the phase diagram of QCD matter is shown in Fig. 1-7. The horizontal axis shows the baryon chemical potential and the vertical axis is the system's temperature. It has been suggested that matter having low temperatures but high values of μ_b may be found inside of neutron stars. This could include a predicted form of matter known as a color superconductor. Looking at lower μ_b , most QCD matter is a hadron gas. This includes the familiar nuclear matter that resides at the center of atoms. At hot temperatures the matter is in a




Figure 1-7: The phase diagram for nuclear matter. Reproduced from Ref. [19].

QGP phase. These two phases are separated by a 1st order phase transition, except for low values of μ_b , where the transition is a smooth cross over. It is predicted that there exists a point on the phase diagram that separates the 1st order and crossover phase transitions. The exact location of this critical point is unknown, but of great interest. The transition between the QGP phase and the phase of matter found in neutron stars is also unknown; it is possible that a QGP exists at the center of neutron stars.

Heavy ion collisions at RHIC and the LHC create a hot QGP in the region around $|\eta| = |\ln(\tan(\theta/2))| \approx 0$, where θ is the polar angle away from the beam pipe. This is where the collision's energy density is largest. This QGP quickly expands and cools off. As the matter cools down, it returns to the hadron gas phase. The approximate trajectories in the phase diagram of a heavy ion collision at RHIC and the LHC are shown by the orange and yellow lines on Fig. 1-7. These collisions can be used to

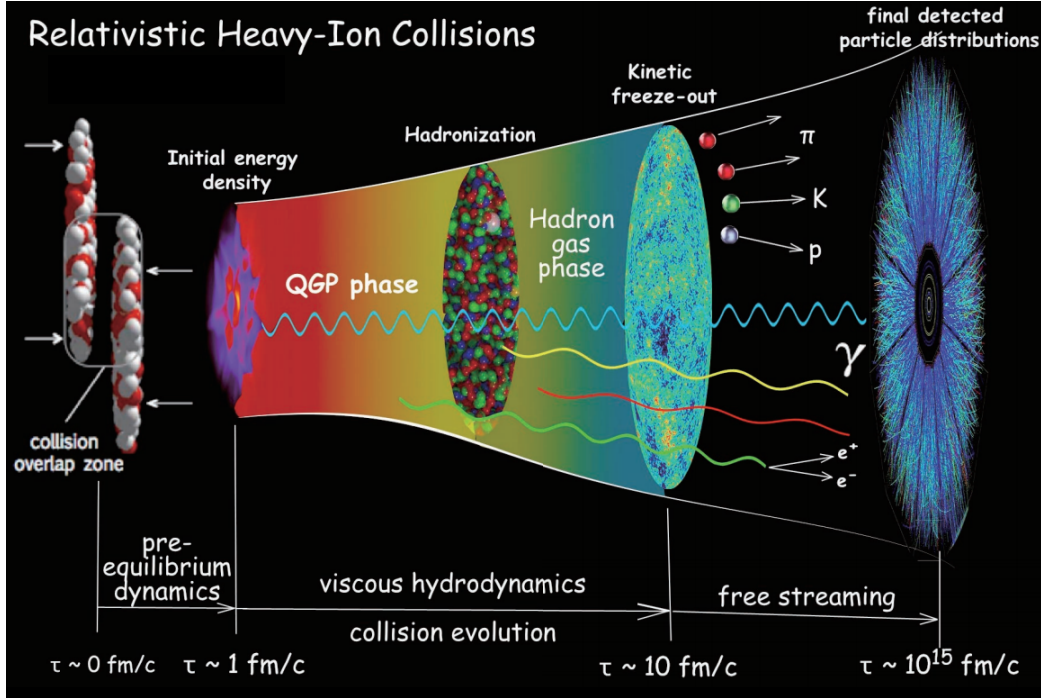


Figure 1-8: An illustration of the temporal evolution of a heavy ion collision. Image reproduced from Ref. [20].

study the unique properties of this strangely fluid QGP. Such an understanding is strongly connected to our knowledge of non-perturbative QCD and the properties of QCD matter.

1.3.1 The Evolution of a Heavy Ion Collision

A heavy ion collision produces a complex evolving system. Figure 1-8 shows a cartoon sketching the lifetime of a heavy ion collision. Initially, two large nuclei (usually gold or lead) moving at nearly the speed of light approach each other. Although these nuclei are roughly spherical when at rest, they are deformed into flat disks by relativistic Lorentz contraction. The nuclei do not always hit each other head-on. Sometimes, only a fraction of each nucleus hits each other. Thus, the size of the region where collisions occur varies on an event-by-event basis. As mentioned before, the partons in each nucleus scatter off each other. Importantly, this can sometimes produce very hard QCD scatterings which can be treated perturbatively using QCD

factorization. First, the system equilibrates as particles that were previously causally disconnected from each other begin to interact. Fits to particle spectra and hydrodynamic flow measurements (Section 1.3.3) indicate that this pre-equilibrium stage lasts for a little less than 1 fm/c. This results in a QGP with an energy density profile that is related to the initial geometry of the collision overlap region. This QGP evolves, expands, and cools for the next ≈ 10 fm/c. During this time the dynamics of partons are affected by the presence of the QGP, resulting in some of the unique signatures of the QGP. Once it has cooled enough, the partons in the QGP undergo hadronization and turn into a hadron gas. This gas has a much lower number of degrees of freedom than the QGP phase, because the individual partons of the system are sequestered into hadrons. The hadrons in this gas can still scatter off each other, jumbling up some of the kinematics of particles coming directly from the QGP. As this gas expands, it eventually reaches a ‘chemical freeze out’ point where the particle composition becomes fixed because the hadrons stop experiencing a significant amount of inelastic scatterings. Later a ‘kinetic freeze out’ point is reached. Here, the hadron gas does not have a significant number of *elastic* scatterings so the particles kinematics do not change greatly. The particles then stream outward into the particle detector where they are detected to produce the final event. Our task is to use the information present in the detector to unravel the dynamics of the different phases of this collision evolution. Some classic signatures that are studied in heavy ion collisions are listed below.

1.3.2 Jet Quenching

In e^+e^- , ep , and pp collisions high energy partons originating from a hard scattering initiate a parton shower that evolves in vacuum. However, in heavy ion collisions this parton shower will occur inside the bulk of the QGP. Because the QGP consists of color-charged objects, it can interact with the intermediate partons in the shower and modify the shower evolution. As the shower’s partons interact with the QGP, they lose energy. In a perturbative QCD framework, this energy loss can be classified as either collisional or radiative energy loss.

Collisional energy loss interactions are elastic scatterings between a parton from the shower and a parton from the QGP. These interactions tend to transfer energy from the parton shower into relatively low energy QGP partons. Thus, the total energy in the parton shower is decreased, while the energy of the QGP is slightly increased. Because the QGP is a strongly interacting fluid, this extra energy tends to quickly get thermalized and spread among the many partons of the QGP. Collisional energy loss is expected to scale linearly with the path length of the initial hard parton when traversing a large static medium.

Radiative energy loss occurs from inelastic processes where the hard parton radiates a gluon. The pattern of radiative energy loss can also be affected by the presence of a QGP. Soft scatterings with QGP partons can induce additional gluon radiation [21]. Furthermore, they can rotate the color charge of a parton, causing it to immediately start to radiate incoherently with the rest of the parton shower. This increases the amount of radiative loss in the parton shower, as compared to the vacuum evolution [22].

Because the partons in a shower lose more energy in the QGP relative to the vacuum, the average energy of hadrons resulting from the hadronization of these partons will be lower. This means that the presence of a QGP should suppress the yield of high energy hadrons. Thus, a jet that is created by clustering the hadrons observed in an experimental apparatus will have a lower energy than a jet that did not suffer any energy loss. This effect is known as *jet quenching* and is a direct result of parton energy loss [23].

Evidence for jet quenching was first observed at RHIC by examining spectra and correlations of high energy hadrons in gold-gold collisions. In general, it was found that the yields of high energy hadrons were suppressed compared to an appropriately scaled reference of proton-proton collisions. This suppression is discussed in detail in Section 1.4.2. It was also observed that hadrons recoiling from a hard scattering tended to have lower energy [24, 25, 26] than they did in proton-proton collisions. Similar effects have been observed for hadrons by the ALICE Collaboration at the LHC [27].




Figure 1-9: A 2.76 TeV PbPb event display from CMS showing a dijet event in which one jet has been heavily quenched. Figure from Ref. [28].

The full calorimeter coverage of the CMS and ATLAS experiments enable clear measurements of the quenching of very high energy jets at the LHC. This was first seen by examining pairs of back-to-back jets (dijets) recoiling off each other. These jets should have approximately the same p_T because of conservation of momentum, but they were observed to have strong asymmetries because one of the two jets was more strongly quenched than the other [28, 29, 30]. An event display showing a dijet event where one of the two jets nearly disappears is shown in Fig. 1-9. These asymmetries are typically quantified using ratios of the highest (leading) jet p_T and second-highest (subleading) jet p_T in the event. A measurement of one such ratio, $x_J = p_T^{\text{subleading}} / p_T^{\text{leading}}$, is shown in Fig. 1-10 for both PbPb and pp collisions at 2.76 TeV. The PbPb values cluster at significantly lower x_J than the pp ones, indicating that one of the two jets in the dijet pair has been quenching more than its partner.

More recent measurements have explored the average energy loss and internal structure of quenched jets in much finer detail. They show that the energy lost in the quenching process tends to be redistributed to low energy particles at large distances from the jet [31, 32, 33, 34, 35]. These observations have been confirmed




Figure 1-10: A measurement by the ATLAS Collaboration of the p_T ratio of leading and subleading jet in 2.76 TeV PbPb and pp Collisions. The large difference between the two distributions is due to jet energy loss. Figure from Ref. [30].

using measurements of jets recoiling away from a high energy electroweak boson that gives a more accurate tag of the unquenched parton momentum [36, 37, 38, 39, 40, 41]. This enables a measurement of the absolute energy loss. However, these measurements tend to be limited by the statistical power of the data collected, and therefore much remains to be learned about the nature of parton energy loss and jet quenching.

As a control measurement, searches for jet quenching have also been performed in collisions of protons with heavy ions [42, 43, 44]. These studies have concluded that jet quenching does not seem to be a strong effect in these smaller systems, indicating that a QGP may not be present or is not large enough to produce observable effects. This remains an open question, however.

1.3.3 Hydrodynamic Flow

Studies of two-particle correlation functions in heavy ion collisions have revealed significant anisotropies in particle production [45, 46, 47, 48, 49, 50, 51]. Some anisotropies are expected because of conservation of momentum; many particles originating from parton showers will be close to each other in azimuthal angle (ϕ) and pseudorapidity (η) such that $(\Delta\eta, \Delta\phi) \approx (0,0)$. Also, particles coming from low- p_T




Figure 1-11: The near side ridge observed in PbPb (left), pPb (middle), and pp collisions (right) by the CMS and ATLAS Collaboration. Taken from Refs. [53, 51, 54].

resonance decays will tend to be back-to-back and have a difference of azimuthal angle $\Delta\phi = \pi$, while high- p_T decays will tend to contribute around $(0,0)$. However, a large anisotropy is also observed for particles having large $\Delta\eta$ but small $\Delta\phi$. The presence of this structure, which can be seen by the green ridge in the front half of Fig. 1-11, causes a fit to the data to have a large second order Fourier coefficient. The presence of this component is known as ‘elliptic flow.’ The observation elliptic flow at RHIC prompted the conclusion that the QGP flows like a liquid undergoing hydrodynamic expansion. The main idea is that the initial energy deposit from the collision is highly nonuniform and has significant spatial anisotropy that depends on the impact parameter. The QGP then responds to these anisotropies by trying to minimize pressure gradients, just like a fluid would. The result is that the initial spatial anisotropies are translated into momentum anisotropies in the final state. Thus, when measuring the two particle correlation function for particles at a given p_T , a large structure that is roughly independent of η can be found [52].

The presence of a elliptic flow is now known to be quite common in heavy ion collisions. Perhaps surprisingly, however, it has also been observed in collisions of protons and deuterons with heavy ions [55, 56, 57, 53], as well as high multiplicity pp collisions [54]. This has prompted an interpretation that a small droplet of QGP is formed in these types of collisions. Other theories, such as the presence of initial state correlations due to a color-glass condensate [58] have also been proposed to explain these correlations.

1.3.4 Thermal Photons

Because the quarks in the QGP are electrically charged, it can emit thermal photons as it expands. Because these photons have no color charge they can then pass through the QGP relatively undisturbed. Measurements by the PHENIX [59] and ALICE [60] Collaborations indicate a sizeable excess of photons not arising from hadron decays and having $p_T < 2$ GeV in head-on heavy ion collisions. An example of the ALICE measurement is shown in Fig. 1-12. The extra photons can clearly be seen by comparing the red points to the orange prediction. Such an excess cannot be described by perturbative QCD models and has been interpreted as thermal radiation from the QGP fireball. Using these data, models have been used to calculate an effective temperature of the QGP. The result is $T_{eff} = 221 \pm 19(\text{stat}) \pm 19(\text{syst})$ for the PHENIX data and $T_{eff} = 304 \pm 51(\text{stat} + \text{syst})$ for the ALICE data. The validity of these estimations is subject to concerns about the contribution of photons produced by charged particles after the evolution of the QGP fireball has ended [61]. However, if this estimate is correct, the QGP is around 300,000 times hotter than the center of the Sun and may be the hottest form of matter in the universe.

1.3.5 Strangeness Enhancement

At a temperature of a few hundred MeV, the average energy of the partons in the QGP is large enough create a significant number of strange quark anti-quark pairs via gluon splitting. These strange quarks then participate in the general hadronization of all the partons produced by the QGP. The result is that the final state hadrons tend to have a larger strange quark content as compared to proton proton collisions. This effect is particularly strong for hadrons that have more than one strange (anti-)quark, such as the ϕ meson, as well as the Ξ and Ω baryons. This phenomenon is known as *strangeness enhancement* [62]. It was first observed at the SPS [63, 64]. Later experiments at RHIC [65] and the LHC [66] have confirmed this effect. Figure 1-13 shows a compilation of measurements of the ratio of K_s^0 meson and Λ , Ξ and Ω baryon yields to that of π^\pm as a function of collision size. Larger values are observed




Figure 1-12: A measurement by the ALICE Collaboration of the yields of photons coming from 2.76 TeV PbPb collisions. Taken from Ref. [60].

for heavy ion collisions.

Recently, a weaker signature of this effect has also been observed in pPb [68] and high multiplicity pp [67] collisions. This could indicate that a small amount of QGP is created in these systems, or that there is also some other mechanism that contributes to an enhancement of strangeness production.

1.3.6 Quarkonia Suppression

Quarkonia are mesons comprised of a bound heavy quark anti-quark pair. When placed in a QGP, the strong interaction between the quark pair is reduced in strength

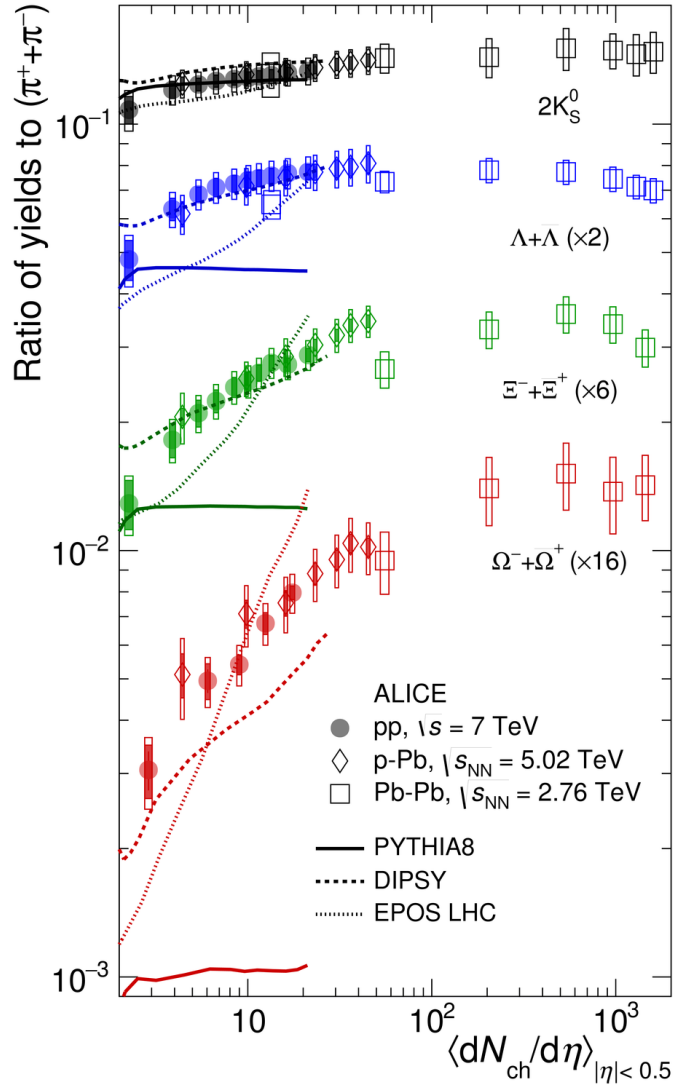


Figure 1-13: The ALICE Collaborations measurement of strange-hadron-to-pion ratios, showing the enhancement of strange baryon production as a function of collision size. Taken from Ref. [67].

by a process similar to Debye screening, where the QGP medium becomes polarized. This screening effect is expected to increase as the temperature of the QGP rises. If the effect is strong enough, the quarkonium may dissociate. Additionally, quarkonia states with a larger angular momentum have lower binding energy. Thus, the excited states are expected to melt before the ground state [69]. The suppression of the J/Ψ meson, composed out of a charm quark pair, has been well documented at the SPS [70, 71], RHIC [72, 73], and LHC [74, 75, 76]. Suppression of the $\Upsilon(1S)$ and its




Figure 1-14: A measurement by the CMS Collaboration of the relative yields of the $\Upsilon(1S)$ and $\Upsilon(2S)$ mesons in PbPb and pp collisions. Taken from Ref. [77].

excited state the $\Upsilon(2S)$, both of which contain a bottom quark pair, has also been observed [74, 77]. As can be seen in Fig. 1-14, the yield of $\Upsilon(2S)$ mesons relative to the ground state is smaller in PbPb collisions, indicating a sequential melting of these mesons.

1.3.7 Cold Nuclear Matter Effects

Not all effects observed in heavy ion collisions are caused by the presence of the QGP. They can instead be caused by the large number of nucleons present in the initial state. These effects are known as *cold nuclear matter effects*. It can be difficult to disentangle the contributions of these effects in collisions where a QGP is produced, because they become entwined with QGP signatures. Thus, these effects are usually studied by colliding a small particle, i.e. a proton or deuteron, with a large one such as gold or lead. In these collisions, a large QGP was not expected to be formed, such that observed phenomena could be attributed to cold nuclear matter effects. However, the universal validity of this approach has been called into question by the observation of QGP-like flow phenomena in small systems.

One of the most well-established cold nuclear matter effects is the modification of the nucleon PDF when it is in the presence of the other nucleons within a nucleus. This modification of the PDF causes a change in the kinematics of the hard scatterings produced in the collision. A special set of nuclear parton distribution functions (nPDFs) describe how large this effect is [78]. The nPDFs are usually calculated using a global fit of hard particle and/or jet production in pA or lepton-ion collisions. The exact details and data sets used in this fit procedure can vary, leading to multiple nPDF sets [79, 80, 81, 82]. An example of the modification in the nCTEQ15 nPDF set is shown in Fig 1-15 for gluons, as well as u , d , and s sea quarks. In general, the modification can be split into four kinematic regions which depend on the x of the parton and Q^2 of the process considered. At low values of x , a reduction in the nPDF is seen because nucleons at the surface of the nucleus can scatter the incoming parton flux, casting a ‘shadow’ on the nucleons at the center of the nucleus. This reduction is therefore known as nuclear *shadowing*. At slightly larger values of x , the nPDF tends to increase until it is larger than unity. This region compensates for the lost momentum fraction in the shadowing region and is known as *anti-shadowing*. At even higher values of x , the nPDF dips below unity again. This is known as the EMC effect. Finally, at the largest values of x it is possible for the nPDF to significantly exceed unity as a result of a parton inside a nucleon getting some additional momentum from the Fermi motion of the nucleon as it travels in the nucleus. These nPDF effects are most important when examining the cross sections and yields of hard processes in heavy ion collisions.

A second cold nuclear matter effect known as the Cronin effect has been observed in pA collisions [83, 84]. This effect describes a slight enhancement (on the order of 10% at the LHC and 40% at RHIC) in the production of high p_T hadrons due to an incoming parton scattering off multiple nucleons and therefore receiving a larger transfer of transverse momentum [85]. This effect peaks at around $p_T = 3$ GeV at RHIC and the LHC and is larger for baryons than for mesons. This effect is distinct from nPDF anti-shadowing, because the nPDF describes the probability of a parton participating in a scattering, while the Cronin effect describes a modification to the




Figure 1-15: The relative modification of the free proton PDF in the lead ion at $Q = 10$ GeV, as determined in nCTEQ15 [80]. The result is shown for gluons, and u , d , and s sea quarks.

actual scattering amplitude itself.

1.4 Charged Particle Nuclear Modification Factors

One quantitative method of measuring the effects of the presence of a QGP and/or cold nuclear matter effects is to compare the yield or spectrum of a given particle to a reference of proton proton collisions. The proton collisions are assumed to have no nuclear effects. If one assumes a heavy ion collision is simply a superposition of multiple proton-proton collisions without any other effects, then scaling a proton-proton spectrum by the appropriate number of nucleon-nucleon collisions (N_{coll}) should give the same result as a heavy ion collision. (It should be noted that this assumption breaks down for low- Q^2 processes where a parton can interact coherently with the entire nucleus rather than a single parton inside one of the nucleons.) This comparison

can be quantified by taking a ratio of the heavy ion and the scaled pp spectra

$$R_{AA}(p_T) = \frac{1}{\langle N_{\text{coll}} \rangle} \frac{dN^{AA}/dp_T}{dN^{pp}/dp_T}. \quad (1.7)$$

The quantity R_{AA} is known as the nuclear modification factor. Although the assumption that the heavy ion spectrum scales with N_{coll} is not valid at low p_T , measurements of electroweak boson production in heavy ion collisions indicate it is reasonable for hard processes [86, 87, 88, 89, 90, 91, 92]. The exact values of the scaling factor N_{coll} can be calculated using a Glauber model of heavy ion collision, which is described in Section 4.2. An equivalent definition of the nuclear modification factor replaces the proton proton per-event yield with a differential cross section, and N_{coll} with T_{AA} , where

$$T_{AA} = N_{\text{coll}}/\sigma_{NN}^{inel}. \quad (1.8)$$

Here σ_{NN}^{inel} is the total inelastic proton proton cross section. This substitution is done because T_{AA} can usually be calculated more accurately than N_{coll} . With this replacement, the equivalent definition of R_{AA} is

$$R_{AA}(p_T) = \frac{1}{T_{AA}} \frac{dN^{AA}/dp_T}{d\sigma^{pp}/dp_T}. \quad (1.9)$$

The nuclear modification of charged particles can be measured accurately to fairly high p_T in most experiments. Charged particles are produced in abundance by most QCD processes at hadron colliders and are much more common than the production of heavy objects like the electroweak bosons. Furthermore, high p_T charged particles come almost exclusively from hard scatterings and the resulting parton shower. This means that the high p_T charged particle nuclear modification factor is very sensitive to parton energy loss and jet quenching. At lower p_T , the strength of phenomena affecting hadrons, such as strangeness enhancement and hydrodynamic flow, are also probed by the nuclear modification factor.

In this thesis, the nuclear modification factors of charged particles in PbPb collisions are measured at a ($\sqrt{s_{NN}}$) of 5.02 TeV. The same quantity is measured in XeXe

collisions at a center of mass energy of 5.44 TeV, and the results are compared. To have a reference that takes into account the potential result of cold nuclear matter effects not related to QGP production, the nuclear modification factor in pPb collisions at 5.02 TeV is also measured. The measurements are conducted using the CMS experiment at the LHC.

The following two sections describe the theoretical and experimental understanding of the charged particle nuclear modification factor roughly as of late 2014, when the analyses described in this thesis began.

1.4.1 Theoretical Models

A full description of the QGP is a challenging topic theoretically because the dynamics of a strongly-coupled QCD medium involve low- Q^2 processes and are very nonperturbative in nature. However, significant progress has been made trying to connect the dynamics of perturbative QCD found in parton showers with the soft interactions of the QGP [93]. Measurements of the charged particle nuclear modification factor at high p_T are able to constrain these calculations of in-medium parton energy loss. To facilitate comparison of data to theory, some of the different models used to describe parton energy loss are described below. Some methods are very theoretical, while others are more phenomenological in nature.

AdS/CFT Correspondence

The Hybrid model [94] uses the AdS/CFT correspondence [95] to map the problem of a jet evolving in a strongly-coupled medium into one of a string falling into a 5 dimensional black hole. This gravitational dual can be solved to give an expression for the energy loss per unit length (dE/dx) of a parton:

$$\frac{1}{E_{in}} \frac{dE}{dx} = -\frac{4}{\pi} \frac{x^2}{x_{stop}^2} \frac{1}{\sqrt{x_{stop}^2 - x^2}}. \quad (1.10)$$

Here E_{in} is the initial parton energy, and x_{stop} is given by

$$x_{stop} = \frac{1}{2\kappa_{SC}} \frac{E_{in}^{1/3}}{T^{4/3}}. \quad (1.11)$$

The parameter κ_{SC} is the only free parameter of the model, and is determined by fitting measurements of the jet R_{AA} .

This energy loss expression is implemented by taking an existing simulated parton shower (calculated using perturbative QCD methods) and letting each parton lose energy between splittings. The new set of partons are then run through the hadronization procedure again to give the final output. As its name implies, the Hybrid model combines a perturbative QCD shower and a strongly-coupled calculation in a way that attempts to take advantage of each method's strengths. This method has been proven to be able to describe many features of heavy ion data.

BDMPS

The Baier-Dokshitzer-Mueller-Peigne-Schiff model of parton energy loss is focused mainly on describing the spectrum of radiated gluons in a medium-modified parton shower [96, 97]. The medium is modeled as a set of colored static scattering centers. Because these scattering centers are static, energy cannot be transferred to the medium via collisional processes; these centers only serve to induce gluon radiation. Most applications of the BDMPS formalism assume that the parton interacts via a series of multiple soft radiations. In this class of models, the medium can be fully characterized by a *jet transport coefficient* $\hat{q} = d\langle k_T^2 \rangle / dL$ that describes the average squared transverse momentum (k_T) that a parton exchanges with the medium per unit path length. This model can be related to the charged particle nuclear modification factor by [98]

$$R_{AA}(p_T) = \int_0^\infty d\epsilon \left(1 + \frac{\epsilon}{p_T} \right)^{-n} P(\epsilon). \quad (1.12)$$

In this equation ϵ is a quantity of energy lost by the parton while traversing the medium, n is the spectral index of the hadron spectrum in pp collisions, and $P(\epsilon)$

is a probability distribution describing the chance to lose energy ϵ . This distribution is known as a quenching weight and can be calculated from the BDMPS gluon spectrum [99].

GLV

The Gyulassy-Levai-Vitev (GLV) formalism [100] introduces the concept of an opacity expansion, in which radiative effects are expanded in terms of the number of scatterings a parton experiences. Typically terms with only one or a few scatterings are included. This formalism naturally takes into account interference between vacuum and medium-induced radiations, unlike the BDMPS formalism. The original implementation of GLV also modeled the medium as a series of heavy static scattering centers having a density ρ . It was later expanded to include dynamic scattering centers by Djordjevic and Heinz [101]. This development, along with calculations by Djordjevic of the collisional energy loss using a similar opacity expansion procedure [102], allow for the creation of models for LHC data [103, 104].

The CUJET model [105] also uses a dynamic GLV energy loss kernel, with a different formulation for collisional energy loss [106]. It also modifies the energy loss when the QGP is near the critical transition temperature in order to incorporate results from lattice QCD calculations, such as the formation of quasi-particle degrees of freedom.

Soft Collinear Effective Theory

Both the BDMP and GLV frameworks assume that hard splitting processes are somewhat factorized from the induced radiations coming from the presence of the medium, and therefore the hard parton never changes its identity. However, it is possible that hard scatterings could cause the transfer a large part of its momentum to another parton that also undergoes energy loss. Supporting such hard scatterings motivates the use of a different theoretical framework.

Soft Collinear Effective Theory (SCET) attempts to circumvent the divergences of QCD at low energy scales by factorizing soft and collinear gluon modes in the

evolution of the parton shower [107]. This is done by defining collinear, ‘soft,’ and ‘ultrasoft’ gluon fields that describe dynamics at different momentum scales. To include the effects of a QGP medium, an additional gluon mode describing Glauber gluons is added. This model is known as SCET_G [108]. Glauber gluons are gluons having momenta transverse to the direction of the hard parton; these are the gluons that are most likely to result in significant broadening of the parton momentum. Within this framework, it is possible to define in-medium fragmentation functions that evolve in a manner similar to DGLAP evolution. These fragmentation functions can support the presence of hard scatterings in the parton shower, and can be used for estimating the charged particle R_{AA} [109].

Linear Boltzmann Transport

In many of the simpler parton energy loss models, the fate of energy radiated from the leading parton is not closely tabulated. The Linear Boltzmann Transport (LBT) model uses a linearized Boltzmann equation to keep track of the energy flow within a parton shower and its resulting medium excitations. This equation describes the evolution of parton phase-space distributions f_i , which are a function of the parton four-momenta p_i [93]. It can be written as

$$p_1 \cdot \partial f_1(p_1) = - \int \frac{d^3 p_2}{2E_2(2\pi)^3} \frac{d^3 p_3}{2E_3(2\pi)^3} \frac{d^3 p_4}{2E_4(2\pi)^3} (f_1 f_2 - f_3 f_4) |M_{12 \rightarrow 34}|^2 \times (2\pi)^4 \delta^4(p_1 + p_2 - p_3 - p_4). \quad (1.13)$$

For thermal partons the phase space distributions are given by quantum statistics as $f_i = 1/(e^{(p \cdot u)/T} \pm 1)$, while for shower partons $f_i = (2\pi)^3 \delta^3(\vec{p} - \vec{p}_i) \delta^3(\vec{x} - \vec{x}_i - \vec{v}_i t)$. The $|M_{12 \rightarrow 34}|$ term describes two-to-two elastic scattering amplitudes from perturbative QCD calculations. The final delta function enforces conservation of energy-momentum in all the elastic scatterings considered.

Using this equation, the parton shower is allowed to evolve over a series of time steps, and the effects of elastic collisional energy loss processes are tabulated over time.

Induced radiation effects are simulated by allowing gluon radiation between time steps according to a higher-twist formalism [110]. In the so-called twist-expansion, interaction amplitudes are expanded in terms of λ/Q , where λ is some relatively small energy scale and Q is the momentum transfer of the process [111]. For modeling the QGP medium, a combination of single-parton-exchange ‘twist-2’ matrix elements, as well as those that correspond to the exchange of multiple partons (‘higher-twist’ contributions), are used. This method is somewhat similar to the GLV approach because most numeric implementations only consider a single scattering center per emitted gluon (although this scattering can involve the exchange of multiple partons). Unlike the GLV method, it only uses the first moment of the scattering Q^2 distribution so that the medium can be described by a \hat{q} parameter. Because the LBT model keeps track of the evolving energies of the partons in a shower, it is able to account for the changing of the strong coupling and back reaction effects from excitations of the medium [112].

QGP-PDF

The formulation of Bianchi et. al. uses a rather different approach to parton energy loss [113]. In this model a PDF for the QGP is defined, where the momentum fraction variable x for a parton in the QGP is given by

$$x = \hat{q}L/(2M_N E). \tag{1.14}$$

In this formulation, M_N is the mass of a nucleon, E is the jet energy, \hat{q} is the jet transport coefficient, and L is the path length in the QGP. With this definition, the QGP-PDF is parameterized as

$$G(x, \mu^2 = 1 \text{ GeV}^2) = Nx^\alpha(1-x)^\beta \tag{1.15}$$

where α and β are parameters that are fit to data. Both DGLAP and BFKL (which evolves PDFs in x) evolution can then be used to calculate the dependence of the jet transport coefficient, \hat{q} , on Q^2 . Once this dependence is known, it can be combined

with a hydrodynamic simulation of the medium to allow the calculation of event-by-event energy loss and R_{AA} [114]. In this model it is found that \hat{q} varies sizeably when comparing RHIC and LHC energies.

BBMG

The BBMG energy loss formalism of Betz, Gyulassy, and Torrieri [115, 116] describes a broad class of parton energy loss models that have a parton energy loss per unit length, dE/dx , that can be described as

$$\frac{dE}{dx} = -\kappa E^a(x) x^z T^{2+z-a} \xi_q \Gamma_{flow}. \quad (1.16)$$

Here $E(x)$ is the jet energy, x is the distance the parton has traveled in the medium, T is the local medium temperature, and ξ_q describes local energy loss fluctuations. The parameter κ is a dimensionless constant that is fit to data, usually a measurement of a nuclear modification factor. Γ_{flow} is a factor that varies with the difference between the jet direction and the direction of the local hydrodynamic flow field of the medium. The parameters a and z can be chosen to give various scenarios, but a “pQCD scenario” ($a = 0$, $z = 1$) is used in the models discussed here. Thus, this model gives a roughly quadratic path-length dependence (from integrating both sides of Equation 1.16 in x). Physics observables can be calculated by simulating the temperature and hydrodynamic evolution of the medium, and then evaluating the total energy lost by a jet traversing the medium with Equation 1.16. This model is rather phenomenological in that it does not attempt to model the underlying QCD processes that cause parton energy loss. It instead focuses on how the evolution of the medium itself can affect a parton shower.

Quenching Weights

A phenomenological description of jet quenching can be achieved by defining the K factor of the medium created in a particular type of collision as

$$K = \frac{\hat{q}}{2\epsilon^{3/4}}. \quad (1.17)$$

The quantity ϵ is the QGP energy density, and $2\epsilon^{3/4}$ is considered to be the ideal value of \hat{q} . Thus, K quantifies a given medium's deviation from the ideal transport coefficient. Using a quenching weight $P(2K\epsilon^{3/4})$ [97], the effective jet transport coefficient can be applied to partons traversing a hydrodynamically evolving medium to calculate the rate of energy loss. In principle, the quantity K is fit to data. The model of Andrés et al. uses charged particle nuclear modification factor data at RHIC and the LHC to try to constrain K and make predictions at higher center of mass energies [117].

1.4.2 Experimental Status

Experimentally, the nuclear modification factor for light hadrons is a touchstone measurement that has been performed many times for a variety of collision systems and energies. Here these measurements are broken up into two categories. The first contains pA and dA collisions, which attempt to probe potential cold nuclear modification effects, and the second is for AA collisions, which seek to quantify QGP effects. A summary of historical measurements of nuclear modification factors for light hadron species is shown in Table 1.1.

pA and dA collisions

High energy hadronic collisions involving heavy nuclei were first conducted by using proton beams incident on a fixed target. Many of these experiments focused on measuring the spectral ratios of various particle species, and few compared directly to a pp reference to form what we now refer to as a nuclear modification factor. One of the

first intriguing comparisons of heavy ion collisions to a pp reference was done in 1975 by Cronin et. al. [83] using the 300 GeV proton beam at the National Accelerator Lab (NAL) with tungsten, beryllium, and titanium targets. Particle production was measured at high transverse momentum (up to 9 GeV), and compared to pp spectra measured at the CERN Intersecting Storage Rings (ISR). It was observed that the slopes of the spectra were fairly similar at high p_T after accounting for the additional number of nucleons in the target nucleus. This experiment also established the presence of the ‘Cronin’ effect (discussed in Section 1.3.7) for collisions having $\sqrt{s_{NN}} \approx 20$ GeV, where the yields of baryons seem to be enhanced relative to those of mesons.

In the early 1980s, nuclear deep inelastic scattering experiments compared the structure functions of nucleons within different sized nuclei and found intriguing differences, such as the EMC effect [147, 148], which are now taken as part of the nuclear parton distribution function. Thus, the existence of two significant cold nuclear matter effects was known before any the first studies of QGP production in nuclear collisions.

Once RHIC came online around 2000, these cold nuclear matter effects needed to be understood in detail in order to confirm that a QGP was indeed being created. To this end, a run of dAu collisions and pp collisions at $\sqrt{s_{NN}} = 200$ GeV was performed. All four experiments at RHIC measured the charged hadron nuclear modification factor in these collisions, and obtained results similar to the BRAHMS data shown by the green points in Fig 1-16. The excess of charged particles around $p_T = 3-5$ GeV indicated a large amount of Cronin enhancement in these collisions. However, the trend of the data seems to indicate that this enhancement gets smaller as the p_T increases past 4 GeV. Similar measurements were also conducted by the PHENIX Collaboration for π^0 and η mesons, indicating that there does not seem to be a strong high p_T suppression of any light hadron production relative to pp collisions.

In 2013, a run of pPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV was conducted at the LHC. The large increase in collision energy, along with the presence of detectors having full calorimeter coverage, make these collisions a unique opportunity to examine the behavior of the pA nuclear modification factor at very high p_T . In particular, scattered




Figure 1-16: An example of the $\sqrt{s_{NN}} = 200$ GeV measurements of the nuclear modification factor in AuAu and dAu collisions at RHIC. Figure taken from the BRAHMS Collaboration [149].

partons having an energy of a few hundred GeV are expected to be produced from the anti-shadowing region of the nuclear PDF. Measurements of the *jet* nuclear modification factor by the CMS and ATLAS Collaborations, shown in Fig. 1-17, confirmed that a slight anti-shadowing of around $(10 \pm 10)\%$ is indeed present in these collisions at midrapidity [44, 43]. Later measurements of the dijet pseudorapidity distribution by CMS, displayed in Fig. 1-18, also showed this effect with a higher precision [150]. Because of the subsequent parton fragmentation and hadronization, it was expected that this anti-shadowing enhancement should also show up in a measurement of the charged particle nuclear modification factor in the p_T range of 10–100 GeV.

The CMS Collaboration’s measurement of the charged particle nuclear modification factor in 5.02 TeV pPb collisions [126] is shown in Fig. 1-19. A similar measurement by the ALICE Collaboration is also displayed [123] in blue. The ATLAS Collaboration also showed a *preliminary* result that was very similar to the CMS data points, but the final result has the data points above 20 GeV removed [125]. The data under $p_T = 5$ GeV is largely consistent with what was observed at RHIC, despite the large difference in center of mass energy. However, starting at around 20 GeV a




Figure 1-17: The nuclear modification factor of jets in pPb collisions, as measured with an extrapolated pp reference [44]. A measurement by ATLAS is also shown [43].

large bump structure is observed in the CMS data. Although some anti-shadowing enhancement was expected in this kinematic range, the observed enhancement goes up to 1.4. Theoretical models, such as the one shown by the red line which uses the EPS09 nPDF [151], only predicted a 10% enhancement. As previously noted, the jet nuclear modification factor indicated an excess of around $(10 \pm 10)\%$. Thus, it appeared that nPDF effects could not explain this large bump.

This puzzling result seemed to imply one of two things. Firstly, the fragmentation function which connects the nPDF-modified hard parton scattering with the final state hadron spectrum (via Eqn. 1.6) could itself be modified by some previously-unknown cold nuclear matter effect. If the pPb fragmentation function somehow caused jets to fragment into hard hadrons 0% more frequently than in pp collisions, then the discrepancy between the jet and charged particle nuclear modification factors would be fully explained. This possibility is explored further in this thesis by measuring the pPb jet fragmentation function and comparing it to the same quantity in pp collisions. The other potential explanation deals with how the denominator in the nuclear modification factor ratio is calculated. At the time of these measurements, no 5.02 TeV pp reference data existed, so the reference quantities were constructed by extrapolating or interpolating existing pp measurements at 2.76 and 7 TeV. The




Figure 1-18: Comparisons by the CMS Collaboration of the dijet pseudorapidity distribution in 5.02 TeV pPb and pp collisions. In the mid-rapidity region a strong anti-shadowing effect is seen for all p_T bins examined. Figure from Ref. [150].

jet nuclear modification factor used a extrapolation procedure based on simulation, while the charged particle measurement used a data-driven interpolation method. It is possible that one or both of these methods deviate from the true value of the pp spectrum at 5.02 TeV, leading to the apparent discrepancy. In 2015 a dedicated 5.02 TeV pp data set was collected to resolve this issue by making direct measurements of the required reference quantities. These data are also examined here and the conclusions of these measurements are revisited.

AA collisions

Early measurements of nuclear collisions at the ISR [127] and SPS [128, 129, 130, 131] were limited to center of mass energies of around 20 GeV and could only produce particles having $p_T \approx 5$ GeV. Nonetheless, significant evidence including strangeness enhancement and quarkonia suppression pointed towards the production of a QGP in these collisions. The nuclear modification factors of charged and neutral pions were also measured in these collisions. A summary of the various R_{AA} can be see in Fig. 1-20. In general a suppression is observed at low p_T where particle yields are expected to




Figure 1-19: The nuclear modification factor for charged particles in pPb collisions, as measured with an interpolated pp reference [126]. A model prediction [151] and a result from ALICE [123] are also shown.

scale with the number of participating nucleons instead of the total number of binary collisions. A rising trend then approaches unity in the region of 3–4 GeV. This line shape was thought to be produced by a significant amount of Cronin enhancement, but has also been interpreted as originating from a combination of anti-shadowing and moderate energy loss effects [152]. The interpretation of these results was also plagued by differing parameterizations of the chosen pp reference. Thus, it remained uncertain if jet quenching was a large effect at the SPS.

At RHIC, the nuclear modification factor was measured for 200 GeV AuAu collisions by four experimental collaborations in a variety of channels including charged hadrons, π^0 , and η mesons [118, 132, 133, 134, 120, 135, 136, 137]. All four collaborations found a large *suppression* of R_{AA} at p_T around 7 GeV. An example of the characteristic line shape at RHIC is shown by the red points in Fig. 1-16. In general, a suppression and rising trend is seen at low p_T , similar to what is seen in dAu and lower energy results. However, a peak forms around 2–3 GeV and the R_{AA} begins to decrease again. This additional suppression at high p_T indicates the presence of parton energy loss effects in this kinematic regime. The strength of the suppression is expected to be related to the strength of parton energy loss. To probe




Figure 1-20: A summary of some of the R_{AA} measurements in the pre-RHIC era. Summary figure taken from [153].

this further, similar experiments were undertaken using smaller CuCu collisions at 200 GeV [138, 139, 140, 141]. It was expected that the different collision geometry would produce a different suppression pattern in the nuclear modification factor, allowing the study of how parton energy loss changes with the size of the collision system. In general the CuCu data was found to be suppressed less than the AuAu data in the high p_T region, indicating that the smaller QGP produced in CuCu collisions is able to quench partons less. However, these measurements at RHIC were fairly limited in the maximum p_T probed; only one measurement managed to examine the p_T range above 10 GeV.

Recent measurements at the LHC have managed to extend the p_T range of nuclear modification measurements by nearly an order of magnitude [142, 143, 144, 145, 146]. Figure 1-21 shows measurements done by the ALICE, ATLAS, and CMS Collaborations for PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. Note that the plot is on a logarithmic scale, so that the oscillatory shape is visually enhanced. A similar shape is seen as compared to the RHIC data under $p_T = 10$ GeV, but now a rising trend above 10 GeV can now be observed up to around 100 GeV. This rising trend indicates that the parton energy loss effects become less strong at higher parton energies, indicating that the QGP becomes more transparent to hard partons. Interestingly, the CMS data




Figure 1-21: A compilation of the LHC charged particle R_{AA} measurements made at $\sqrt{s_{NN}} = 2.76$ TeV. Figure taken from Ref. [145].

seem to indicate that the R_{AA} might plateau at a value of 0.5 around $p_T = 80$ GeV, while the ATLAS data indicate that it continues to rise at even larger p_T . Whether the QGP continues to become transparent at very large p_T values remains an open question.

Figure 1-22 is a summary of the nuclear modification factor for π^0 mesons across a large range of center of mass energies, from the SPS to the LHC. A trend of the data becoming more suppressed as the energy increases can be observed, particularly at large p_T . Interestingly, the 200 GeV RHIC and 2.76 TeV LHC data are very close to each other around $p_T = 8$ GeV. The similarity of the suppression is understood as an effect of the hardening of the particle spectra at higher collision energies. As the slopes of the spectra become smaller, the nuclear modification factor will be less suppressed assuming a similar amount of energy loss. Thus, the total suppression of the nuclear modification factor depends on both the strength of energy loss and the center of mass energy. It seems that these effects roughly cancel out to make the suppression observed in the RHIC and LHC data similar. It remains to be seen how these two effects interact at even higher collision energies.




Figure 1-22: A summary of how the $\pi^0 R_{AA}$ evolves as a function of center of mass energy, from SPS to LHC energies. Figure taken from Ref. [143].

1.4.3 Questions Addressed in this Work

During the LHC Run 2 period between 2015 and 2018, unprecedentedly large data sets of PbPb and XeXe collisions were collected. These collisions were gathered at a new heavy-ion record center of mass energy of 5.02 TeV (PbPb) and 5.44 TeV (XeXe). Furthermore, these data were accompanied with a dedicated pp reference run at 5.02 TeV, allowing the calculation of the charged particle nuclear modification factors for these systems. These reference data also allow the reexamination of 5.02 TeV pPb collisions that were gathered in 2013. Due to synergies with LHC's high-energy physics program, the CMS detector is well-suited for these measurements because of its excellent charged-particle reconstruction capabilities, fast data acquisition, and full calorimetric coverage. This thesis focuses on charged particle measurements using the 2013 pPb data, the 2015 PbPb and pp data, and the 2017 XeXe data.

The next chapters focus on the experimental tools and techniques used to make charged particle measurements with the CMS detector. Particular detail will be spent on the tools used for the reconstruction of charged particles, as they are an essential component of the subsequent measurements.

Next, the apparent discrepancy between existing pPb nuclear modification factor

measurements is tested with a measurement of the jet fragmentation function in pPb collisions. Using previously gathered pp data at collision energies of 2.76 and 7 TeV, a pp reference fragmentation function at 5.02 TeV is constructed for comparison. A strongly modified fragmentation function would imply some new physics mechanism in pPb collisions that has not been yet predicted, while an unmodified one would indicate some issue with existing measurements. The dependence of these measurements on a interpolated or extrapolated pp reference is also scrutinized by updating the pPb charged particle nuclear modification factor to use measured pp data. With these measurements the importance of cold nuclear matter effects on charged particle nuclear modification factors are explored.

Next, the charged particle R_{AA} is measured in 5.02 TeV PbPb collisions. With these data the dependence of parton energy loss on $\sqrt{s_{NN}}$ can be explored by comparing to previous results at lower energies. Additionally, the size of the PbPb data set allows probing R_{AA} at very large values of p_T , allowing clarification on the slope of R_{AA} at $p_T > 80$ GeV. Such a measurement is directly linked to the opacity of the QGP to very high-momentum partons in the medium.

Finally, the charged particle R_{AA} is also measured for 5.44 TeV XeXe collisions. Such a measurement is important to provide data at an ‘intermediate’ collision system size at LHC energies. Xenon ions are smaller than lead ions, so less parton energy loss is expected in head-on collisions because a smaller QGP should be formed. Detailed comparisons between the XeXe and PbPb nuclear modification factors, after accounting for the difference in $\sqrt{s_{NN}}$, provide information about the path length dependence of energy loss in a QGP medium. The XeXe and PbPb results are also compared to theoretical models with the intent of constraining calculations of parton energy loss.

Table 1.1: Previous measurements of the nuclear modification factor for light particle species. The first grouping consists of pA and dA collisions, while the second grouping is AA collisions.

| Accelerator | Experiment | System | Final state | $\sqrt{s_{NN}}$ (GeV) | max p_T (GeV) | Reference |
|-------------|----------------------|----------------|----------------------|-----------------------|-----------------|---------------------------|
| NAL | Cronin et. al. | pW, pBe, pTi | π^\pm | 18.4, 23.8, 27.4 | 9 | [83] |
| RHIC | BRAHMS | dAu | h^\pm | 200 | 5.5 | [118] |
| RHIC | PHENIX | dAu | h^\pm, π^0, η | 200 | 6 | [119, 120] |
| RHIC | PHOBOS | dAu | h^\pm | 200 | 6 | [121] |
| RHIC | STAR | dAu | h^\pm | 200 | 10 | [122] |
| LHC | ALICE | pPb | h^\pm, π^0, η | 5020 | 50 | [123, 124] |
| LHC | ATLAS | pPb | h^\pm | 5020 | 22 | [125] |
| LHC | CMS | pPb | h^\pm | 5020 | 120 | [126] |
| ISR | BCMIOR Collaboration | $\alpha\alpha$ | π^0 | 21 | 9 | [127] |
| SPS | NA49 | PbPb | π^\pm | 17.3 | 2 | [128] |
| SPS | WA98 | PbPb | π^0 | 17.3 | 2.5 | [129] |
| SPS | CERES/NA45 | PbAu | π^\pm | 17.3 | 3.5 | [130] |
| SPS | WA80 | SS, SAu | π^0 | 19.4 | 3.5 | [131] |
| RHIC | BRAHMS | AuAu | h^\pm | 200 | 5 | [118] |
| RHIC | PHENIX | AuAu | h^\pm, π^0, η | 39, 62.4, 200 | 20 | [132, 133, 134, 120, 135] |
| RHIC | PHOBOS | AuAu | h^\pm | 62.4, 200 | 4.5 | [136] |
| RHIC | STAR | AuAu | h^\pm | 200 | 10 | [137] |
| RHIC | BRAHMS | CuCu | π^\pm | 200 | 2 | [138] |
| RHIC | PHENIX | CuCu | π^0 | 22.4, 62.4, 200 | 17 | [139] |
| RHIC | PHOBOS | CuCu | h^\pm | 62.4, 200 | 7 | [140] |
| RHIC | STAR | CuCu | π^\pm | 200 | 8 | [141] |
| LHC | ALICE | PbPb | h^\pm, π^0, η | 2760 | 50 | [142, 143, 144] |
| LHC | ATLAS | PbPb | h^\pm | 2760 | 120 | [145] |
| LHC | CMS | PbPb | h^\pm | 2760 | 100 | [146] |

Chapter 2

The CMS Detector

The Compact Muon Solenoid (CMS) experiment is a versatile particle detector situated 100 meters underground near the town of Cessy, France. CMS examines collisions of protons, xenon, or lead nuclei delivered by the Large Hadron Collider (LHC), the world's most energetic particle collider [154]. Designed as a general purpose detector, CMS is comprised of a host of subdetectors assembled in a nearly hermetic configuration to capture as many particles emerging from the LHC interaction point as possible. These subdetectors are arranged in concentric layers to allow precise measurements of particle properties, including their momentum and charge, or energy.

CMS derives its name from the superconducting solenoid that maintains a uniform 3.8 Tesla magnetic field through the core of CMS, parallel to the LHC beam pipe. This solenoid is very large, having an inner bore of 5.9 meters and length of 12.9 meters, in order to accommodate four major subdetectors. Lying closest to the magnet is the hadronic calorimeter (HCAL), used to measure the energy of particles such as neutrons that interact mainly hadronically. Inside of the HCAL is the electromagnetic calorimeter (ECAL), which measures the energy of electromagnetically interacting particles like electrons and photons. The strip tracker is nested inside the ECAL, and is used to reconstruct the momentum and trajectory of charged particles as they propagate out from the interaction point towards the calorimeters. At the very center of CMS lies a pixel tracker, which uses its fine granularity to precisely identify the origin of the charged particles which are then subsequently tracked by the strip




Figure 2-1: A cutaway view of CMS, showing the positions of the magnet and different subdetectors [155].

tracker.

Other subdetectors are present outside of the solenoid. In the iron magnetic return yoke lying around the solenoid are many muon stations, used to identify high energy muons which manage to pass through all the previously mentioned subdetectors. Two forward calorimeters, CASTOR and the Zero Degree Calorimeter, lie very close to the beam line, but far downstream from the interaction point. A third forward calorimeter, the hadronic forward calorimeter (HF), is very close to the magnetic return yoke and covers the pseudorapidity gap between the beam pipe and the start of the HCAL acceptance. A cut-away view of the CMS detector can be seen in Fig. 2-1.

The coordinate system used in CMS is arranged with the x -axis pointed towards the center of the LHC. The y -axis is oriented upward in the vertical direction, and the z -axis is defined such that the coordinate system is right handed. Azimuthal

angles, denoted with the variable ϕ , are measured starting from the x -axis in the xy plane and radial distances, r , are also measured in this plane. Polar angles, θ , are defined in the rz plane, such that a particle's pseudorapidity is $\eta = -\ln(\tan(\theta/2))$. Components of a particle's momentum transverse to the z direction are denoted as p_T .

A complete description of CMS can be found in Ref. [155].

2.1 Pixel Tracker

The CMS pixel tracker is a silicon subdetector designed to sit close to the CMS interaction point. The subdetector covers the pseudorapidity region of $|\eta| < 2.4$. It consists of three 53 cm long cylindrical layers which lie 4.4, 7.4, and 10.2 cm away from the interaction point, as well as four forward disks. These disks are placed at $z = \pm 34.5$ and $z = \pm 46.5$ cm and have inner and outer radii of 6 and 15 cm. This arrangement, shown by the red lines in Fig. 2-2, was chosen to ensure that at least three measurements are made along a charged particle's trajectory within nearly the entire pseudorapidity range instrumented.

The focus of the subdetector is to provide high precision three-dimensional measurements in the r , ϕ , and z coordinates. Towards this end, the various layers are instrumented with arrays of silicon pixel detectors having a pixel size of $100 \times 150 \mu\text{m}^2$. This leads to an average pixel occupancy of 10^{-4} per bunch crossing during nominal pp conditions at the LHC. A total of 64 million pixels are used to cover the 1.06 m^2 area of the subdetector. This excellent granularity allows the detector to still function adequately in very high multiplicity environments, such as those seen in head-on PbPb collisions.

At midrapidity each pixel sensors' electric field, which is responsible for collecting the liberated charge produced by a passing particle, is at approximately right angles to the CMS magnetic field. This causes the liberated charges to drift sideways and be deposited in more than one pixel. The Lorentz angle characterizing this drift has been measured to be approximately 22° [156]. The Lorentz drift, combined with

measurements of the amount of collected charge in neighboring pixels, is used to improve the spatial resolution of pixel hits from approximately $30\mu\text{m}$ (the inherent resolution due to the physical size of the pixels) to values between 15 and $20\mu\text{m}$. The Lorentz drift in the forward disks is much smaller than in the barrel due to the electric field being approximately parallel to the CMS magnetic field. To compensate for this, the pixel modules on the forward disks are rotated by 20° in order to enhance charge sharing due to charged particles entering the pixel modules at an angle.

Extracting hit information from the subdetector requires a multi-level electronics readout scheme. At the lowest level, groups of 52×80 pixels are attached to Read-Out Chips (ROCs) for data acquisition. Each ROC amplifies the charge collected by each pixel and applies a zero-suppression to pixels by requiring at least 90 electrons collected by the pixel on average (the exact threshold can be configured for each individual pixel). Only pixels passing this requirement have their analog signal forwarded to the next step in the readout chain: the Token Bit Manager (TBM) chips. These chips query the ROCs when a trigger is received, and assemble all hit information from the ROCs for a given event. Each TBM is responsible for either 8 or 16 (21 or 24) ROCs in the barrel (forward) layers. The signals are then forwarded to one of 40 Front End Drivers (FEDs). The FEDs are groups of 36 optical inputs, each taking a signal from a TBM, which are connected to Analog-to-Digital Converters (ADCs). From this point forward, the data stream is fully digital. The FEDs also assemble error information transferred from the ROCs, and are able to flag potential problems with the readout chain or defective sensors. Typically, these errors will result in sensors being purposefully ignored for offline analysis.

The three precise measurements provided by the pixel detector allow accurate reconstruction of event vertices. This is important when taking high-luminosity data, which often has non-negligible pileup conditions resulting in multiple primary interactions per bunch crossing. Without precise vertexing these interactions could look like a single, more active collision. Furthermore, the pixel detector allows the reconstruction of secondary vertices, allowing the study of short-lived intermediary particles like B and D mesons. In the CMS tracking algorithms, groupings of three or more pixel

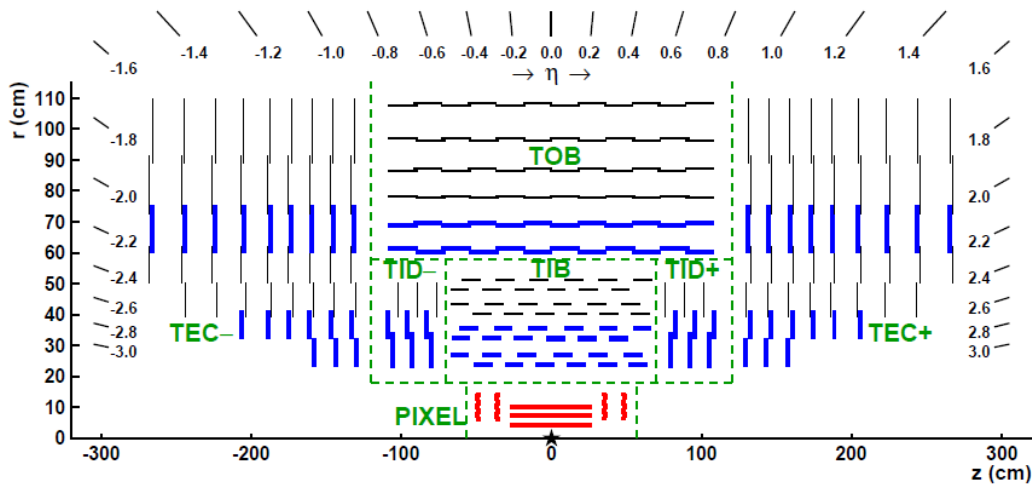


Figure 2-2: A schematic of the positions of the CMS tracker layers and disks in half of the tracker. The thick blue lines indicate sensors which give "stereo" hit information. The red lines at the bottom are the pixel detector [157].

hits that are compatible with a helix trajectory are used to provide an initial estimate of a particle's path, and can even be used as stand-alone 'pixel tracks' at low track p_T .

Between 2016 and 2017, the pixel detector was upgraded to its "phase one" configuration by replacing it with a new detector having a fourth pixel barrel layer that lies within 3 cm of the beam pipe. Another set of two endcap disks was also added. A comparison of the original pixel detector, along with its phase one upgrade, is shown in Fig. 2-3. This upgrade increased the number of active channels to 124 million and significantly reduced the amount of material in the detector [158]. This is the detector that was used during XeXe data-taking.

The pixel detector is an essential component for the studies done in this thesis. Its high granularity measurements are needed to initiate reconstruction algorithms to calculate the trajectories of the charged particles passing through CMS. Additionally, the precise vertexing capabilities it provides are crucial for constraining reconstruction algorithms so that they run in a reasonable amount of time for heavy ion collisions.




Figure 2-3: A comparison of the original and phase one pixel detectors in CMS. The left panel shows a side view with the phase one detector on top and the original detector on bottom. An isometric view is shown in the right panel. Figure taken from [158].

2.2 Strip Tracker

The CMS Silicon strip tracker sits outside the pixel detector and is designed to provide tracking information over a much larger volume. Like the pixel tracker, it provides information out to $|\eta| = 2.4$. It is arranged into four subsystems: the Tracker Inner Barrel (TIB), Tracker Outer Barrel (TOB), Tracker Inner Disks (TID), and Tracker EndCap (TEC). The TIB consists of four layers, having radii of 26, 34, 42, and 50 cm, and covering the region $|z| < 60$ cm. The silicon strips have a spacing, or pitch, of 80 or 120 μm , and a thickness of 320 μm . The TOB has six more layers having radii of 61, 69, 78, 87, 97, and 108 cm, which extend further along the beam line to cover $|z| < 110$ cm. The sensor thickness was chosen to be 500 μm in order to increase the signal to noise ratio of the sensors which are further away from the beam line. The strip pitches were also increased to 120 and 180 μm . The TEC comprises 9 disks arranged with their axes parallel to the beam pipe, covering the range in $|z|$ from 120 to 280 cm and r from 22 to 114 cm. The first three disks have sensors of thickness of 320 μm , while the last six use 500 μm silicon. A small gap between the TIB and TEC, at $80 < |z| < 90$ cm and $20 < r < 50$ cm, is filled with the TID. This contains three wheels of 320 μm thick sensors. The strip pitch in the TEC and TID varies from 95 to 180 μm . The two layers (disks) in the TIB and TOB (TID

and TEC) closest to the beam line are outfitted with “stereo” sensors made of two strip detectors glued to each other with the strips slightly misaligned by 100 mrad. This allows these sensors to provide 3-dimensional measurement by comparing the hit positions of the two sensors. The fifth wheel in the TEC is also comprised of these sensors. The rest of the tracker uses “mono” sensors providing 2D hit information. Although the spatial hit resolution depends on the pitch of the silicon strips used, typical hit resolutions are between 25 and 50 μm in all three dimensions. There are a total of 22400 sensors in the strip tracker, giving a total of 9.6 million channels covering an active area of 200 m^2 . A summary of the tracker layout can be seen in Fig. 2-2.

A simulation of the amount of material in the tracker is shown in Fig. 2-4. The units are given in terms of radiation lengths, X_0 , defined as the distance an electron will travel on average before only having $1/e$ of its original energy. The large increase in material budget around $|\eta| = 1.2$ leads to more scattering of particles off detector components. This adversely affects the tracking performance for heavy ion collisions in this pseudorapidity window. The relatively small material budget for $|\eta| < 1$, however, gives good tracking performance at midrapidity.

The precise spatial positions and orientations of the elements comprising the tracker, known collectively as the tracker alignment, are monitored by an array of infrared lasers which are directed at various modules in the tracker. The signals generated by the laser impinging on the sensor are read out to obtain the position of each sensor to an accuracy on the order of 100 μm . The alignment is also checked offline by using reconstructed tracks from collisions and from cosmic rays.

The readout chain for the strip detector is somewhat simpler than that of the pixel detector. A custom integrated circuit, known as an APV, collects the signals of 128 strips. The outputs of two APVs are multiplexed onto an optical link which transfers the data away from the detector to one of 450 FEDs in the adjacent service cavern. Each FED accepts 96 optical links and performs analog to digital conversion, pedestal subtraction, and common mode noise subtraction. In pp running, clusters of adjacent hit strips are output, but in the heavy ion data collection mode the information from




Figure 2-4: The simulated material budget of the CMS tracking system as a function of pseudorapidity [157].

every channel is output.

Typical detector occupancies in the first layer (disk) of the TIB (TID) are 2-3% during nominal pp operation. Occupancies in the outer parts of the tracker are on the order of 1%. Unlike the pixel detector, the occupancies in heavy ion collisions are large enough to cause significant distortions in the electrical readout of the strips. An example of this effect is illustrated by the trace of 768 adjacent silicon strips shown in Fig. 2-5. The gray line shows the raw signal, and the green line shows the signal after removing a constant pedestal offset. The sharp peaks in the green line correspond to particle hits, and should lie above a flat baseline. For most of the strips in this case the baseline (shown in red) is indeed flat and is reproduced using flat segments that are 128 strips in length. However a large charge deposit around strip number 380 causes a distortion in the baseline for strips 257 to 384. For these strips, a flat segment does not accurately reproduce the baseline and a more complicated “baseline restoration” algorithm is employed to fit the shape of the distorted baseline. The




Figure 2-5: An example of the readout from a set of 768 adjacent silicon strips. The raw signal is shown by the gray line, while the pedestal-subtracted data is given by the green line. The baseline (red) is distorted for the middle set of 128 strips, and must be reconstructed with a special algorithm. The final clusters after baseline subtraction are shown by the pink and blue lines.

final cluster measurements, after subtracting the red baseline from the green line, are shown by the blue and pink peaks.

Unfortunately, the baseline restoration algorithm can not be run on the FED readout electronics and must be applied at a later time; the information from every strip in the detector, corresponding to the green line, must be read out instead of just the pink and blue clusters. Reading out every strip quickly saturates the bandwidth available to record data, limiting the tracker readout rate to around 3 kHz. To alleviate this limitation, a more compact FED readout format was commissioned specifically for heavy ion data-taking in 2015.

A diagram showing a summary of the new and old output formats is shown in Fig. 2-6. In general, the data is stored between some header and footer information, which remains unchanged between the two formats. The previous data format placed




Figure 2-6: A schematic comparing two different readout formats used in the CMS tracker. Each box outlined in black represents 1 byte. Grey fields indicate information from a strip. The 16-bit format wastes 6 bits of space per strip encoded, as indicated by the red boxes. A special 10-bit format for heavy ion data exchanges this for a small padding region (purple) before the footer of the data payload.

the information from one strip into the space of 16 bits (2 bytes). This ensured that the data stream always consisted of an integer number of bytes. However, each strip outputs a number between 0 and 1023 and thus only requires 10 bits to store its full information. This means that 6 bits in each 16-bit word contain no useful information, as indicated by the red boxes in the figure. The new data format simply removes this useless information, and appends the 10 bits of each strip directly to the end of the previous one. Because the data stream must fit into an integer number of bytes, between 0 and 6 bits of padding are then added to the end of the data payload before the footer section. In this way, the 10-bit data format is able to store 4 strips into 5 bytes instead of the 8 bytes required by the 16-bit format. A special firmware implementation using this heavy ion readout format was deployed and validated on all CMS tracker FEDs with emulated data. Additionally, new software was written and validated so that the data could be successfully unpacked for offline processing. These changes were successful and pushed the readout limitation of CMS up to around 5 kHz in 2015.

The strip tracker allows for tracks found in the pixel detector to be significantly extended in length, L . If multiple scattering effects are ignored, the transverse momentum, p_T , resolution of a track is a function of the track length, the magnetic field

strength B , the intrinsic hit resolution σ , and the number of measurements made along the track, N [159]:

$$\frac{\delta p_T}{p_T} \propto \frac{\sigma p_T}{BL^2} \sqrt{\frac{720}{N+4}}. \quad (2.1)$$

Thus, the extra high-resolution hits and longer tracks provided by the strip tracker, along with the large magnetic field from the solenoid, allow for excellent track momentum resolutions. This advantage is exploited for nearly all of the physics objects used by CMS and is crucial for the measurements of large momentum charged particles done in this thesis. The exact performance of the reconstruction algorithms using the pixel and strip trackers are discussed in Section 3.2.

2.3 Electromagnetic Calorimeter

The ECAL barrel is placed directly outside of the strip tracker and has a radius of 1.29 m and pseudorapidity coverage of $|\eta| < 1.479$. It consists of a matrix of 61200 PbWO_4 crystals, each having dimensions of 22×22 mm and a depth of 23 cm, corresponding to 25.8 radiation lengths. Each crystal subtends 1° in azimuth. The barrel is capped on each side with an endcap covering the range $1.479 < |\eta| < 3$. Each endcap is formed from 7324 crystals which are $29 \times 29 \times 220$ mm in size. In both the barrel and the endcap, the crystals are projectively oriented towards the CMS interaction point so that particles entering the front face of a crystal are less likely to cross a boundary into another crystal. When an electron or photon hits an ECAL crystal an electromagnetic shower begins, exciting the atoms in the crystal. The atoms then emit scintillation light that is detected by an avalanche photodiode and converted into an electrical signal. Using the signals from adjacent crystals, the position of a high energy shower can be determined to within an accuracy of 0.5 mm. This allows precise matching of ECAL showers to tracks found in the tracking detectors for electron reconstruction. The amount of scintillation light collected by each photodiode is proportional to the incident particle's energy. Because PbWO_4 crystals are sensitive to small temperature changes and can lose optical transparency when irradiated, the energy response of the

ECAL crystals is monitored in real-time by periodically shining a calibration laser into them. After applying appropriate corrections, the ECAL energy resolution is parameterized as:

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2, \quad (2.2)$$

with a stochastic term, S , a noise term, N , and a constant term C . Estimates of these parameters from electron beam tests are $S = 0.0363 \pm 0.1\sqrt{\text{MeV}}$, $N = 124 \text{ MeV}$, and $C = (0.26 \pm 0.01)\%$. Above 500 GeV it is possible that electromagnetic showers leak out of the back of the calorimeter, but these energies are rarely reached in heavy ion collisions.

The ECAL crystals are read out in groups of 25. The front end electronics sum the signals of each crystal in a group and send this information to the CMS trigger system, before adding the data to a buffer. If the trigger system decides to accept the event a selective readout algorithm is executed. This procedure outputs the measurements coming from each group of crystals above a dynamic energy threshold that changes based on the crystals' proximity to other high energy deposits.

Minimally ionizing particles like pions and muons can also leave energy in the ECAL but have a high probability to penetrate through the full ECAL. Neutral hadrons will also not leave a significant amount of energy in the ECAL. Thus, another calorimeter is placed behind the ECAL.

2.4 Hadronic Calorimeter

The HCAL lies between the ECAL and the CMS solenoid and is designed to stop any hadrons penetrating the ECAL, including neutral particles like neutrons. Like the ECAL, it consists of a barrel detector covering $|\eta| < 1.4$, having an inner radius of 1.8 m and an outer radius of 2.9 m. The barrel is made of 18 absorber layers of brass and stainless steel, interleaved with 17 layers of plastic scintillator. The absorber layers are between 4 and 7 cm in thickness and the scintillators are 0.4 cm in thickness. The scintillator tiles have a segmentation of $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$ and are arranged

to form projective towers that point towards the interaction point. An extra set of scintillating tiles is placed behind the solenoid to catch any showers that are not contained in the main body of the HCAL. This addition brings the total HCAL depth to 11 interaction lengths at midrapidity. Two endcaps cover the region $1.3 < |\eta| < 3$. They are constructed in a similar fashion as the barrel HCAL, but the granularity of the scintillator tiles is progressively increased as a function of pseudorapidity.

As particles traverse the dense absorption layers, they have a chance to interact and cause a particle shower. The constituents of this shower will enter into a layer of scintillator and cause the emission of scintillation light. This scintillation light is gathered by wavelength shifting fibers and converted by a hybrid photodiode into an electronic signal. This signal is handled in a fashion similar to what is done in the ECAL, with a trigger signal being sent while the data is buffered, waiting for readout.

Using Eqn. 2.2, the HCAL has resolution parameters of $S = 1.25\sqrt{\text{GeV}}$, $N = 5.6 \text{ GeV}$, and $C = 3.3\%$.

The hadronic forward calorimeter (HF) fills in the gap in rapidity coverage between $2.9 < |\eta| < 5$ and is instrumental in characterizing the forward activity of heavy ion collisions. It consists of layers of steel, read out by quartz fibers which penetrate the absorber and run parallel to the beam pipe. These materials were chosen for their radiation hardness, as this forward region sees very large radiation fluences. Cherenkov light produced by charged particles traversing the quartz fibers is channeled to readout electronics at the back of the calorimeter. No instrumented material sits in front of the HF, and it therefore sees a mix of electromagnetically- and hadronically-interacting particles. To attempt to distinguish the relative energy contribution of these two different types of particles, some fibers only extend partially into the calorimeter, while others go all the way through. Electromagnetically-interacting particles tend to start a particle shower in the calorimeter at a relatively shallower depth and therefore leave a larger signal the long fibers.

In this thesis, the ECAL and HCAL are primarily used to validate the energy measurement of high energy charged particle candidates that are found in the tracker. If these candidates are indeed real particles, they should leave a large deposit of energy




Figure 2-7: A schematic of the CMS trigger system [160].

in the two calorimeters. If an energy deposit is not found, this increases the likelihood that the charged particle candidate is not real. The ECAL and HCAL are also used to reconstruct jet objects, which are needed for a measurement of the jet fragmentation function. Although the measurements completed here could in principle be executed using a detector that does not have full calorimeter coverage, the inclusion of these detectors greatly increases the robustness of the analyses.

2.5 Trigger and Data Acquisition

The collision rate delivered by the LHC is far too high for CMS to record all collision events. Therefore CMS is instrumented with a sophisticated trigger system [161] that quickly parses the detector information from each event and decides whether or not to store it. Making this decision accurately is crucial, as failing events are thrown away and lost forever. The system has two stages: the level one (L1) and high level trigger (HLT) systems. A diagram of the system can be seen in Fig. 2-7.

The L1 system is hardware-based to reject events failing certain loose selections

very quickly. The system is allocated a maximum of $3.2 \mu\text{s}$ to decide whether or not an event is accepted, but most of this time is consumed by the transit of the signals from the detector electronics to the L1 hardware. The actual logic decision must be computed in less than $1 \mu\text{s}$. The logical circuitry is constructed out of multiple custom Field-Programmable Gate Arrays (FPGAs) and Application Specific Integrated Circuits (ASICs). This system searches for energy deposits in the calorimeters or hits in the muon chambers. Information from the tracker is not used at this point. While the decision is being computed, all detector information is held in pipelined buffers. If an event is accepted these buffers are notified, and the full event data is collated into a complete package that is forwarded to the HLT.

In 2015, the PbPb collision rate of the LHC for CMS was around 15 kHz. However, the tracker was only able to read out data at around 5 kHz so the L1 system had to provide a reduction factor of around 3. In pp running, the collision rate is much higher and the required rejection rate can be as high as 300.

The HLT consists of a farm of consumer-grade computing nodes running sophisticated analysis algorithms. These algorithms are a near-copy of the full CMS offline reconstruction, with slight modifications to increase their speed. This produces high quality calibrated physics objects which can then be used to make an accurate determination of an event's value. At this point selections can be made on nearly any physics object or observable, including ones using tracker information, as long as the required computational resources are reasonable. In the 2015 PbPb run, the HLT decision reduced the data rate by a factor of ≈ 10 to fit into the available HLT output bandwidth and offline storage space.

A combination of desired L1 and HLT decisions is known simply as a 'trigger.' Some 'minimum-bias' (MB) triggers are quite simple, taking any event that passes an extremely loose set of selection criteria. These triggers are useful for evaluating detector efficiencies and the effectiveness of more complicated triggers. An example of a more complicated trigger is one that selects events having a high- p_T calorimeter cluster. This trigger would look for a large calorimeter deposit at the L1 stage, and then perform a more complete clustering of energy deposits in the HLT to decide if

the energy in any given area of the event is larger than some threshold.

Sometimes events that would pass a trigger are still voluntarily thrown away at a predefined rate to save bandwidth for other more interesting events. This ‘prescaling’ process allows physicists to fine-tune exactly how many events passing each trigger are collected. In practice, triggers for very rare events are never prescaled (unprescaled) while relatively common events are heavily prescaled. The exact prescale rate applied to each trigger is recorded in a database as a function of time, allowing a calculation of how many triggers were voluntarily discarded.

All the data used in this thesis was gathered with some form of trigger, and therefore an understanding of the performance of the trigger system is needed before one can analyze the data. Performance metrics like the total trigger efficiency for various data sets are discussed in Section 4.5.

Events surviving the prescale decision and passing a trigger are sent to the Tier-0 computing center at CERN where they are then fully reconstructed using complex algorithms. Copies of the raw input data, as well as the processed output events, are then shipped to data centers around the world for use and archival.

Chapter 3

Event Reconstruction

Data from CMS are recorded as large lists of electronic signals. This huge amount of raw information is not extremely useful in its unprocessed state and must be refined into objects that have a more physical interpretation, i.e. particles. The process of refining the raw data into higher-level physics objects is known as reconstruction. An exhaustive list of all reconstruction algorithms used in CMS would easily surpass the scope of this document. It is beneficial to just focus on techniques used for measurements of charged particle production in heavy ion collisions. Firstly, the determination of a heavy ion event's centrality is explained. Then the reconstruction of charged particles and their point of origin, known as 'tracking' and 'vertexing,' is discussed. Finally, a variety of jet-finding techniques are explored. The resulting objects can be seen in event displays such as Fig. 3-1.

3.1 Centrality

The impact parameter (b) of a collision of heavy nuclei is a vitally important quantity for determining the dynamics of the collision. Firstly, collisions with small b tend to have more nucleons participating because of the larger region of overlap between the two nuclei. Additionally, the geometry of this overlapping region also depends on b . For example, collisions having $b = 0$ would produce a roughly circular overlapping region, but larger impact parameters can give rise to regions that are lenticular-shaped.




Figure 3-1: A PbPb event collected by CMS in 2015 [162], highlighting the objects resulting from reconstruction. The green lines are tracks originating from a common vertex, while the red and blue towers are calorimeter deposits that can be clustered into jets.

These initial geometry differences are vitally important for explaining effects such as hydrodynamic flow [163]. Unfortunately, the collision impact parameter cannot be measured directly and can only be inferred using the particles captured by the experimental apparatus. During this inference procedure it is assumed that there is a monotonic increase in the number of particles produced (or total amount of energy seen in the detector) as b decreases. Events are therefore usually ordered by either the particle multiplicity or energy observed in the detector, and this ordering is thought to correspond to an ordering of events by their value of b as well. After this ordering procedure, the degree of overlap between the two nuclei in any one collision is usually expressed as a percentage quantile of a sample containing MB collisions. These quantiles are usually referred to as a collision's 'centrality'. Events having a smaller b typically are given lower quantile values. Thus, collisions having a centrality between 0–5% would be the 5% of collisions having the largest degree of overlap in a sample. These type of collisions are referred to as 'central.' On the other hand, collisions in

the 90–100% centrality range have a large impact parameter and represent ‘glancing’ blows between the two nuclei. These are referred to as ‘peripheral’ collisions.

In CMS, an event’s centrality is reconstructed using the HF subdetectors. The total transverse energy, $E_T = E/\cosh\eta$, of all towers in this system is computed for every event in a MB sample. An example of this distribution in 2.76 TeV PbPb data is given by the solid black line in Fig. 3-2. The sharp drop at around 125 TeV corresponds to the most violent, $b = 0$, collisions. Thresholds in E_T are then calculated to separate the distribution into 200 sections having an equal number of counts. Thresholds at larger HF values correspond to lower centrality percentiles and more central events. These centrality bins can be combined to contain a larger portion of the total hadronic cross section. In Fig. 3-2, the thresholds separating various centrality bins are shown by the dotted black vertical lines. Once the centrality thresholds have been calibrated, any event can be classified into a centrality bin based on the total HF energy.

If the MB trigger used to calibrate the centrality is not fully efficient, the centrality distribution will be biased because some of the total cross section is not contained in the HF distribution. This can be solved by estimating the fraction of the total hadronic cross section that is missing and adjusting the centrality thresholds to only contain the remaining fraction of the cross section. The estimation of the missing fraction of events is typically done using comparisons to simulated events and can lead to significant uncertainties. A similar procedure can also be applied to account for potential contamination of a MB sample by electromagnetic processes that are not part of the total hadronic cross section.

3.2 Tracking

As charged particles travel outward from the center of the detector they leave deposits of electric charge in the tracking detector. These deposits of charge, or hits, are recorded and can be used to reconstruct the trajectory of the passing charged particle. In general, these trajectories are shaped like helixes because of the large uniform




Figure 3-2: The HF distribution in 2.76 TeV PbPb data, illustrating how centrality quantiles are determined [28].

magnetic field at the center of CMS. A precise calculation of the radius of curvature (R) of this helix allows a determination of the charged particle's p_T using the formula:

$$p_T \approx 0.3qBR. \quad (3.1)$$

Here p_T is measured in GeV, q is the particle's electric charge in units of e , B is the magnetic field in Tesla (3.8 for CMS), and R is measured in meters. Determining the origin of the helix allows a reconstruction of the location of the original particle point of origin. Likewise, the endpoint of a helix usually can be matched to a calorimeter energy deposit created by the same particle. For these reasons the procedure of finding charged particle trajectories, known as 'tracking', is of utmost importance for understanding the dynamics of an event. The reconstructed trajectories themselves are simply called 'tracks'. In CMS tracks are described using the perigee convention [164] having 5 parameters. These parameters are the distances of closest approach (DCAs)

in the xy plane and the z direction as compared to a reference point, the polar and azimuthal angles (θ and ϕ , respectively) of the momentum vector, and the curvature ρ of the helix.

In general, the process of finding tracks from a set of hits is similar to a game of connect-the-dots. One simply tries to fit a helix through a series of hits in subsequent layers. Fig. 3-3 illustrates this by showing the hits (red circles) and corresponding tracks (radial black lines) for a toy 4-layer detector. In this example, the tracker is being viewed from the side, such that helix trajectories appear as segments of a circle. Finding all the tracks is fairly straightforward when there are only a few tracks in the event, as shown by the left column. However once the number of tracks increases, calculating all the trajectories quickly becomes difficult because each hit in a single layer can be matched to multiple hits in the next layer. This is illustrated by the right column. The number of combinations in a heavy ion event is so massive that tracking becomes extremely challenging, and any tracking algorithm becomes imperfect. These algorithms also tend to take a significant amount of computing time to run. Thus, the goal of tracking is to try to maximize the performance of the reconstruction algorithm, while keeping the computational resources required within a tolerable limit.

It should be emphasized here that a track and a charged particle are two distinct objects. A charged particle is a physical object that travels through the detector. A track is the output of the tracking algorithm, and may or may not correspond to the trajectory of a real charged particle. Some tracks may be misidentified and not correspond to any real particle.

3.2.1 The CMS Tracking Strategy

Tracking in CMS involves five major steps: hit reconstruction, track seeding, track propagation and fitting, track selection, and merging. Hit reconstruction is the process by which electronic signals from many adjacent sensors are combined to give information about the local position of a passing particle. A typical event in CMS will have many tracker hits and the number of hit combinations quickly becomes




Figure 3-3: An illustration of how hits in a toy 4-layer detector (top row) correspond to tracks (bottom row). Hits are shown by red circles. Tracks are black lines radiating from the center of the concentric rings, which are detector layers. For low detector occupancies (left column), it is fairly simple to find tracks, while the problem becomes very difficult as the occupancy grows (right column).

intractable. To alleviate this, short candidate tracks are found in the pixel detector during the seeding process. Only considering the pixels simplifies the combinatorial problem. These candidate tracks are then extended through the volume of the strip tracker during the propagation and fitting step. This significantly increases the length of the tracks and improves the resolution of the track kinematics. The tracks then have some light selections applied to reduce the amount of misreconstructed tracks. The track seeding and propagation procedures take a large amount of computational resources, and can not be reasonably run over all possible combinations of pixel hits at the same time. This is handled by first running these steps over only a subset of hits corresponding to tracks that are relatively ‘easy’ to find. The selected output tracks then have their hits removed from the total list of hits considered, and the

procedure is repeated again with a larger search window during the seeding. This is repeated multiple times. Each set of seeding, propagation, and selection steps run in this way is known as ‘iteration.’ This model of running multiple tracking passes is known as ‘iterative tracking.’ Once all iterations have been run, the output tracks are merged together into one track collection which can be used for analysis. More details on the CMS iterative tracking model can be found in Ref. [157].

3.2.2 Hit reconstruction

The first step in the tracking procedure is to construct the hits which are used as inputs. In the pixel detector, adjacent pixels having a signal above the noise threshold are grouped together to create a pixel cluster. The charge distribution in each cluster is then compared to a database of cluster charge templates [165] produced by the a PIXELAV simulation [166]. An example of the average charge collection profile of a set of pixels calculated in PIXELAV is compared to test beam data for four different bias voltages is shown in Fig. 3-4. The agreement between the simulation and data is excellent. The PIXELAV simulation is able consider details such as radiation damage levels, temperature variations, and the Lorentz drift inside the pixels. This allows the clustering to adapt to various detector conditions with a simple update of the template database. Based on the best match of a real pixel cluster with a template, the true position of a hit is estimated. In the strip detector, the adjacent strips having a signal above the noise threshold are also grouped together into clusters. The cluster positions and their uncertainties are calculated from a charge-weighted average of all the strips in the cluster. Corrections are applied for the Lorentz angle and inefficiencies in the sensor’s charge collection.

3.2.3 Track Seeding

Track seeds provide an initial estimation of the trajectory and origin of the charged particles in an event. A seeding algorithm depends on a list of ‘seeding layers’ and a ‘tracking region’.

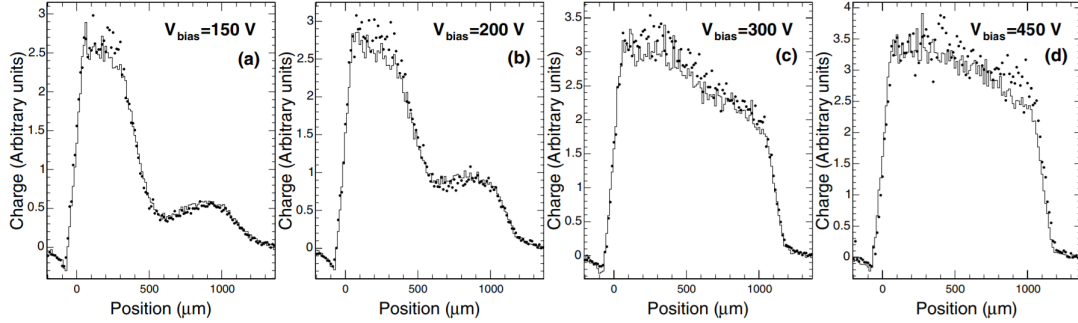


Figure 3-4: Pixel charge collection templates are shown for four different bias voltages. The calculation from PIXELAV is shown by the solid line, and a measurement with a test beam is given by the black dots [165].

The seeding layers are simply which layers of the tracker to consider when building seeds. To constrain the trajectory of a track, at least three three-dimensional hits are needed. Most seeding iterations therefore require three hits in the pixel detector, but some can require hits in both the pixels and the strips. With the addition of an additional pixel layer in the Phase 1 tracker upgrade in 2016, seeding with 4 pixel hits became possible. The addition of an extra hit further constrains the initial trajectory estimation of the seed and helps reduce combinatorial backgrounds. Finally, it is possible to build a seed with only two pixel hits by assuming that the track originates from the luminous region of the interaction point. This luminous region is commonly referred to as the beam spot. At the beginning of a run, a very broad estimation for the beam spot is used. As tracks and vertices are reconstructed, the beam spot position is refined to a much more precise quantity using the procedures described in Ref [157].

The tracking region is the other major parameter in a seeding algorithm. It describes a three-dimensional region in space that the trajectory of a track must originate from. Additionally, it constrains the p_T of tracks the seeding algorithm attempts to find. In general, iterations having smaller tracking regions and higher p_T thresholds take significantly less computational time. Thus, these tracking regions are the main parameters that are modified to create a version of the CMS tracking that is fast enough to run in heavy ion events. Most tracking regions take the form of

a three-dimensional cylinder, centered on a primary vertex (described in Section 3.3) or the beam spot and extending down the beam line. The main parameters are half the cylinder’s length (z_{region}), the origin radius (R_{origin}), and the minimum transverse momentum (p_T^{min}).

If a cylindrical tracking region is centered on a vertex, this vertex is usually built from the tracks found in the first iteration. In the case of the 2015 PbPb data, the vertex comes from a special work flow that only uses the pixel detector (described in Section 3.3). Other types of tracking regions include ‘jet seeded’ ones, which only consider the areas inside the cores of high- p_T jets, and ‘point seeded’ ones that cover a specific area in the (η, ϕ) plane. Point-seeded tracking regions are typically used to fill in specific detector areas that are unable to create other types of seeds because of dead detector modules. Scaling tracking regions shrink their size as the tracker occupancy increases. This prevents the computational time of the tracking from getting out of control.

Once the seeding layer and tracking regions are specified, the seeding algorithm can be run. The most basic seeding algorithm searches for pairs of hits in the inner two seeding layers that are compatible with a track originating from the tracking region. If there are still more seeding layers to check, an additional search is done to find a third hit in the next seeding layers. The final set of hits must be compatible with a track originating from the tracking region to be accepted. The shape of the charge distribution in the hits must also be compatible with the trajectory of the seed, i.e. if the seed enters the layer at a perpendicular angle then the hit should have a narrow charge profile.

With the phase 1 tracker, the original three layer pixel detector in CMS was replaced with a new 4 layer one. The previous procedure can be easily extended to create 4-hit seeds, but instead a faster ‘cellular automaton’ (CA) seeding is used [167]. In this framework, adjacent hits in two different layers are paired together. This pairing procedure is the same as the first step of the previously described seeding algorithm. One hit can be paired to multiple different hits in other layers. A pairing of two hits is referred to as a ‘cell.’ Cells are considered ‘connected’ if they share a




Figure 3-5: An illustration of the CA algorithm in a toy 5-layer detector [167]. (Left) Hits in all layers, shown by the blue dots, are connected to form cells, represented by the black lines. (Right) An iterative procedure gives the cells scores. Here black corresponds to a score of 0, red to 1, orange to 2, green to 3, and blue to 4. Seeds are then found by starting at the outside and moving inward.

hit. Connected cells are ‘neighbors’ if the difference in polar and azimuthal angles, $\Delta\theta$ and $\Delta\phi$, with respect to the beam spot is less than some threshold. Typical cut values are around $\Delta\theta < 0.002$ and $\Delta\phi < 0.2$. An example diagram of the hits and cells in a toy 5-layer detector is shown in the left panel of Fig. 3-5. Each cell has a score that is originally initialized to 0. An iterative procedure then begins. All cells are looped over, and a cell’s score is increased by one if it has a neighbor of equal score that is also closer to the center of the detector. This procedure is repeated until no cells change their score in subsequent iterations. The outcome of such a procedure is illustrated in the right panel of Fig. 3-5. The cell scores are shown as colors, with black representing 0, red being 1, orange being 2, green being 3, and blue being 4. Scores generally increase the further away from the center of the detector a cell is, as long as a continuous chain of inner cells exists to form a pathway to the center. Track candidates are then formed by selecting outer cells with a large score and following the attached chain of cells back towards the center of the detector. Ambiguities from seeds that share hits can be resolved with a simple χ^2 test of the seed’s hits to an expected track trajectory.

The order of the seeding iterations is optimized to reduce the computational time while providing good tracking performance. In general, the first iterations require three or four hits in the pixel tracker. These hits need to be compatible with a

track with p_T above ≈ 1 GeV and must also have a small DCA with respect to the beam spot. This increases the probability that any tracks found will originate from the primary interaction. Once this iteration is finished, subsequent iterations either relax the p_T requirement, the DCA requirement, or the number of pixel hits required. Finally, iterations that use seeding layers from the strip detector are also included. Hits from tracks found in previous iterations are removed from the set of hits used for seeding in nearly all cases. The only exception to this rule is the ‘Jet Core’ iteration, which always runs on the full set of tracker hits.

A summary of the different iterations and their parameters for all the data sets used here is given in Tables 3.1 and 3.2. The exact tracking setup evolved from year to year, with recent years have a significantly more complicated tracking.

Table 3.1: The track seeding iterations for the pp and pPb data used in this work. Here ‘BS’ and ‘PV’ refer to tracking regions centered on the beam spot and primary vertex, respectively. Quantities given in units of σ refer to a number of standard deviations of the primary vertex or beam spot uncertainty in the z direction.

| Number | Name | Layers | Tracking Region | z_{region} | R_{origin} (cm) | p_T^{min} (GeV) |
|----------------------------------|------------------------|-----------------|-----------------|--------------|-------------------|-------------------|
| 7 TeV pp (2011) | | | | | | |
| 0 | Initial step | 3 pixels | BS | 3σ | 0.2 | 0.8 |
| 1 | Pixel Pair step | 2 pixels | PV | 0.2 cm | 0.2 | 0.6 |
| 2 | Low p_T Triplet step | 3 pixels | BS | 3.3σ | 0.2 | 0.075 |
| 3 | Mixed Triplet step | 3 pixels+strips | BS | 10 cm | 1.2 | 0.35 |
| 4 | PixelLess step | 2 strips | BS | 10 cm | 2.0 | 0.5 |
| 5 | TOB/TEC step | 2 strips | BS | 30 cm | 5.0 | 0.6 |
| 2.76 TeV pp, 5.02 TeV pPb (2013) | | | | | | |
| 0 | Initial step | 3 pixels | BS | 4σ | 0.02 | 0.6 |
| 1 | Low p_T Triplet step | 3 pixels | BS | 4σ | 0.02 | 0.2 |
| 2 | Pixel Pair step | 2 pixels | PV | 0.03 cm | 0.015 | 0.6 |
| 3 | Detached Triplet step | 3 pixels | BS | 15 cm | 1.5 | 0.3 |
| 4 | Mixed Triplet step | 3 pixels+strips | BS | 10 cm | 1.5 | 0.6 |
| 5 | PixelLess step | 2 strips | BS | 10 cm | 2.0 | 0.7 |
| 6 | TOB/TEC step | 2 strips | BS | 30 cm | 6.0 | 0.6 |
| 5.02 TeV pp (2015) | | | | | | |
| 0 | Initial step | 3 pixels | BS | 4σ | 0.02 | 0.6 |
| 1 | Detached Triplet step | 3 pixels | BS | 15 cm | 1.5 | 0.3 |
| 2 | Low p_T Triplet step | 3 pixels | BS | 4σ | 0.02 | 0.2 |
| 3 | Pixel Pair step | 2 pixels | PV | 0.03 cm | 0.015 | 0.6 |
| 4 | Mixed Triplet step | 3 pixels+strips | BS | 10 cm | 1.5 | 0.6 |
| 5 | PixelLess step | 3 strips | BS | 12 cm | 1.0 | 0.4 |
| 6 | TOB/TEC step | 3 strips | BS | 30 cm | 6.0 | 0.6 |
| 7 | Jet Core step | 2 pixels+strips | jet cores | – | – | 10 |

Table 3.2: The track seeding iterations for the PbPb and XeXe data used in this work. Here ‘BS’ and ‘PV’ refer to tracking regions centered on the beam spot and primary vertex, respectively. Quantities given in units of σ refer to a number of standard deviations of the primary vertex or beam spot uncertainty in the z direction. $\Delta\theta$ and $\Delta\phi$ are provided when Cellular Automaton seeding is employed.

| Number | Name | Layers | Tracking Region | z_{region} | R_{origin} (cm) | p_T^{min} (GeV) | $\Delta\theta$ | $\Delta\phi$ |
|----------------------|---------------------------|-----------------|-----------------|--------------|-------------------|-------------------|----------------|--------------|
| 5.02 TeV PbPb (2015) | | | | | | | | |
| 0 | Initial step | 3 pixels | PV | 0.2 cm | 0.2 | 1.5 | – | – |
| 1 | Detached Triplet step | 3 pixels | PV | 15 cm | 0.5 | 0.9 | – | – |
| 2 | Low p_T Triplet step | 3 pixels | PV | 4σ | 0.02 | 0.4 | – | – |
| 3 | Pixel Pair step | 2 pixels | PV | 4σ | 0.005 | 1.0 | – | – |
| 4 | Jet Core step | 3 pixels+strips | jet cores | – | – | 15 | – | – |
| 5.44 TeV XeXe (2017) | | | | | | | | |
| 0 | Initial step | 4 pixels | BS | 4σ | 0.02 | 0.6 | 0.0012 | 0.2 |
| 1 | Low p_T Quadruplet step | 4 pixels | PV | 0.5 cm | 0.02 | 0.49 | 0.0017 | 0.3 |
| 2 | High p_T Triplet step | 3 pixels | PV | 0.2 cm | 0.02 | 0.7 | 0.004 | 0.07 |
| 3 | Low p_T Triplet step | 3 pixels | PV | 4σ | 0.02 | 0.49 | 0.002 | 0.05 |
| 4 | Detached Quadruplet step | 4 pixels | PV | 3.75 cm | 1.5 | 0.9 | 0.0011 | – |
| 5 | Detached Triplet step | 3 pixels | PV | 2.5 cm | 1.5 | 0.9 | 0.002 | 0.05 |
| 6 | Pixel Pair step | 2 pixels | point-seeded | 0.03 cm | 0.015 | 0.6 | – | – |
| 7 | Mixed Triplet step | 3 pixels+strips | scaling PV | 3.75 cm | 1.5 | 0.4 | – | – |
| 8 | PixelLess step | 3 strips | scaling PV | 3 cm | 1.0 | 2.0 | – | – |
| 9 | TOB/TEC step | 3 strips | scaling PV | 5 cm | 3.5 | 2.0 | – | – |
| 10 | Jet Core step | 2 pixels+strips | jet cores | – | – | 5.0 | – | – |

3.2.4 Track Propagation and Fitting

Once seeds have been produced, they must be extended through the tracker volume by adding additional hits. This is done using a Kalman Filter algorithm [168, 169, 170]. The Kalman filter is the optimal solution for a discrete dynamic system which has some state described by variables x_k that evolve in a linear fashion between steps. The solutions are equivalent to least-squares fits of the measurements of the variable x_k over a given amount of steps. One can think of the track-finding problem as a dynamic system in which a track's trajectory evolves as more measurements (tracker hits) are added onto the track.

In the following, the arguments of Ref. [168] are followed. Consider a track candidate that is already known to cross k layers at positions \mathbf{x}_k , where \mathbf{x} is a vector of x , y , and z coordinates along with associated momenta. The variables \mathbf{x}_k have covariances \mathbf{C}_k associated with them. The desired goal is to predict the state and covariance at the $k + 1$ th layer given what is already known after k layers. Denote these quantities as \mathbf{x}_{k+1}^k and \mathbf{C}_{k+1}^k . Additionally, denote a linear track propagator \mathbf{F}_k that describes how the variables \mathbf{x}_k transform to \mathbf{x}_{k+1} in the absence of multiple scattering effects. Parameterize the effects of multiple scattering with variable \mathbf{w}_k , such that each w_k is a random Gaussian centered around zero. Let $\text{cov}(\mathbf{w}_k) = \mathbf{Q}_k$. With this notation, the predicted variables are

$$\mathbf{x}_{k+1}^k = \mathbf{F}_k \mathbf{x}_k \tag{3.2}$$

and updated covariances are

$$\mathbf{C}_{k+1}^k = \mathbf{F}_k \mathbf{C}_k \mathbf{F}_k^T + \mathbf{Q}_k. \tag{3.3}$$

Note how the addition of multiple scattering does not affect the predicted vector at layer $k + 1$, but merely increases the covariance. This is due to the assumption that w_k is centered around zero. Now denote quantities that are measured in the detector (positions of hits at the k th layer) as \mathbf{m}_k . The track parameters can be related to the

measurements using a matrix \mathbf{H}_k , along with a vector \mathbf{v}_k having covariance \mathbf{V}_k that describes a Gaussian uncertainty associated with each measurement:

$$\mathbf{m}_{k+1} = \mathbf{H}_{k+1}\mathbf{x}_{k+1} + \mathbf{v}_{k+1}. \quad (3.4)$$

Assume \mathbf{v}_k are also Gaussian random variables centered around 0. Then the residuals, \mathbf{r}_{k+1}^k , and their covariance, \mathbf{R}_{k+1}^k at $k+1$ using information from the previous k layers are given by

$$\mathbf{r}_{k+1}^k = \mathbf{m}_{k+1} - \mathbf{H}_{k+1}\mathbf{x}_{k+1}^k \quad (3.5)$$

and

$$\mathbf{R}_{k+1}^k = \mathbf{H}_{k+1}\mathbf{C}_{k+1}^k\mathbf{H}_{k+1}^T + \mathbf{V}_{k+1}. \quad (3.6)$$

Thus, the uncertainties \mathbf{V}_k in the measurements also affect the uncertainties of the residuals.

Now assume that there is a new measurement on layer $k+1$ that can be added to the track. The estimated parameters using the previous layer information, \mathbf{x}_{k+1}^k , can be updated to a new estimation that includes an extra layer \mathbf{x}_{k+1} . This is done using the Kalman gain matrix:

$$\mathbf{K}_{k+1} = \mathbf{C}_{k+1}^k\mathbf{H}_{k+1}^T(\mathbf{R}_{k+1}^k)^{-1}. \quad (3.7)$$

With this definition, the updated values after adding a measurement to the track are

$$\mathbf{x}_{k+1} = \mathbf{x}_{k+1}^k + \mathbf{K}_{k+1}(\mathbf{m}_{k+1} - \mathbf{H}_{k+1}\mathbf{x}_{k+1}^k) \quad (3.8)$$

and

$$\mathbf{C}_{k+1} = (\mathbf{I} - \mathbf{K}_{k+1}\mathbf{H}_{k+1})\mathbf{C}_{k+1}^k. \quad (3.9)$$

Because it contains the inverse of \mathbf{R}_{k+1}^k , \mathbf{K}_{k+1} quantifies how much weight should be given to the new measurement when updating the track's parameters. If the residuals are large compared to the covariance of the track, the entries in \mathbf{K}_{k+1} will be small,

and therefore the updated parameters in Eqn. 3.8 will be very close to the parameters without the measurement added. If, on the other hand, the new residuals are small then \mathbf{x}_{k+1} will favor the new measurement rather than the predicted location from the propagation step. After addition of the new measurement, this procedure can then repeat by propagating to layer $k + 2$ and so on and so forth. After all measurements have been added, a final fit to all the hits can be performed to update that quantities that were calculated before all the hits were added to the track.

A cartoon representation of this process is shown in Fig. 3-6. An initial estimation of the track's trajectory is shown by the blue arrow originating from layer 0. The uncertainty in this trajectory is shown by the light blue shaded region. On layer 1 a new hit (red circle) is found that is compatible with the track candidate. A new trajectory is calculated and used to propagate to layer 2 and so on. As more hits are added to the track the uncertainty on the trajectory decreases. An example track that could be fit to these hits is shown by the dashed green line.

In CMS track candidates are propagated to the next layer using an analytic propagator that assumes an ideal magnetic field. Although this is not most accurate propagation, it is very fast. Once a candidate has been propagated to a new layer, a search for tracker modules compatible with the track trajectory is initiated. Modules are considered compatible if the trajectory is within 3σ of a module boundary. Up to two modules can be found per tracker layer, because the sensors in each layer have a slight overlap that allows one track to hit two modules. Once modules are identified, clusters in those modules are checked against the trajectory of the track. If a cluster is compatible with the track candidate, as determined by a χ^2 test, the cluster is appended to the track. Multiple potential clusters can be appended to a track candidate and each forms a new independent candidate. However, an upper limit on the number of track candidates coming from a single seed is put in place to prevent exponential growth of track candidates. This limit is usually set to around 5. If this limit is exceeded, candidates that have the largest χ^2 are deleted. These new track candidates then have their trajectories updated and are propagated to the next layer. This procedure can produce duplicates of the same track if two seeds roughly




Figure 3-6: A cartoon of a Kalman filter procedure. Hits are shown as red circles. The predicted trajectory and its uncertainty are shown by the blue arrow and the shaded triangular regions. A example of a track producing these hits is shown by the dashed green line.

point towards the same track, so a cleaning step is also applied at the end of the procedure. All pairs of track candidates are compared, and if the fraction of shared hits

$$f_{shared} = N_{shared}^{hits} / \min(N_1^{hits}, N_2^{hits}) \quad (3.10)$$

is greater than 19%, the candidate with less hits is discarded. If the candidates have the same number of hits, the track with a larger χ^2 is removed.

Once a list of hits is associated with each track candidate, a full Kalman filter fit and smoothing procedure is initialized. The filter is first run outward from the track seed. A second filter is then run backwards from the outer hits. The optimal track parameters at any point along the track can then be found as a weighted average

of the results of these two filters; the first depends only on the hits found ‘inside’ the point of interest while the second depends on hits found ‘outside’ the point of interest. During the propagation step of the filters, a precise Runge-Kutta propagator is used to predict the position of the track at the next layer. This propagator uses a very accurate magnetic field map that was measured before the LHC delivered any collisions.

3.2.5 Track Selection

Not all tracks are the result of true charged particles. A light selection is applied on the output of the Kalman filter to reduce the amount of misreconstructed tracks in the sample. In data gathered before 2015, this is done by cutting on the track χ^2/dof , DCAs, number of hits, and p_T resolution. Starting in 2015, the selection uses a Boosted Decision Tree (BDT) [171] implemented using the ROOT multivariate toolkit (TMVA) [172]. A BDT is a set of shallow tree structures that iteratively apply cuts on input variables and return either -1 or 1. These trees are trained on a simulated sample and learn how to classify tracks as good or bad. The BDT takes an array of track quality variables as input and returns a single number that is the average of the results of many different trees. An output closer to 1 indicates that the track has a higher quality. The variables included in the BDT are the p_T and its uncertainty, η , the number of hits, number of layers, χ^2 , number of degrees of freedom and number of layers without a hit. Iterations searching for tracks coming from the primary interaction also include the track DCAs in both the transverse and longitudinal directions. The BDTs used for the 2015 pp and 2017 XeXe data were trained on a cocktail of simulated pp events having different physics processes. The BDT for the 2015 PbPb run was trained on simulated PbPb events.

Based on the output of the BDT, a track is classified into one of four categories. The first category consists of very poor tracks. These tracks are simply deleted. The next three categories are called — in order of increasing track quality — *loose*, *tight*, and *high purity*. Any tracks that are at least *loose* are kept in the sample. Only tracks that are *high purity* have their hits removed before a new iteration is run. This

means that a *loose* or *tight* track may be found again in a subsequent iteration. In practice, only *high purity* tracks are used for analysis. However the *high purity* cut itself is not a very stringent requirement despite its name. Thus, additional cuts are usually applied to the tracks for offline analysis.

3.2.6 Track Merging

Because it is possible for the same track to be found by multiple tracking iterations, a merging procedure is needed to identify and remove duplicate tracks. This is done using the same procedure described Section 3.2.4. The number of shared hits between two tracks, given by Equation 3.10, is required to be greater than 19% to be a duplicate track. This 19% threshold was determined empirically and is configurable. The two tracks are then given a score S based on the number of pixels hit, the number of strips hit, the number of missing hits, and the track χ^2 :

$$S = 20N_{pix} + 10N_{strips} - 5N_{missing} - \chi^2. \quad (3.11)$$

The track having the higher score is kept. If the track scores are within 0.01 of each other, the track with a better BDT classification is kept. If the tracks also have the same BDT classification, the track coming from the earliest iteration is kept. After this procedure a full collection of tracks is written out for use in analysis.

3.2.7 Tracking Performance

The tracking performance is quantified in three major ways. The first is the resolution of the track properties. A track is only useful if its properties correspond closely to the kinematics of the charged particle that created the track. The second performance criteria is the tracking efficiency. This is the probability that a charged particle will be successfully reconstructed as a track, and should be as high as possible. The last quantity of interest is the misreconstruction rate. If a large number of tracks do not correspond to a real charged particle then the sample is contaminated. This contamination, as well as tracking inefficiencies, must be corrected for.




Figure 3-7: The track p_T resolution for XeXe collisions for two different track selections. The ‘All Trk Selections’ points refer to the selections given for XeXe data in Table 5.1.

Resolution

The track p_T resolution is excellent in CMS because of the large magnetic field. A plot of the resolution as a function of track p_T can be seen in Fig. 3-7 for XeXe collisions. For most of the p_T range of interest, the resolution is around 1%. This flat behavior is caused by multiple scattering in the material of the tracker. As the p_T increases, the resolution also slowly gets worse. At 400 GeV, the highest p_T measured in these analyses, the resolution is about 5%. The resolution of analysis-quality tracks was found to not change significantly with the detector occupancy; a similar resolution is achieved in pp and PbPb events. The DCA resolution in both the transverse and longitudinal directions is shown in Fig. 3-8 for PbPb collisions. Low- p_T tracks tend to have resolutions on the order of a few hundred μm . However, the resolution decreases to around 50 μm as the track p_T increases. The DCA resolution also does not strongly depend on detector occupancy. The pointing resolution in η and ϕ are also excellent. The variation of these parameters from reconstruction effects is negligible for the analyses presented here.




Figure 3-8: The DCA resolutions for 2015 PbPb tracks.

Efficiency

The tracking efficiency is the probability to reconstruct a real charged particle as a track. In general the tracking efficiency can be degraded by many factors including dead detector modules, areas with a large amount of detector material, large detector occupancies, and algorithmic inefficiencies. Thus, the efficiency is a complex function of the kinematics of the tracks being searched for. It is usually parameterized as a function of some combination of p_T , η , and ϕ . Understanding the centrality-dependence is also very important for heavy ion collisions.

The efficiency also depends on choices from the analyzer. If a tight set of track selections is used, less tracks are used in the analysis and the efficiency will be lower. However, the misreconstruction rate will also be lower too. Because the misreconstruction rate is naturally fairly low in pp and pPb collisions, a highly efficient working point is chosen when analyzing these data. However, the misreconstruction rate gets




Figure 3-9: (Left) The tracking efficiency for 2015 5.02 TeV pp collisions as a function of p_T . (Right) The efficiency for inclusive 2015 PbPb collisions. Typical analysis selections have been applied.

very large in central heavy ion collisions when the detector occupancy is very large. To suppress this contamination a much tighter selection is used. This also naturally makes the tracking efficiency much lower.

Examples of the tracking efficiency in the range $|\eta| < 1$ as a function of p_T are given in Fig. 3-9. The left panel shows the typical efficiency for 2015 5.02 TeV pp collisions. Typical analysis selections have been applied. It is around 90% for most of the range examined. There is a sudden dip around 20 GeV, corresponding to the start of an additional requirement that necessitates the track to be matched to a calorimeter deposit. The downward trend at very high p_T is caused by the tendency for secondary particles produced by interaction with detector material to align with the trajectory of the initial charged particle. This effect becomes more pronounced for high- p_T particles and can confuse the Kalman filter [157]. The efficiency for inclusive 2015 PbPb collisions is shown in the right panel of Fig. 3-9. In general, the efficiency is lower because of the higher detector occupancy and tighter selection requirements. This is particularly pronounced under 1 GeV, where the efficiency rapidly drops.

An example of the centrality dependence of the tracking efficiency is shown for 2017 XeXe collisions at 5.44 TeV in Fig. 3-10. The colored bands around the lines represent the statistical uncertainties of the efficiency. At high p_T the tracking efficiency is around 0.7 but rapidly decays at low p_T . This effect is most pronounced in the




Figure 3-10: The tracking efficiency for 5.44 TeV XeXe collisions. The efficiency is given for six different centrality selections. The colored bands show the statistical uncertainty [173].

0–5% centrality bin. The slope of the efficiency loss becomes largest around $p_T = 0.7$ GeV. This p_T corresponds to the minimum momentum a particle needs to successfully cross all layers in the tracker and hit the calorimeter. Thus, particles below this threshold are highly curved and can form multiple loops in the tracker volume. It is difficult for a Kalman filter to reconstruct such a trajectory consistently.

Misreconstruction Rate

The misreconstruction rate (sometimes known as the fake rate f) is the fraction of reconstructed tracks that do not correspond to a real charged particle. They are usually produced by random combinations of hits and real track fragments that happen to roughly align. Naturally, if there are more hits in the detector then the probability of this happening increases. The misreconstruction rate is quite low in pp




Figure 3-11: (Left) The quantity $1/(1-f)$ for 2015 5.02 TeV pp collisions as a function of p_T . (Right) The same quantity for inclusive 2015 PbPb collisions. Typical analysis selections have been applied.

and pPb collisions but can be very large for central PbPb events.

The quantity $1/(1-f)$ for 2015 5.02 TeV pp collisions is shown in the left panel of Fig. 3-11. For small values of f , this is approximately equal to $1+f$. The misreconstruction rate is less than 3% for the entire p_T range examined. The same quantity for inclusive 5.02 TeV PbPb collisions is shown in the right panel of the same figure. This rate is a little bit higher but is still fairly low because of the stringent track selections applied in this case. Furthermore, most of the misreconstructed tracks are usually concentrated in the 0–10% centrality range. This can be seen in Fig. 3-12, which shows $1-f$ for 5.44 TeV XeXe collisions for six different centrality selections, as evaluated in three different types of simulation. In the 0–5% centrality bin, the misreconstruction rate can be as high as 35% at $p_T < 0.7$ GeV, resulting in large systematic uncertainties in this p_T region. It quickly falls to more tolerable levels as the track p_T increases, however.

The effects of tracking inefficiency and misreconstruction must be corrected for at analysis level. These corrections will be discussed in more detail later.

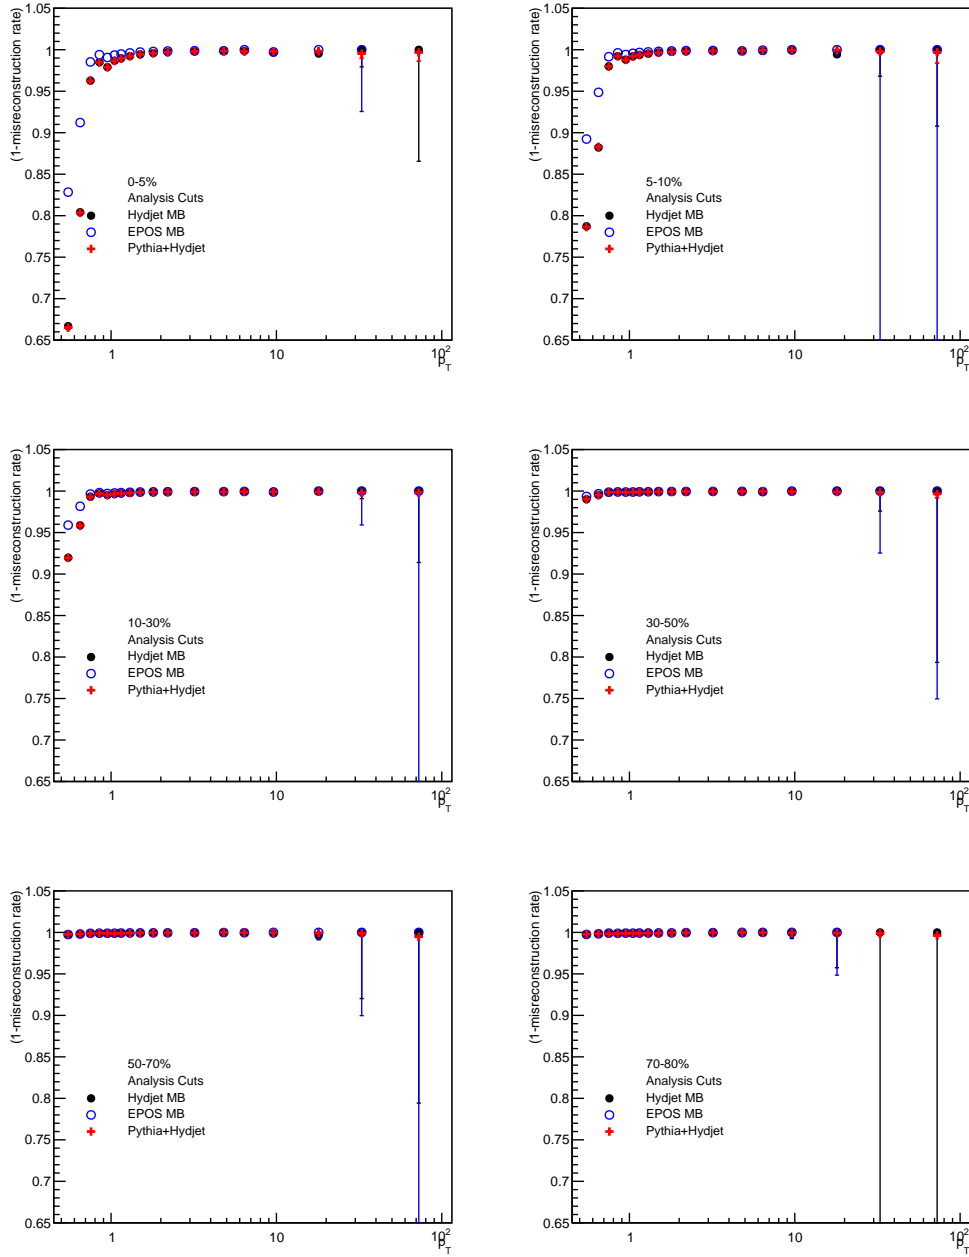


Figure 3-12: The value $(1 - f)$ for six different centrality selections in 5.44 TeV XeXe collisions. The results from three different types of simulation are shown. The fake rate becomes very large for central collisions.

3.3 Vertexing

A vertex is a common origin of two or more tracks that indicates the spatial position of a physical process that results in multiple daughters. A ‘primary’ vertex is one created by the initial interaction between the protons or heavy nuclei being accelerated by the LHC. ‘Secondary’ vertices originate from the decays of unstable particles produced by the primary interaction or subsequent decays. Here we focus on the reconstruction of primary vertices.

In general, there can be more than one primary vertex in an event. This effect is known as pileup. In practice, this is only relevant for pp collisions, where it is possible to have tens of interactions in every event. In this scenario, it is important to be able to correctly associate which tracks come from the same primary vertex, and therefore the same physics process. In heavy ion collisions this ambiguity does not exist as there is usually only one true primary interaction. However, the position of the primary vertex is still a useful quantity; it allows the rejection of secondary tracks based on their distance of closest approach to the vertex.

Primary vertexing is done in three steps. Firstly, the input tracks are selected. Secondly, the tracks are grouped together into clusters that correspond to the tracks originating from the same interaction. Finally, a fitter is run on the tracks in each cluster to calculate the final primary vertex properties. The full details of the CMS vertexing algorithms can be found in Ref. [157].

3.3.1 Input Tracks

Input tracks are selected by requiring them to be consistent with originating from a primary interaction. This is done by requiring the transverse distance of closest approach (d_0) significance with respect to the beam spot center to be less than 5 standard deviations (σ). In XeXe collisions this was reduced to only 3σ to save computational time. Tracks must have at least 2 pixel hits and 5 strip hits associated with them. They must also have a fit $\chi^2 < 20$.

In PbPb collisions a special vertexing using pixel tracks instead of full tracks is

Table 3.3: The track selections used during primary vertex reconstruction.

| | pp, pPb | XeXe | PbPb |
|--------------|-------------|-------------|-------------|
| track type | full | full | pixel |
| d_0 DCA | $< 5\sigma$ | $< 3\sigma$ | $< 2\sigma$ |
| χ^2 | < 20 | < 20 | < 5 |
| N_{pixels} | ≥ 2 | ≥ 2 | ≥ 2 |
| N_{strips} | ≥ 5 | ≥ 5 | - |

used to save computational time. First, the approximate longitudinal position is estimated by scanning along z and calculating the number of pixel clusters that are compatible with an event at that position. The z value that has the maximum number of compatible clusters is taken as the nominal position for the following steps. Next, a dedicated pixel tracking iteration is run in the region around the initial z estimate. If a large number of pixel clusters are detected, this dedicated iteration will shrink its tracking region with the goal of reducing computational time while still producing enough pixel tracks for a vertex to be found. Like the full track case, the pixel tracks are then subjected to a series of selections. The track d_0 significance must be below 3σ , the $\chi^2 < 5$, and at least two pixel hits are required. The pixel tracks are then fed into the next step of the vertexing. A summary of the track selections used in the vertexing is given in Table 3.3.

3.3.2 Track Clustering

After tracks are selected they must be grouped together into clusters corresponding to separate primary interactions. In XeXe and PbPb collisions, where the pileup is very low and only one interaction is expected on average, a simple ‘gap’ clustering algorithm is used. The tracks are sorted by their DCA in z to the beam spot, and adjacent tracks are clustered together. Any tracks that have a z separation larger than $z_{sep} = 1$ cm are assigned to different clusters. This simple algorithm does a good job of finding a single large vertex but tends to merge separate vertices that are positioned close together. For this reason a more complicated Deterministic Annealing (DA) algorithm [174] is used in pp and pPb collisions.

The DA algorithm attempts to minimize the global χ^2 of the system of tracks and

vertices. It does this in a way that makes strong contact with statistical mechanics, by modeling tracks and vertices as a system with a large number of degrees of freedom. This system is then allowed to relax into a minimum energy state by reducing its temperature. To achieve this, a pairing of a track and a vertex is assigned an ‘energy’:

$$E_{ij} = \frac{(z_i^T - z_k^V)^2}{\sigma_z^2}, \quad (3.12)$$

where z_i^T is the z DCA to the beam spot of the i th track, σ_z is the associated uncertainty, and z_k^V is the z position of the k th vertex. For a given ‘temperature’ T , the free energy is given by:

$$F = E - TS = \sum_i^{N_{trk}} \sum_k^{N_{vtx}} p_{ik} E_{ik} - T \sum_i^{N_{trk}} \sum_k^{N_{vtx}} p_{ik} \log p_{ik}. \quad (3.13)$$

Here p_{ik} are the probabilities that track i is legitimately associated with vertex k . Thus, the set of sums on the right hand of Eq. 3.13 corresponds to the Shannon entropy of the system. The probabilities p_{ik} are related to the track-vertex pair energies and temperature via a Boltzmann distribution:

$$p_{ik} = \frac{w_i \rho_k \exp(-E_{ik}/T)}{\sum_{i'}^{N_{trk}} \sum_{k'}^{N_{vtx}} w_{i'} \rho_{k'} \exp(-E_{i'k'}/T)}. \quad (3.14)$$

The ρ_k above represent a relative weight given to each vertex, subject to the constraint that $\sum_k^{N_{vtx}} \rho_k = 1$. The terms w_i are constant weights given to each track based on their d_0 DCA to the beam spot; these do not change during the execution of the algorithm.

The algorithm is initialized by setting T to a large number and postulating only a single vertex, such that $N_{vtx} = \rho_1 = 1$. A value of z_1^V is then calculated which minimizes F . An iterative cooling procedure then begins. First T is reduced by a ‘cooling factor’ of 0.6. The previous iteration’s vertex positions, z_k^V , are reoptimized to calculate a new minimum value of F . If possible, F is further minimized by replacing any k th vertex with two nearby vertices, such that their weights sum to ρ_k . Thus, as the temperature gets lower, the probability of finding multiple vertices increases.

The cooling procedure is continued until $T < T_{min} = 4$. This cutoff value was chosen to give a reasonable compromise between the algorithm's power to resolve nearby vertices, while not splitting up genuine vertices into multiple smaller ones. At this point, outlier tracks that are more than 4σ from a given vertex are heavily down-weighted by adding a factor of $\exp(-\mu^2/T)$, with $\mu = 4$, to the vertex sum in the denominator of Eq. 3.14. This addition also discourages the formation of new vertices at this point. The cooling process is then continued until $T < 1$. The tracks are then assigned to the remaining collection of k vertices by evaluating p_{ik} and choosing the vertex that has more than a 50% chance to be associated with the track. The output of the procedure is a list of vertices, their z positions, and a list of tracks that are attached to each vertex.

3.3.3 Fitting

The final step of the vertexing process is fitting a full vertex from each track cluster. This is accomplished with an adaptive vertex fitter [175], which operates in a similar fashion to the DA clustering algorithm. It defines a χ^2 with regards to the compatibility of the cluster's tracks with the vertex's 3D position and then slowly lowers an effective temperature to converge towards a fit. The fitter also gives little weight to outlier tracks that are more than a few σ away from the vertex position. The final vertex has a x , y , and z position along with an associated covariance matrix. Fit quality variables such as the χ^2 and number of degrees of freedom are also available.

3.3.4 Resolution

An example of the primary vertex spatial resolutions are given in Fig. 3-13 for 5.02 TeV pp collisions. The resolution in x is shown by the red data points, while the resolution in z is shown by the black points. In general, as the number of tracks in the event increases the vertex fitter is able to more precisely converge on a vertex position. As the number of tracks gets large, the resolutions converge to around $30 \mu\text{m}$ in x and $35 \mu\text{m}$ in z .




Figure 3-13: The primary vertex spatial resolutions in x and z for pp collisions at 5.02 TeV as a function of the number of tracks in the event.

Similar resolutions are shown in the left panel of Fig. 3-14 for PbPb collisions as a function of centrality. As the centrality decreases from 100%, the resolutions quickly fall to values similar to those of pp collisions, around 20 μm . At around 50% centrality the track selection starts constraining the size of the pixel tracking region used to produce the inputs to the vertexing. This is what causes the small wiggle in the resolution. As the centrality is further lowered, the resolution is held relatively constant because the number of tracks used in the vertex fit is also held relatively constant to save on computational time. The slight rising trend seen for central events is caused by the average quality of the input tracks degrading as the detector occupancy grows.

The right panel of Fig. 3-14 shows the resolutions for XeXe collisions, where many aspects of the PbPb algorithm were improved at the cost of computational time. At large centralities, the resolution quickly falls with the increasing number of input tracks. Unlike PbPb collisions, the number of input tracks is not held constant at a certain point and the resolution keeps improving as more tracks are added. The final resolutions in central events are better than 5 μm . This large improvement is caused by the increased number of input tracks and the use of full tracks that have a smaller σ_z than pixel tracks. The pixel detector for the 2017 XeXe data also has an




Figure 3-14: The primary vertex spatial resolutions in x and z as a function of centrality for PbPb and XeXe collisions.

additional layer closer to the interaction point to further improve the resolution.

3.4 Particle Flow and Calorimeter Matching

The CMS detector is composed of many separate high-granularity subdetectors. Many of these components make redundant measurements of the energy and/or momentum of a particle. The redundancy of these measurements can be used to improve the global understanding of an event's properties. Furthermore, some subdetectors have weaknesses that can be mitigated by using information from other components. For example, calorimeters cannot distinguish between two charged particles that are very close to each other and enter the same calorimeter cell. However, the tracker will be able to distinguish these particles from each other. Likewise, it is relatively easy for the tracker to produce a high energy misidentified charged particle from a random combination of hits, but it is not easy to produce a large fake energy deposit in a calorimeter. CMS uses a complex algorithm to assemble relevant information from all subdetectors into an optimized global event description. The algorithm is known as Particle Flow (PF) reconstruction. The PF algorithm took years to develop and has been extensively validated in pp collisions. The discussion in this section is restricted

to the matching of tracks from the tracker to calorimeter deposits in the ECAL and HCAL. This linking is of crucial importance because it gives a reliable data-driven way to verify the validity of reconstructed high- p_T hadrons. A full accounting of PF reconstruction is given in Ref. [176].

3.4.1 Input Candidates

The first step of the algorithm is the reconstruction of individual objects in each subdetector separately. The reconstruction of tracks in the tracker has already been described earlier in this chapter. In the calorimeters, a special clustering algorithm is used to group sets of hits together. The clustering behavior differs depending on the η region of interest, but it suffices to focus on the barrel region for the purposes of this thesis. Firstly, cluster seeds are found by looking for calorimeter cells that have an energy above a seeding threshold. The cell energy must also be above the energies of all 8 neighboring cells. These seeds are then grown into topological clusters by including neighboring cells that have an energy greater than two times the calorimeter's noise threshold. Fig. 3-15 displays a section of HCAL with the cells represented by a set of gray hollow rectangles. The smaller shaded rectangles inside each cell represent the size of the energy deposit it contains. In this example, the darker gray boxes are cluster seeds and the surrounding cells make up the rest of the topological cluster.

The center of each individual cluster is calculated by modeling a topological cluster having N seeds as a superposition of N 2-dimensional Gaussian energy depositions. The width of this Gaussian profile is kept fixed to a value depending on the type of calorimeter being used. A maximum likelihood fit to the energy distribution inside the topological cluster is performed to determine the amplitude and spatial center of each Gaussian. These parameters are then taken to correspond to the energy and position of N separate clusters. It should be noted that the position of a cluster does not have to correspond to the center of a calorimeter cell. In Fig. 3-15 the positions of the two reconstructed clusters are shown by the red dots. This procedure is done separately for both the ECAL and the HCAL.




Figure 3-15: A schematic of a set of two PF HCAL clusters. The dark gray squares are the two cluster seeds, surrounded by smaller cluster deposits. The red dots are the result of the cluster center calculation. The bottom of the green lines are where two tracks enter the HCAL and leave the deposit of energy [176].

3.4.2 Linking

The algorithm then enters a linking step, which attempts to assemble relevant information from various subdetectors into a global event description. First muons, electrons, and isolated photons are identified and their corresponding particle flow elements are removed from consideration for the rest of the algorithm. Next, tracks having a large p_T uncertainty are masked, as their energy is probably measured more accurately in the calorimeters. This also reduces the effect of misreconstructed tracks on the following steps. The remaining tracks are extrapolated from the tracker to the maximum expected distance of an electron's electromagnetic shower inside the ECAL. If the track passes through an ECAL cluster's cell (or a cell that directly neighbors a cluster's cell), it is associated with that cluster. This association is described by a link distance: the distance in the (η, ϕ) plane between the cluster center and the

track's extrapolated endpoint. If an ECAL cluster is associated with more than one track, only the link with the shorter distance is retained.

A similar procedure is carried out in the HCAL. Tracks are extrapolated to a depth of one interaction length. If a *track* is linked to more than one HCAL cluster, only the link with a shorter distance is kept. This is the opposite behavior of the ECAL linking because the HCAL granularity is much poorer. The trajectories of two tracks, represented by green lines, are shown in Fig. 3-15. Each track is linked to the nearest HCAL cluster (red point).

3.4.3 Energy Calibration

After this step, the energy of the calorimeter clusters associated to tracks is recalibrated under the assumption that the particle creating the cluster was a charged hadron. The recalibrated energies are given by

$$E_{ECAL}^{calib} = b(E)f(\eta)E_{ECAL} \quad (3.15)$$

and

$$E_{HCAL}^{calib} = c(E)g(\eta)E_{HCAL}. \quad (3.16)$$

The functions $b(E)$ and $c(E)$ are calculated in simulated events as a function of the true hadron energy, E . These functions can be seen in Fig. 3-16 and are generally close to unity at high energies. The calibration coefficient $c(E)$ changes between two forms depending on if the linked track leaves deposits in just the HCAL (blue line) versus the ECAL and HCAL (green line). The divergent behavior of these functions below 10 GeV occurs because of inefficiencies in the clustering and linking algorithms for charged hadrons having less than 10 GeV of energy. Although the true hadron energy E is not known in data, the functions evolve slowly at high energies and can therefore be evaluated at the reconstructed value of E . The functions $f(\eta)$ and $g(\eta)$ are determined by fitting 2nd order polynomials to the η dependence of the energy response in simulated events.




Figure 3-16: The calibration coefficients used to fine tune the PF cluster energy [176].

Finally, the calibrated cluster energies are compared to the attached tracks' energy. If they are compatible with each other, nothing extra is done. If the cluster energy is in excess of the track's, the calorimeter cluster is split into two. One part of the cluster is still linked to the track, while the excess energy is shunted into another cluster that can be interpreted as a neutral particle. If the total track energy is more than 3σ above the cluster energy, the tracks attached to the cluster are unlinked in descending order of energy. Once the energy excess disappears or all tracks have been unlinked the procedure stops. At this point a collection of tracks with attached calorimeter deposits is now available for analysis.

3.4.4 Calorimeter Matching Performance

The transverse energy of a deposit matched to a track is

$$E_T = \frac{E_{ECAL}^{calib} + E_{HCAL}^{calib}}{\cosh \eta} \quad (3.17)$$

and can be directly compared to a track's p_T to help reject misidentified tracks. A plot of the ratio E_T/p_T for high- p_T tracks in a sample of simulated PbPb PYTHIA+HYDJET is shown as a function of centrality in the left panel of Fig. 3-17. Most of the tracks have a ratio near unity, indicating a successful match to the calorimeters. There is another band tracks around $E_T/p_T = 0$ corresponding to fake tracks that have no associated energy deposit. To reject these tracks a cut requiring $E_T/p_T > 0.5$, shown by the red line on Fig. 3-17, can be applied. The right panel on the same figure shows the relative efficiency loss for genuine tracks that is caused by the application of this cut. For low- p_T tracks the calorimeter matching algorithm does not operate reliably and a large amount of real tracks are removed. However, above ≈ 10 GeV the matching is very efficient, with only a 2% relative efficiency loss. This corresponds roughly to the same energy range that the energy calibrations shown in Fig. 3-16 diverge. To be conservative the calorimeter matching requirement is only applied for tracks having a measured $p_T > 20$ GeV.

3.5 Jets

Jets are experimental proxies for initially hard-scattered partons and their resulting parton showers. However, a jet's kinematics will not always match the parton's kinematics exactly; some of its daughter particles could not be included in the jet, or extra particles could be erroneously added into the jet. This means that the properties of a measured jet depend on exact the procedure used to cluster the particles together in the first place. Any such jet definition should have a few desirable properties. Firstly, it should be relatively simple and easy to implement. Secondly, it should not be too computationally expensive to allow processing large amounts of experimental data.




Figure 3-17: (Right) The distribution of E_T/p_T as a function of centrality in PbPb PYTHIA+HYDJET. (Left) The relative efficiency from applying a calorimeter-matching required as a function of the track p_T in the same sample .

Thirdly and most importantly, it should be theoretically tractable to allow comparisons between experiment and theory. As discussed in Section 1.2.3, the probability of gluon emission in a parton shower diverges as the gluon’s energy goes to zero (infrared limit), or as its emission angle goes to zero (collinear limit). This means any reasonable jet definition should not be sensitive to the addition of a very soft gluon or a very collimated radiation. Observables fulfilling these two criteria are known as infrared and collinear (IRC) safe.

A class of jet-finding algorithms known as sequential combination algorithms fulfill the three requirements above. They define a distance measure d_{ij} between all the constituents to be clustered:

$$d_{ij} = d_{ji} = \min(p_{T_i}^{2p}, p_{T_j}^{2p}) \frac{\Delta R_{ij}^2}{R^2}. \quad (3.18)$$

Here $\Delta R_{ij}^2 = (\eta_i - \eta_j)^2 + (\phi_i - \phi_j)^2$, and R is a distance parameters (or jet radius) that determines roughly how large the jet’s angular opening is. The parameter p can be chosen to emphasize different particle properties when clustering. The most popular choice is $p = -1$, which produces ‘anti- k_t ’ jets [177]. These jets tend to have their

hard constituents clustered together first. They also tend to be roughly circular in shape, making ‘anti- k_t ’ jets fairly intuitive to understand from an experimental point of view. A distance to the beam line is also defined:

$$d_{iB} = p_{Ti}^{2p}. \quad (3.19)$$

During clustering, the smallest distance measure is calculated for all the pairs of constituents. If it is a d_{ij} , the two constituents have their 4-momenta added together and are re-added to the list of constituents as one combined constituent. If the distance measure is a d_{iB} , then constituent i is removed from the list of constituents and is added to the list of jets produced. New distance measures are computed, and the procedure is iterated until no constituents remain. This can produce a large number of very low- p_T jets, but typically only jets above some given threshold are used for final analysis. This algorithmic procedure is efficiently implemented in the FASTJET software package [178].

A visualization of the output of an anti- k_t algorithm is shown in Fig. 3-18. The algorithm tends to favor producing roughly circular jets that emphasize large energy deposits at their center. The jets are not always circular, however.

In PbPb collisions, jets are created using energy deposits in the ECAL and HCAL. The jets are clustered using an anti- k_t algorithm with $R = 0.4$. Because there is a significant amount of energy coming from the underlying event instead of the hard scattering in the event, the jets’ final p_T is modified using a ‘pileup subtraction’ algorithm [179]. (This technique was initially developed for mitigating the effects of pileup in pp events, but also naturally works for the large backgrounds in heavy ion events.) This subtraction works by assuming that the underlying event contributes a uniform energy density to the jet. In CMS, this subtraction is implemented by first calculating the average transverse energy E_T in strips of η along the calorimeters. The average transverse energy plus one standard deviation $E_T + \sigma(E_T)$ is subtracted from all the calorimeter towers. Negative towers are then zeroed out. At this point a jet clustering is run on the remaining nonzero towers to produce a set of jets. How-




Figure 3-18: The output of an anti- k_t jet clustering algorithm on a toy event [177].

ever, these jets are necessarily over-subtracted because the jets themselves increase the average E_T in the event. To compensate for this, the entire procedure is repeated after masking the areas around the previously-found jets. This alleviates the auto-correlation between the presence of a jet and the calculation of the underlying event level. The output of this second clustering is the final jet collection. More details of the procedure can be found in references [28, 180, 181].

Using simulated PYTHIA+HYDJET samples, it was checked that this procedure gives a correct estimation of the true jet momentum on average, albeit with some finite resolution. The PbPb jet momenta are used here for combining data from different triggers, but the final result is fairly insensitive to the exact momentum calibration.

The jets used in 2.76 TeV pp and 5.02 TeV pPb data are clustered with the anti- k_t algorithm using $R = 0.3$. Instead of using simple calorimeter deposits, CMS particle

flow candidates are clustered. Having a precise jet momentum calibration is important for the measurement of the jet fragmentation function in a given p_T bin, so a rigorous set of data-driven jet energy corrections is used. These corrections are derived by examining dijet and γ -jet events. The details are explained in CMS studies of the dijet balance [42] and dijet eta distribution [150] in pPb events.

The jets used for the 7 TeV pp analysis are $R = 0.3$ anti- k_t particle flow jets with pileup subtraction. The procedure for correcting their energy is detailed in Ref [182].

Chapter 4

Data and Monte Carlo Samples

Conducting an analysis at the LHC involves processing many massive data samples. In general, these samples can be categorized into 2 classes: data and Monte Carlo (MC). Data are simply the real information streaming from the detector. MC samples are simulations that are designed to try to reproduce some or many aspects of the data. MC is usually used for calibrations, detector performance, and feasibility studies. Some specific MC samples are of special use to heavy ion physics. They implement a Glauber model of nuclear collisions [183] (described in Section 4.2), allowing a connection to be made between the final observed particles and the initial geometric orientation of the colliding nuclei.

4.1 Monte Carlo Generators

A common question when reconstructing events is ‘how do I know my algorithm is working as intended?’ To answer this question, computer simulations that mimic the response of a particle detector to a known list of particles are used. Output of the reconstruction can be compared to the input particle list to see if things are working as expected. In practice, reconstruction algorithms are not perfect and these comparisons are crucially important for understanding the efficiency and resolution of the final physics objects produced.

Although any list of particles can be used as an input into a detector simulation,

it is desirable to have something that looks very similar to real data. To generate these lists, physicists use theoretically-motivated MC techniques. These MC generators typically have multiple free parameters which can be tuned to give a good description of data. Each MC generator has its own strengths, weaknesses, and scope of applicability. For example, some generators only attempt to describe pp collisions, while others allow the generation of both pp and heavy ion events. Furthermore, the models of different generators can be combined in a process called ‘embedding.’ In this technique, the list of particles from an event created with one generator are simply appended to the list from another generator. This is useful for creating events having a hard pp-like scattering coming from a heavy ion background.

The major MC generators used in this work are given below:

- **PYTHIA:** This generator is the major workhorse of modern high-energy physics. It is able to simulate the full pp collision cross section across a wide range of energies. It also allows users to request a specific hard process or class of processes (i.e. jet production) to easily create a sample of rare events. After generating a hard scattering, a parton undergoes fragmentation to create a parton shower. In PYTHIA this calculation is done using Sudakov form factors [184] and results in a parton shower that is ordered in p_T . The resulting partons are then turned into colorless hadrons using the Lund string hadronization model [185]. PYTHIA’s physics model also includes a sophisticated handling of multiple parton interactions, which are crucial for describing the soft component of pp collisions. Simulating both initial and final state radiation, as well as color reconnection of outgoing partons with the beam remnants, is also supported. This generator tends to reproduce pp collisions quite well but is not used to simulate heavy ion collisions. Some criticisms of PYTHIA include its failure to reproduce azimuthal anisotropies in high-multiplicity pp collisions and an underprediction of strange hadron yields. Here we use PYTHIA 6 [186], tune Z2 [187] and PYTHIA 8 [188], tune CUETP8M1 [189]. It should be emphasized that, while PYTHIA 6 and PYTHIA 8 have very similar physics models, they are not the exact same and do not give precisely the same results.

- **HERWIG++:** This is a major competitor to PYTHIA as a general-purpose pp collision MC. Unlike PYTHIA, it produces parton showers that are angular-ordered (as opposed to p_T ordered) because of its treatment of color coherence effects. Additionally, a cluster hadronization model is used instead of string hadronization [190]. In this work, we use tune EE5C [189] as a cross-check of the results obtained using PYTHIA.
- **HIJING:** In pPb collisions, the incoming proton can collide with multiple nucleons in the lead ion. HIJING simulates this with a model that modifies the scattered parton's kinematics after each additional scattering. This results in the production of a significant amount of minijets (defined as jets that are too low in energy to be experimentally resolved, but still energetic enough to be calculated with perturbative QCD techniques). Glauber-based initial nucleon positions and nuclear shadowing effects are also included in the model. Much of the framework for calculating the results of each individual scattering uses old PYTHIA 5 kernels and hadronization is done with a similar Lund string model [191].
- **HYDJET:** Nuclear collisions are modeled in HYDJET as a hydrodynamically evolving soft background that expands until it reaches a given freeze out temperature. Hard scatterings are also present, but the resulting partons are quenched using a modified version of PYTHIA known as PYQUEN [192]. The minimum hard scattering energy scale is set to zero, so that high-energy jet production in MB events is relatively rare. To model high-energy processes, PYTHIA signal events are typically embedded into HYDJET background events (these samples are called PYTHIA+HYDJET). Internal CMS tunes that have been tuned to PbPb data are used.
- **EPOS:** This generator uses a two-part model where a high energy-density 'core' follows a hydrodynamic evolution until a specified freeze out temperature, while a low energy-density 'corona' evolves in a vacuum-like manner [193]. In this way, it naturally scales between vacuum-like pp collisions, which contain

mostly corona, and the hydrodynamic behavior of heavy ion collisions that have a large core. It does a much better job of reproducing strange hadron yields than HYDJET. The LHC tune is used here [194].

- AMPT: This model uses HIJING to generate the initial conditions for an ion collision. Parton scatterings in the bulk of the medium are then modeled using pQCD scattering cross sections after including screening masses. Hadronization is done with the Lund string model, but this can be changed to a quark coalescence model if a ‘string melting’ setting is activated. Finally, the resulting hadrons are allowed to undergo rescattering. This model has been able to qualitatively reproduce strange hadron spectra and hydrodynamic flow measurements at RHIC and the LHC [195].
- STARLIGHT: At the LHC extremely strong fluxes of virtual photons are created in the regions directly surrounding around the accelerated particles. This is especially true for heavy ion collisions, where the probability of a nucleus interacting with a photon scales with the fourth power of the atomic number ($82^4 \approx 5 \times 10^7$ for lead). Because of this, the cross section for electromagnetic interactions becomes orders of magnitude larger than the hadronic cross section [196]. Unless one is specifically studying electromagnetic processes, typically only hadronic events are used for a measurement. In this work, only the hadronic cross section is examined. However, it is difficult to distinguish some of these electromagnetic events from very peripheral hadronic events, so a model is required to estimate what each of these event classes look like in a detector. Their contribution to the final result can then be corrected for. STARLIGHT [197] produces these electromagnetic processes and allows for a calculation of the selection efficiency for hadronic events. Version 2.2 is used here.

Particles coming from these MC generators are fed into a GEANT 4 [198] simulation of the CMS detector. This simulation calculates where and how much energy is deposited into CMS’s subdetectors and outputs simulated electronic signals that look

similar to real data. From this point onward, the data and MC samples can be treated as nearly being the same for the purposes of studying reconstruction performance.

4.2 Glauber Model

Heavy nuclei are simply bigger than protons and therefore have larger production cross sections for various particles. When examining hard processes, it makes sense to consider the nucleus as a collection of individual nucleons. Furthermore, a nucleus is a non-uniform quantum object, and the spatial positions of individual nucleons within the nucleus are not static. This causes variations in the initial geometry of a collision, even when b is held constant. The cross section of a hard process will roughly scale with the number of binary nucleon-nucleon collisions (N_{coll}) that occur for any given geometric state. To compare heavy ion and pp reference data sets fairly, this effect must be accounted for by dividing the heavy ion spectra by N_{coll} , as seen in the definition of R_{AA} (Equation 1.7). As N_{coll} is not known directly, it is typically calculated using a simple model known as the Monte Carlo Glauber model [183]. The model is implemented in the TGLAUBERM C v3.2 [199, 200, 201] software package. The Glauber model is an essential tool for understanding the initial geometry of heavy ion collisions, and how a collision's centrality related to other geometric quantities.

In this model, a nucleus is constructed by randomly generating an appropriate number of nucleons for the given nuclear species according to a spatial distribution ρ given by a modified Fermi distribution:

$$\rho(r, \theta) = \rho_0 \frac{1}{1 + \exp \frac{r - R(1 + \beta_2 Y_{20})}{a}}. \quad (4.1)$$

Here r is the radial coordinate, θ is the polar angle, and ρ_0 is a normalization factor that is irrelevant for the purposes of the Glauber calculation. The nuclear radius (R), skin depth (a), and deformation parameter (β_2) describe the shape of the nucleus. These input parameters are constrained by low energy measurements of nuclear

structure. The spherical harmonic Y_{20} is:

$$Y_{20} = \sqrt{\frac{5}{16\pi}}(3\cos^2(\theta) - 1). \quad (4.2)$$

When building a nucleus, if the centers of two nucleons are within a distance d_{min} of each other, the position of one of the nucleons is regenerated. This prevents two nucleons being inside of one another, simulating the effects of a hard-core repulsion between nucleons.

An impact parameter b is then randomly chosen out of the distribution $dN/db \propto b$, up to a maximum value of $b = 20$ fm. The centers of the incoming nuclei are then positioned at $x = \pm b/2$, and are assumed to be moving along the z direction. Nucleons in the two nuclei are assumed to collide with each other if the distance between their centers is less than the “ball diameter” $D = \sqrt{\sigma_{NN}^{inel}/\pi}$. Here σ_{NN}^{inel} is the nucleon-nucleon inelastic cross section at the appropriate center of mass energy. It is also an input parameter that is constrained by measurements of the proton-proton inelastic cross section.

The relevant input parameters for Glauber simulation of pPb, PbPb, and XeXe collisions are given in Table 4.1. Some of the parameters for lead ions are different between pPb and PbPb collisions because of improvements to the model that occurred between the pPb and PbPb analysis. The value for σ_{NN}^{inel} in XeXe collisions is lower than that of PbPb collisions, despite the latter having a lower center of mass energy. This is because improved fits describing the $\sqrt{s_{NN}}$ dependence of σ_{NN}^{inel} were available for the XeXe analysis. Finally, it should be noted that β_2 for lead is negligible because it is a spherical double-magic nucleus [202].

Visual examples of the results of the MC Glauber procedure are shown in Fig. 4-1. In this figure the nucleons not participating in the collision are represented as yellow and orange circles, depending on which nucleus they belong to. Nucleons hitting at least one other nucleon are colored pink. All nucleons are drawn with a radius of $D/2$. The average nuclear radius R of each nucleus are shown by the dotted circles circumscribing most of the nucleons. The top left panel shows a PbPb collision of

Table 4.1: The Glauber model settings used for modeling the heavy nuclei in pPb, PbPb, and XeXe collisions.

| Parameter | Units | pPb | PbPb | XeXe |
|----------------------|-------|-------------------------|-------------------------|-----------------------|
| $\sqrt{s_{NN}}$ | TeV | 5.02 | 5.02 | 5.44 |
| R | fm | 6.62 ± 0.13 [199] | 6.62 ± 0.06 [203] | 5.36 ± 0.1 [200] |
| a | fm | 0.546 ± 0.055 [199] | 0.546 ± 0.010 [203] | 0.59 ± 0.07 [200] |
| d_{min} | fm | 0.4 ± 0.4 [199] | 0.4 ± 0.4 [199] | 0.4 ± 0.4 [199] |
| σ_{NN}^{inel} | mb | 70 ± 5 [199] | 70 ± 5 [199] | 68.4 ± 0.5 [201] |
| β_2 | | — | — | 0.18 ± 0.02 [200] |

around 50% centrality. The top right panel is a XeXe collision. Even though the xenon nuclei are smaller, the number of colliding nucleons is similar to that of the PbPb collision because the centrality is slightly more central. The bottom left panel shows a pPb collision, where a single proton punches through the center of a lead nucleus. Finally, the bottom right panel shows a very peripheral PbPb collision, illustrating that nuclear collisions should be very similar to proton proton collisions in the limit that b is large enough that only one nucleon from each nucleus collides.

Many events can be generated quickly using this MC procedure, allowing for the calculation of geometric quantities. One of the simplest quantities is the number of nucleons that participate in the collision (the pink nucleons in Fig. 4-1). This quantity, N_{part} , is a good representation of the size of the overlapping region of the collision and is roughly proportional to the total collision multiplicity across a wide range of center of mass energies [204, 205].

Because two nucleons can have similar x and y positions if their z coordinate is different, it is possible for one incoming nucleon to strike multiple nucleons in the other nucleus. A large total number of collisions, N_{coll} , between nucleons increases the probability of a hard process happening in the event. Thus, the production of high- p_T objects such as jets scales with N_{coll} . This quantity is part of the definition of the ‘nuclear overlap function’ or ‘nuclear thickness function’, defined in Equation 1.8. The nuclear thickness function has units of inverse area and describes the ‘effective’ proton-proton luminosity represented by one nucleon-nucleon event [183]. Although observables can be represented using either N_{coll} or T_{AA} , it is generally preferable to use the latter because it is less sensitive to uncertainties in σ_{NN}^{inel} .




Figure 4-1: Results from Glauber MC simulations, showing the initial state geometry for different collision systems. Yellow and orange circles represent spectator nucleons from each nucleus, while pink circles are wounded nucleons. The collision systems are PbPb (Top left), XeXe (Top right), pPb (Bottom left), and PbPb (Bottom right). Produced with TGLAUBERM C v3.2 [199, 200, 201].

Finally, the average path length of a hard parton traveling through a QGP produced by the collision can be calculated. This is done by selecting a random nucleon undergoing a collision and drawing a ray in a random direction in azimuth. The total length of this ray that passes through any nucleons is taken as the average path length [201]. This procedure is very simple, ignoring QGP expansion and both transverse and longitudinal dynamics. However, it is still useful when comparing the energy loss of partons originating from collisions of differing size, such as PbPb and XeXe collisions.

The results obtained from this Glauber model are calculated in bins of centrality calibrated using b . To account for bin-to-bin centrality fluctuations caused by the HF detector resolution, the results are smeared using simulated EPOS events [183]. This produces quantities for centrality bins that are directly comparable to what is

measured in data.

The average values of N_{part} , N_{coll} , and T_{AA} for pPb, XeXe, and PbPb collisions are given in Table 4.2. The event selection efficiency for pPb and PbPb collisions is very high, so its uncertainty has been folded into the quantities reported here. However, the uncertainty on the XeXe event selection efficiency is much larger. It is not included in the quantities reported in this table and therefore must be treated as its own source of uncertainty in subsequent analysis.

Table 4.2: The values of $\langle N_{\text{part}} \rangle$, $\langle N_{\text{coll}} \rangle$, T_{AA} , and their uncertainties, for $\sqrt{s_{NN}} = 5.44$ TeV XeXe collisions and 5.02 TeV PbPb collisions in the centrality ranges used. The relevant values for inclusive PbPb at 5.02 TeV collisions are also given.

| Centrality | $\langle N_{\text{part}} \rangle$ | | $\langle N_{\text{coll}} \rangle$ | | T_{AA} [mb $^{-1}$] | |
|------------|-----------------------------------|------------------------|-----------------------------------|-----------------------|------------------------|----------------------------|
| | XeXe | PbPb | XeXe | PbPb | XeXe | PbPb |
| 0-5% | 236.1 \pm 1.3 | 384.3 $^{+1.8}_{-2.0}$ | 930 \pm 51 | 1820 $^{+130}_{-140}$ | 13.60 \pm 0.74 | 26.0 $^{+0.5}_{-0.8}$ |
| 5-10% | 206.3 \pm 1.7 | 333.3 $^{+3.0}_{-3.2}$ | 732 \pm 44 | 1430 $^{+100}_{-110}$ | 10.70 \pm 0.65 | 20.5 $^{+0.4}_{-0.6}$ |
| 10-30% | 141.2 \pm 1.8 | 226.7 $^{+5.2}_{-5.3}$ | 407 \pm 30 | 805 $^{+55}_{-58}$ | 5.94 \pm 0.44 | 11.5 $^{+0.3}_{-0.4}$ |
| 30-50% | 68.5 \pm 2.2 | 109.2 $^{+4.3}_{-4.2}$ | 135 \pm 15 | 267 $^{+20}_{-20}$ | 1.97 \pm 0.22 | 3.82 $^{+0.21}_{-0.21}$ |
| 50-70% | 27.2 \pm 1.6 | 42.2 $^{+3.0}_{-2.9}$ | 35.3 \pm 4.8 | 65.4 $^{+7.0}_{-6.6}$ | 0.517 \pm 0.071 | 0.934 $^{+0.096}_{-0.089}$ |
| 70-80% | 10.55 \pm 0.78 | — | 9.8 \pm 1.4 | — | 0.143 \pm 0.020 | — |
| 70-90% | — | 11.1 $^{+1.3}_{-1.2}$ | — | 10.7 $^{+1.7}_{-1.5}$ | — | 0.152 $^{+0.024}_{-0.021}$ |
| 0-10% | 221.2 \pm 1.5 | 358.8 $^{+2.4}_{-2.6}$ | 831 \pm 47 | 1630 $^{+120}_{-120}$ | 12.10 \pm 0.69 | 23.2 $^{+0.4}_{-0.7}$ |
| PbPb | — | — | 6.9 \pm 0.3 | — | 0.0983 \pm 0.0044 | — |

Table 4.3: The data samples used in this thesis.

| year | System | $\sqrt{s_{NN}}$ (TeV) | $\langle\mu\rangle$ | \mathcal{L}^{int} |
|------|--------|-----------------------|---------------------|----------------------------|
| 2011 | pp | 7.00 | ≈ 6 | 2.5 fb^{-1} |
| 2013 | pp | 2.76 | 10^{-5} | 5.3 pb^{-1} |
| 2013 | pPb | 5.02 | <0.05 | 35 nb^{-1} [206] |
| 2015 | pp | 5.02 | 0.9 | 27.4 pb^{-1} |
| 2015 | PbPb | 5.02 | <0.003 | $404 \mu\text{b}^{-1}$ |
| 2017 | XeXe | 5.44 | <0.02 | $3.42 \mu\text{b}^{-1}$ |

4.3 Data Samples

It should be kept in mind that the various data sets used in this thesis were gathered concurrently with some analyses activities. Therefore, some analysis decisions were motivated by the lack of an appropriate data set at the time. The data sets used in this work are listed in Table 4.3.

To build a pp reference spectrum for the pPb fragmentation function analysis, 7 TeV data from the first half of 2011 is used. These data were selected because they have a lower average pileup, $\langle\mu\rangle$, than the data gathered later in 2011. Data from 2013, taken at 2.76 TeV are also used for the reference, but this was a special reference run so the pileup is negligible. The pPb data used were also gathered in 2013. It should be noted that the pPb fragmentation function analysis was conducted before the collection of an appropriate pp data set at 5.02 TeV.

The pp reference data at 5.02 TeV from 2015 were gathered directly before the PbPb run, ensuring that detector conditions were very similar between the two data sets. These are the data used in the analysis of the PbPb nuclear modification factor.

Finally, a short test run of XeXe collisions at 5.44 TeV was conducted in 2017. Because of the low interaction rate, all the data were gathered with a single MB trigger, resulting in the collection of around 19 million events. Because no pp run at 5.44 TeV is foreseen, a pp reference at the appropriate energy must be constructed to build the XeXe nuclear modification factor.

The amount of data collected is quantified using the integrated luminosity, \mathcal{L}^{int} , of the sample. This is defined as the ratio of the number of events of a given process,

divided by the cross section for that process:

$$\mathcal{L}^{int} = \frac{N^{process}}{\sigma^{process}}. \quad (4.3)$$

The instantaneous luminosity \mathcal{L} is calibrated in CMS using the Van der Meer (VdM) scan technique [207, 206] where the rate of a process is monitored as the LHC beams are scanned across each other. A large amount of 7 TeV pp data is available because these data are also used for high-priority high energy physics analyses. Less pp data were taken at 2.76 and 5.02 TeV because these data were collected during special runs requested for the heavy ion program. The pPb and PbPb data sets were both the largest data sets of their type ever collected. The XeXe data were gathered over the course of only a few hours, so the amount of luminosity available for this collision system is much smaller than what is available for PbPb.

A data set's average pileup, $\langle\mu\rangle$, is the average number of independent hadronic interactions that occur every time the LHC beams cross each other. The chance of one particle in the beam interacting is independent of the chances of another particle colliding, so the probability P of having k interactions in a given crossing is given by a Poisson distribution with mean $\langle\mu\rangle$:

$$P(k \text{ collisions}) = \frac{\langle\mu\rangle^k e^{-\langle\mu\rangle}}{k!}. \quad (4.4)$$

Most of the data used in these analyses — including all of the heavy ion data — has a very low $\langle\mu\rangle$ so pileup effects can be neglected. However, the 2011 and 2015 pp data have non-negligible pileup. In triggered data where it is known that $k \neq 0$, the probability of having at least 2 interactions is $P(k \geq 2|k \neq 0) = 0.383$ and $P(k \geq 2|k \neq 0) = 0.985$ for 2015 and 2011 data, respectively. Clearly for these data the pileup must be taken into account.

Unlike all other collision systems used here, pPb collisions have a center of mass reference frame which differs from the laboratory reference frame. This is because the LHC requires both particle beams to have the same rigidity. Thus, the proton

beam has an energy of nearly 4 TeV, while the lead nuclei have an energy of 1.58 TeV per nucleon. This leads to a rapidity boost of 0.465 in the laboratory frame. To compensate for this, pseudorapidities for pPb collisions are reported as $\eta_{CM} = \eta_{lab} - 0.465$. For pp, XeXe, and PbPb collisions $\eta_{CM} = \eta_{lab}$.

Because pPb collisions are an asymmetric system, the directions of the two beams were switched approximately halfway through the data collection period. This allows a comparison between the two directions to investigate systematic uncertainties. The sign of the η direction is therefore swapped for the reversed direction data, allowing both beam directions to be combined into a single result.

Data are only used if they successfully pass a CMS validation procedure that confirms all subdetectors were operational and behaving as expected. This validation is done in chunks of approximately 23 second intervals. The identification numbers of all intervals passing the validation are recorded into a JavaScript Object Notation (JSON) file that is used to filter out bad data blocks.

4.4 Event selection

Not all events recorded by CMS are legitimate hadronic collisions. Sometimes detector noise can cause a trigger to occur during an event that is completely devoid of a collision. Interactions of the incident particle beams with the trace amounts of residual gas around the interaction region can also cause background events. Additionally, the strong photon flux generated by beams of heavy nuclei cause a huge amount of electromagnetically-induced interactions. Most of these interactions lead to particles being generated in only one side of the detector and can be rejected by requiring activity in both sides. However, the probability of two separate interactions occurring at the same time, leading to activity in both sides of the detector, is still non-negligible. These events can be difficult to distinguish from peripheral hadronic collisions. Furthermore, very peripheral collisions can sometimes not generate enough activity in the detector to cause a trigger, leading to a centrality-dependent inefficiency. Thus, before proceeding to more complicated analyses, the data samples are

first cleaned to minimize the effects of electromagnetic contamination, while trying to keep the hadronic efficiency as large as possible.

The trigger conditions for MB data are the first set of basic requirements an event must pass to be considered (because it is not recorded otherwise). For 5.02 TeV pp data, one HF deposit above approximately 1 GeV was required to fire the trigger. In pPb collisions, a single pixel track having $p_T > 0.4$ GeV is required. A single pixel track, as well as a deposit of at least 1 GeV in the HF was required for XeXe collisions. PbPb collisions are required to have a deposit of at least 1 GeV in both sides of the HF detector.

For offline cleaning, all collision systems are first required to have at least one primary vertex within 15 cm of the center of the detector. In pp, pPb, and XeXe collisions, at least 25% of the tracks in the event are required to pass the *high purity* selection [157]. For PbPb collisions the pixel cluster shapes must be compatible with those expected for a single PbPb collision. These requirements help reduce beam-related backgrounds. In pPb collisions, at least one detector element in each of the two HF detectors is required to have a hit of greater than 3 GeV. For XeXe and PbPb collisions, three detector elements above this threshold are required. These selections help reduce electromagnetic contamination.

After these selections, the hadronic efficiency and electromagnetic contamination can be estimated by comparing the data to MC such as HYDJET, EPOS, and STARLIGHT. The pp data used in this work is fully efficient for the kinematic ranges used. The total hadronic efficiency plus any remaining electromagnetic contamination is $(99 \pm 1)\%$ in pPb collisions [126], $(99 \pm 2)\%$ in PbPb collisions [208], and $(95 \pm 3)\%$ in XeXe collisions [173]. The lower efficiency for the XeXe data could be related to the HF detector having experienced radiation damage after the significant amount of high-luminosity pp data gathered in 2016 and 2017. Also, because Xe nuclei are smaller than Pb, a greater fraction of the cross section will have a total HF activity that is below the detector's noise threshold.

4.5 Triggers

High- p_T charged hadrons are contained inside of high- p_T jets by definition. Triggers recording jets can therefore also be used to examine charged particle production. With a large enough p_T threshold, the jet trigger can be run without a prescale, meaning that every jet above a certain energy will be recorded. This is incredibly useful, as high- p_T objects are relatively rare but also probe interesting physics.

The 2.76 TeV pp and 5.02 TeV pPb samples used jet triggers requiring a least one $R = 0.5$ anti- k_t calorimeter jet with p_T above either 40 or 80 GeV. These triggers are referred to as Jet40 and Jet80 triggers, respectively. Turn-on curves showing the efficiency of the trigger as a function of the offline jet p_T in 2.76 TeV pp collisions are shown in Fig. 4-2. The left panel shows the turn-on evaluated in a limited MB data. The right panel shows the turn-on in PYTHIA, where the statistical errors are negligible. In both cases, the triggers are fully efficiency at a p_T that is 20 GeV above the threshold given in the trigger name. Similar turn-on curves are given for 5.02 TeV pPb collisions in Fig. 4-3. The top left panel shows the 40 GeV jet trigger efficiency, as evaluated in data. The entries have been weighted by the average prescale of the trigger, meaning that the data points can go above unity because of statistical fluctuations. The top right panel shows the 80 GeV trigger. The bottom middle panels displays the same turn-on curves, as they are evaluated in MC. Once again, the triggers are fully efficient about 20 GeV above the threshold listed in the trigger name.

The 7 TeV pp sample uses a similar trigger as the 2.76 TeV pp data, with thresholds at 30, 60, and 110 GeV. Trigger efficiencies evaluated in MC are shown in Fig. 4-4. The triggers are fully efficient at 40, 80, and 140 GeV, respectively. It is difficult to evaluate the efficiencies in data because of the lack of an appropriate MB data set for this collision system.

The 5.02 TeV PbPb data had jet triggers searching for $R = 0.4$ anti- k_t pileup-subtracted calorimeter jets above thresholds of 40, 60, 80, and 100 GeV. An analysis-level cross-check also uses data gathered with a trigger that ran a fast version of




Figure 4-2: Turn-on curves for 2.76 TeV pp jet triggers.

the offline tracking in $|\eta| < 1.05$ to find single tracks above thresholds of 12, 18, 24, and 34 GeV. Copies of these triggers with the additional requirement that the event originates from a peripheral PbPb event were implemented to boost the statistical power of the data in peripheral events. The total trigger efficiency for the jet triggers is shown in the top left panel of Fig. 4-5. They are fully efficient at a p_T that is around 20 GeV greater than the threshold in the trigger name. The top right panel of the same figure shows the efficiency of the track trigger as a function of the offline highest track p_T . The turn-on is much tighter. Finally, the bottom middle panel displays the efficiency as a function of centrality for the peripheral-gating used for the peripheral versions of these triggers. The gating is fully efficient before the targeted centrality threshold of each trigger.

The 5.02 TeV pp data set has triggers similar to the PbPb data. The jet trigger looked for $R = 0.4$ anti- k_t calorimeter jets above 40, 60, and 80 GeV. The track trigger was run in the pseudorapidity range $|\eta| < 2.4$ and had p_T requirements of 12, 24, 34, 45, and 53 GeV. The trigger performance of these triggers is similar to, or even better than the PbPb triggers.

For all collision systems, a large amount of MB data is also used. In the case of XeXe collisions, all collision events were gathered with a MB trigger. The various MB trigger criterion have already been described previously as part of the global event selection.

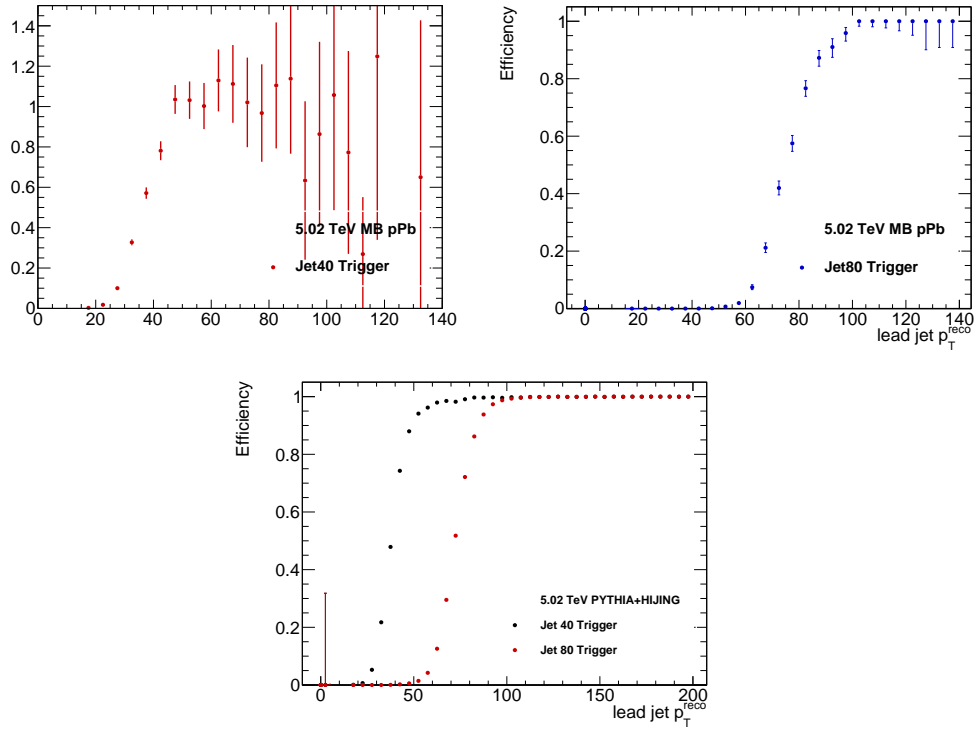


Figure 4-3: Turn-on curves for 5.02 TeV pPb jet triggers.




Figure 4-4: Turn-on curves for 7 TeV pp jet triggers.

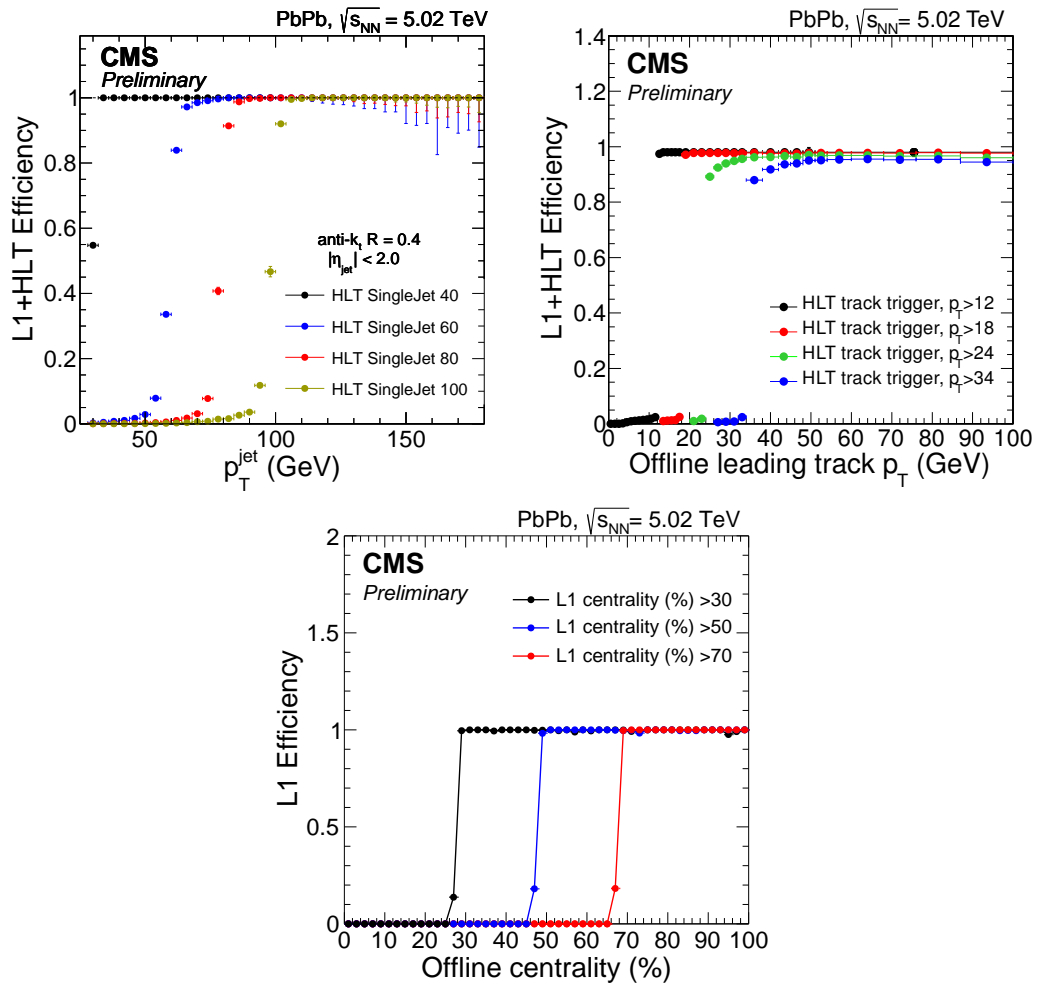


Figure 4-5: Turn-on curves for 5.02 TeV PbPb jet and track triggers, as well as the peripheral centrality gating.

Chapter 5

pPb Fragmentation Function

This chapter focuses on a measurement of the experimental fragmentation function of jets in 5.02 TeV pPb collisions [209]. The experimental fragmentation function is defined as the per-jet differential p_T spectrum of all primary charged particles contained within a jet having a given p_T^{jet} . The temporal context of this analysis is important to remember; it was done in early 2015 before any 5.02 TeV pp reference data had been gathered. At this time the scientific community was struggling to reconcile measurements of jet and charged particle R_{pPb} . This analysis serves as a cross check of the assumptions used for those measurements.

5.1 Analysis Procedure

In the analysis three data sets are used: 2.76 TeV pp data from 2013, 7 TeV pp data from 2011, and 5.02 TeV pPb data from 2013. The pp data sets are used to construct an interpolated reference, as described in Section 5.2. For all three samples, the fragmentation function is measured in five ranges of jet transverse momentum (p_T^{jet}). These ranges are 60–80, 80–100, 100–120, 120–140, and 140–200 GeV.

5.1.1 Jet Selection

In the 2.76 TeV pp and 5.02 TeV pPb data, only events collected with a Jet40 trigger (described in Section 4.5) are used in the range $60 < p_T^{\text{jet}} < 100$. A Jet80 trigger is used for the higher p_T^{jet} ranges. For the 7 TeV pp data, the lowest p_T^{jet} bin is filled using a Jet30 trigger, and the highest p_T^{jet} bin uses a Jet110 trigger. The range between $80 < p_T^{\text{jet}} < 140$ is covered by a Jet60 trigger. These choices ensure the triggers are fully efficient in the jet kinematic range being studied.

First $R = 0.3$ anti- k_t particle flow jets in the range $|\eta_{\text{CM}}^{\text{jet}}| < 1.5$ are selected. This distance parameter was chosen to reduce fluctuations caused by the presence of the underlying event in pPb collisions. The chosen $\eta_{\text{CM}}^{\text{jet}}$ range guarantees that the reconstructed jets are compatible with the acceptance of the tracker after accounting for the jet radius and rapidity boost in pPb collisions. The jets must have an uncalibrated raw energy above 30 GeV. Additionally, between 5% and 95% of their energy must be attributed to charged particles. This selection removes completely neutral jets that could result from calorimeter noise, as well as jets consisting of a single high- p_T particle. The jets have their energy corrected using the procedures outlined in Refs [42, 150].

5.1.2 Track Selection

The measurements presented in this work are done for primary charged particles. These are charged particles having a mean proper lifetime that is larger than 1 cm. Daughter particles from secondary decays are not primary, unless the mother particle had a lifetime that was less than 1 cm. Particles that are produced from interactions with detector material are also not considered primary particles. To enhance the fraction of charged particles that are primary, tracks are required to have a DCA significance for d_0 and z that is less than 3σ . Tracks with poor momentum resolution are removed by requiring the track fit's relative uncertainty in p_T to be less than 10%. The standard *high purity* selections are also applied. Tracks are not explicitly required to match a calorimeter deposit; this is an indirect requirement because only

tracks inside a jet are considered for this analysis. The analysis does not consider tracks below $p_T = 0.5$ GeV. A summary of the track selections used can be found in the first three rows of Table 5.1.

5.1.3 Tracking Correction

For every jet identified, the tracks passing all selections and having

$$\Delta R = \sqrt{(\eta^{\text{jet}} - \eta^{\text{track}})^2 + (\phi^{\text{jet}} - \phi^{\text{track}})^2} < 0.3 \quad (5.1)$$

are tabulated as a function of p_T^{track} . Because anti- k_T jets are not necessarily perfectly circular, it is possible (but rare) that this definition includes particles that are not clustered into the jet, or excludes particles that were clustered. While accumulating all the tracks within a jet, tracks are weighted by a correction factor $C(\eta, \phi, p_T^{\text{track}}, p_T^{\text{jet}})$ to account for the tracking efficiency (ϵ), the misreconstruction rate (f), the rate at which real non-primary tracks contaminate the sample (s), and the probability that a real primary track is reconstructed multiple times (m). The correction factor is

$$C(\eta, \phi, p_T^{\text{track}}, p_T^{\text{jet}}) = \frac{(1-f)(1-s)}{\epsilon(1+m)}. \quad (5.2)$$

These tracking parameters are evaluated using a sample of PYTHIA6 tune Z2 for 2.76 TeV pp collisions, and PYTHIA+HIJING for 5.02 TeV pPb collisions. The corrections from the pPb data are also used for the 7 TeV pp data, as the efficiency was found to not vary much with \sqrt{s} or the detector occupancy for these types of collisions. The track corrections are calculated as a function of the track η , ϕ , and p_T , as well as the associated jet p_T using an iterative tracking correction procedure similar to the one used in Ref. [210]. It is difficult to generate enough simulated events to fully fill this four-dimensional space with sufficient statistical accuracy (especially for high- p_T tracks). Therefore, this procedure assumes that the corrections for tracks with similar kinematic properties should be close to each other. More precisely, it assumes that the correction can be factorized into the product of functions depending

only on one or two of the track parameters:

$$C(\eta, \phi, p_T^{\text{track}}, p_T^{\text{jet}}) = C^1(p_T^{\text{track}})C^2(\eta, \phi)C^3(p_T^{\text{jet}}). \quad (5.3)$$

To begin the calculation, C^1 is evaluated in MC assuming C^2 and C^3 are unity. This correction is applied back to the simulated tracks, and the closure is checked as a function of the track's η and ϕ position. Any nonclosures resulting from changes of the tracking performance with respect to η and ϕ are then corrected by updating C^2 . Finally, both C^1 and C^2 are applied to the MC and the nonclosure is checked versus the p_T of the jet that is closest to the track. C^3 is updated to fix any remaining nonclosures. The procedure of updating C^1 , C^2 , and C^3 to correct any nonclosure in track p_T , η and ϕ , and jet p_T is repeated until the values of all three corrections converge to stable values. Because the curvature of low- p_T and high- p_T tracks is very different, this procedure is run separately for tracks in seven bins of track p_T . The p_T ranges are 0.5–0.55, 0.55–0.65, 0.65–0.80, 0.8–1.0, 1–3, 3–8, and > 8 GeV. This coarse separation prevents the contribution from a large number of low- p_T tracks from overwhelming the high- p_T tracks in the closure tests. After this process, some residual nonclosure remains, as can be seen in Fig. 5-1. The left panel shows the nonclosure for 2.76 TeV pp collisions. The middle and right panels display the nonclosure for 5.02 TeV pPb and 7 TeV pp collisions, respectively. The deviation from unity is less than 2% for almost the entire p_T range studied for this analysis.

For the track selections used, the results of the tracking efficiency calculation are very similar to those shown in the left panel of Fig. 3-9. The small wiggle around 20 GeV is not present, however, because of the track of an explicit calorimeter matching requirement in the analysis. The misreconstruction rate f is less than 2% for most of the p_T range examined, similar to that shown in Fig 3-11. The secondary rate s is also less than 2%. In general, the misreconstruction rate m is negligible for the CMS reconstruction.

To check for any potential effects of pileup in the 7 TeV pp sample, the tracking closure test was repeated after selecting events with a reconstructed number of vertices




Figure 5-1: The tracking nonclosure for 2.76 TeV pp (left), 5.02 TeV pPb (middle), and 7 TeV pp (right) collisions.

either ≤ 4 or ≥ 7 . The corresponding average number of vertices in these samples is around 3.2 and 8.5. The nonclosures can be seen by the red and black points in Fig. 5-2. The dependence on the event pileup is found to be negligible compared to other sources of uncertainty.

After all tracks have been tabulated, the resulting distribution is normalized by the number of jets examined. Each track p_T bin is then divided by its width to form a differential quantity.

5.1.4 Underlying Event Subtraction

At this point, the measured quantity corresponds to the average per-jet yield of all tracks within a each jet cone. However, many of these tracks are not associated with the presence of the hard scattering in the event. They can be produced by pileup interactions or by soft physics processes. The latter background is particular large in pPb collisions. This contamination must be removed from the measurement. It is estimated by repeating the same track tabulation in a $R = 0.3$ cone that is rotated in azimuth $\pi/2$ away from the signal jet's cone. The resulting distribution is then subtracted from the signal distribution. This procedure could potentially be biased by the presence of an azimuthally-dependent background. Hydrodynamic flow in pPb collisions [53, 56, 55] could be one potential mechanism to generate such a




Figure 5-2: The tracking nonclosure for two different event selections in 7 TeV pp collisions. The selections were chosen to study the effect of pileup on the tracking efficiency.

dependence. As a cross check, another method was attempted. In this method, tracks in separate pPb MB event are tabulated in the same cone that is used for the signal event. Because these tracks come from a separate event, they are not correlated with the hard scattering of the signal event by construction. The difference between the particle yields of these two methods is up to 10% for track $p_T = 0.5$ GeV, but quickly shrinks to a negligible amount above 5 GeV. The resulting quantity after this subtraction procedure is the fragmentation function.

The fragmentation functions are also measured with respect to the variable

$$\xi = \ln \frac{p^{\text{jet}}}{p_{\parallel}^{\text{track}}}. \quad (5.4)$$

Here p^{jet} is the total jet momentum (not just the transverse component), and $p_{\parallel}^{\text{track}}$ is the projection of the track momentum onto the jet momentum vector.

5.2 Interpolated Reference

A data-driven approach is used to construct a 5.02 TeV pp reference by interpolating known results measured at 2.76 and 7 TeV. The procedure makes certain assumptions that are justified by the principle of QCD factorization. It assumes the fragmentation function of a quark jet at a given p_T will be equal to that of another quark jet at the same p_T , regardless of the collision center of mass energy. A similar assumption is made about gluon jets. Furthermore, the changing contributions from charm and bottom jets are assumed to be small. Then the total fragmentation function of any jet at some p_T is a linear combination of the quark and gluon fragmentation functions. The coefficients of this linear combination are simply the fractions of quark and gluon jets produced. These quark and gluon fractions do change with the collisions \sqrt{s} . Simulated PYTHIA 6 tune Z2 is used to model these fractions. Although this relies on a simulation, it should be emphasized that the quark and gluon fractions are constrained by the pp parton distribution function, which are in turn constrained by global fits of the world's data [8, 9, 10]. Thus, these fractions are known fairly precisely. The gluon fractions can be seen in Fig. 5-3 for both reconstructed and generator level jet p_T . In general, the fraction of gluon jets goes up as the collision energy gets higher. Furthermore, low- p_T jets are more likely to be gluon-initiated.

For a given jet p_T , the fragmentation functions of the 2.76 and 7 TeV pp data are first calculated. Both jet p_T spectra are reweighted to equal that of the 5.02 TeV pPb data in order to prevent the larger amount of high- p_T jets produced at 7 TeV from biasing the interpolation towards the 7 TeV data. This allows direct comparison of the interpolated function to the 5.02 TeV pPb data. Now let F_2 and F_7 be the 2.76 and 7 TeV pp fragmentation functions after jet spectrum reweighting. Call the quark and gluon fragmentation functions at 5.02 TeV Q and G , respectively. These quantities are all functions of the reconstructed track p_T . The gluon fraction of jets at 2.76 and 7 TeV is represented by r_2 and r_7 , while the quark fraction is simply given by $1 - r$. These two fragmentation functions can then be rewritten in terms of the quark and gluon fragmentation function basis as follows

$$F_2 = r_2 G + (1 - r_2) Q \quad (5.5)$$

$$F_7 = r_7 G + (1 - r_7) Q. \quad (5.6)$$

Solving for Q and G gives

$$G = \frac{(1 - r_7)F_2 - (1 - r_2)F_7}{r_2 - r_7} \quad (5.7)$$

$$Q = \frac{-r_7 F_2 + r_2 F_7}{r_2 - r_7}. \quad (5.8)$$

We can immediately construct the interpolated pp reference at 5.02 TeV, F_5 , using the gluon fraction from PYTHIA, r_5

$$F_5 = r_5 G + (1 - r_5) Q = \frac{(r_5 - r_7)F_2 + (r_2 - r_5)F_7}{r_2 - r_7}. \quad (5.9)$$

Thus, the interpolated fragmentation function is just a weighted average of the two other measured fragmentation functions. The weighting coefficients are given by where the 5.02 TeV gluon fraction lies with respect to the difference of the 2.76 and 7 TeV gluon fractions.

5.3 Systematic Uncertainties

A flat 3.9% systematic uncertainty is quoted for any remaining tracking performance differences between the data and MC samples. This value is justified by comparisons of the yields of the decays $D^{*+} \rightarrow K^- \pi^- \pi^+ \pi^+ \pi^+$ and $D^{*+} \rightarrow K^- \pi^+ \pi^+$ in pp data and MC [211]. Deviations in the ratio of these two processes, as compared to the world average quantity, are proportional to the square of tracking efficiency difference between data and MC. Tracks having the same η_{CM} in pp and pPb collisions will be located in different parts of the tracker because of the pPb rapidity boost. Thus, this uncertainty was conservatively not assumed to cancel in a the ratio of two

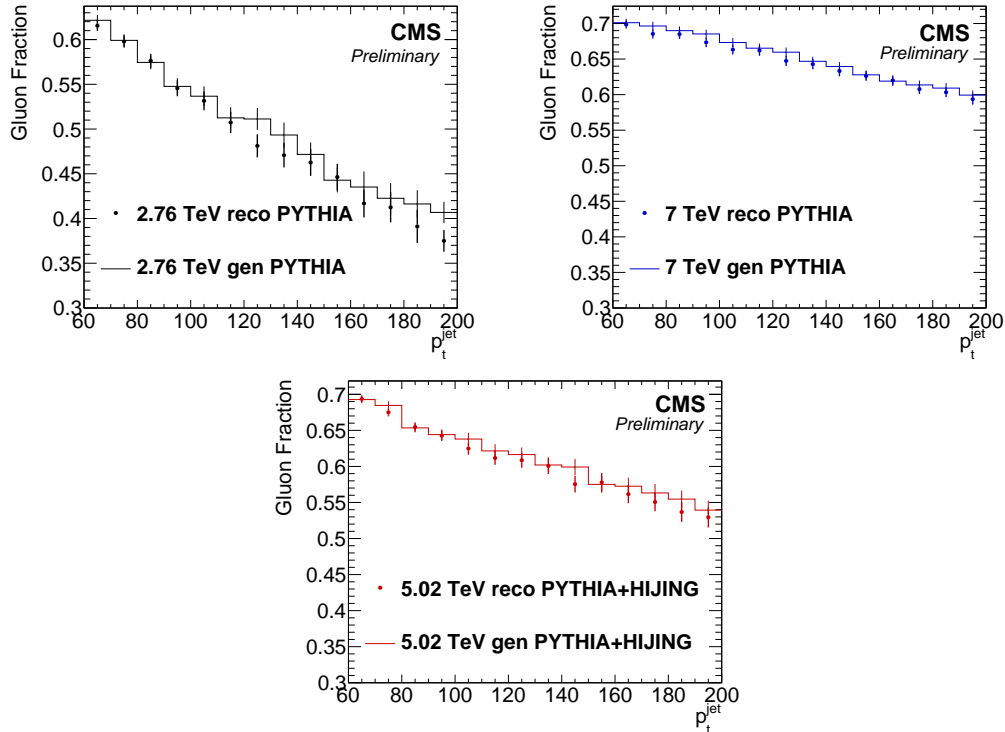


Figure 5-3: The fraction of gluon-initiated jets as a function of either generator-level or reconstruction-level jet p_T , as calculated by PYTHIA.

fragmentation functions.

An underestimation of the jet energy scale calibration (JES) can shift jets having a higher p_T into a lower p_T bin. This migration then makes the fragmentation of all jets appear harder. On the other hand, if the JES is overestimated low p_T jets will be shifted upwards making the fragmentation softer. The JES uncertainties are calculated in a fashion similar to the CMS 7 and 8 TeV global analyses [182], and are therefore η and p_T dependent values. For the kinematic selection used here, the pp jet energy scale uncertainty ranges from (0.7–1.6)% and the pPb uncertainty ranges from (0.8–2.3)%. The JES was varied up and down by its uncertainty for each of the three collision systems. For residual jet energy corrections, separate factors are used for the individual collision systems because the detector response may have changed between the 2011 and 2013 runs. These residual corrections are small, on the order of 1%. The full residual correction is taken as an uncertainty and propagated to the final result. The uncertainty originating from this correction is evaluated asymmetrically. The

hard part of the fragmentation functions is sensitive to variations in the jet energy. This one of the dominant systematic uncertainties for this kinematic region.

The jet energy resolution (JER) is estimated to be 10% of the jet p_T in all three systems used in this analysis [212]. To evaluate the uncertainty from fluctuations in the jet energy from this effect, jet energies were modified by a Gaussian smearing of 5%. This amounts to a 1% increase in the total JER from 10% to 11%. Differences between the smeared result and nominal result are taken as a systematic uncertainty. This uncertainty is on the order of 3%. The data reported here are not unfolded to remove the effect of the JER. However, most of the effects of the JER are expected to cancel out in a ratio of two fragmentation functions.

The main source of uncertainty from the interpolation procedure is its reliance on the correct gluon fractions that are obtained from simulation. This is directly related to the uncertainties of the input proton PDFs. The NNPDF 2.1 LO [213] library provides different parameterizations of the proton PDF that are supposed to represent the uncertainty in the PDF fitting procedure. The analysis was repeated using a set of different parameterizations, and the resulting range of gluon fractions was taken as an uncertainty. This uncertainty is no more than 0.5% when propagated to the final result.

A full test of the analysis machinery, including the interpolation procedure, was completed in simulated MC and compared to the expected generator-level outcome. Any remaining differences are quoted as a systematic uncertainty. This uncertainty includes the small nonclosure remaining from the track correction procedure. This correction is less than 7% for most of the kinematic range analyzed here.

In the analysis, a cut requiring the fraction of a jet's p_T coming from charged particles to be between 5 and 95% is applied. Although this cut helps reject fake jets, it could also possibly remove jets that have large fluctuations in the average amount of produced charge. The full analysis procedure was redone without the charge fraction cut. The difference between the result with and without the jet charge fraction cut was then taken as a systematic uncertainty. This uncertainty grows with the track p_T examined but is typically only a few %.

Halfway through the collection of the pPb data set used in this analysis, the directions of the proton and lead beams circulating in the LHC was reversed. This effectively leads to different detector conditions between the two sets of data. The data set was split into its normal and ‘reversed’ portions. Comparison of these two samples with the nominal result reveals small differences at low jet p_T . Half the difference between the two smaller samples is taken as a systematic uncertainty.

The underlying event subtraction using an azimuthally-rotated cone is checked by using a minimum bias event mixing subtraction method, as described earlier. The ratio of the two methods is calculated and taken as a systematic uncertainty. This uncertainty is up to 10% at very low track p_T , but rapidly shrinks to 0 at a track p_T of around 5 GeV.

Summaries of the various contributions to the systematic uncertainties are shown in chosen bins for the pPb fragmentation function in Table 5.2 and for the fragmentation function ratio in Table 5.3. A graphical summary of the uncertainties for the $140 < p_T^{jet} < 200$ GeV bin is shown for the pPb fragmentation function, the pp reference, and the ratio of the two in the left, right, and middle panels of Fig. 5-4, respectively. The jet related uncertainties are the dominant source for the fragmentation functions at high track p_T . At lower track p_T the uncertainties from the underlying event subtraction are significant. Between these two extremes, the tracking uncertainty dominates. The systematic uncertainties are significantly smaller for the pp collision systems. This causes the pp interpolation to have smaller systematic uncertainties than pPb data.

5.4 Results

The fragmentation functions used as an input into the interpolation procedure are shown in Fig 5-5. The top row of panels are the results for 2.76 TeV pp collisions. The bottom row are from 7 TeV pp collisions. Each panel is a selection of p_T^{jet} , with the rightmost panels being the highest p_T^{jet} bins. The fragmentation functions all have very similar shapes, regardless of the energy of the collision system. This is consistent

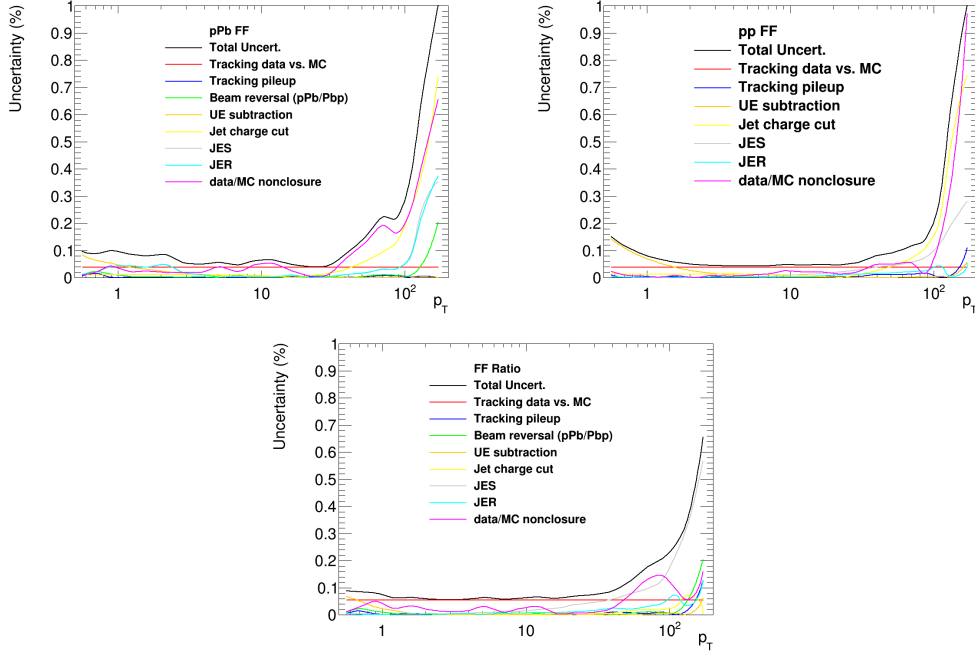


Figure 5-4: A graphical summary of the systematic uncertainty contributions for the 5.02 TeV pPb (left) and pp reference (right) fragmentation functions. (Middle) the uncertainty on the ratio is also shown. These values are for the $140 < p_T^{jet} < 200$ bin.

with the interpolation procedure's assumptions about QCD factorization. As the selected jet p_T increases, the fragmentation function extends to higher p_T^{track} because of the larger kinematic phase space allowed for hard splittings in the jet's parton shower. Likewise, the slope of the fragmentation function becomes very steep as p_T^{track} approaches p_T^{jet} as the kinematics become more constrained. The most likely p_T^{track} of a produced charged particle is around 1.5 GeV, and the probability of producing a particle above ≈ 30 GeV is only around 1% for the highest p_T^{jet} studied here. The fragmentation function for 5.02 TeV pPb collisions is shown in the top row of Fig. 5-6. The results from applying the interpolation procedure to pp data are displayed by the lower row of panels in the same figure. The shapes are very similar to those of the pp data.

Figure 5-7 displays the ratio between the pPb fragmentation function and the interpolated pp reference. The ratio is remarkably flat; all data points are consistent with unity. At low p_T^{track} the data points are always around 5% above unity. Al-

though such an excess is not significant, this is most likely caused by an incomplete subtraction of the underlying event and other soft particle production. The systematic uncertainties become very large for extremely hard jet fragments, as this region is very sensitive to the calibration of the jet energy.

The same result is shown in Fig. 5-8 as a function of ξ . Low values of ξ correspond to $p_T^{track} \approx p_T^{jet}$, while $\xi = 5$ corresponds to $p_T^{track} \approx p_T^{jet}/150$. Thus, the full track range of the analysis is used for the leftmost panel, while only tracks above $p_t \approx 1$ GeV contribute to the rightmost panel. In general the transformation from p_T^{track} to ξ moves the data points around a little, but the conclusion remains the same. No significant enhancement of particle production per jet is observed for 5.02 TeV pPb collisions.

This observation challenges the conclusions made from viewing the CMS charged particle [126] and jet [32] nuclear modification factor analyses together. If the fragmentation function is not enhanced and all high- p_T charged particles are produced within jets then the discrepancy between these two values would appear to be caused by some inconsistent treatment of the data. However, a preliminary reanalysis of the inclusive 5.02 TeV pPb charged particle spectrum indicates that this is not the case. The main suspicion for this discrepancy must then fall to the assumptions made while constructing the various pp references used in these analyses. To further complicate the problem, a *preliminary* report by the ATLAS Collaboration saw a slight modification of the pPb fragmentation of $\approx 10\%$ the same time these results shown [214]. However, ATLAS used yet another method to construct the pp reference. To circumvent the apparent inconsistencies caused by different reference methodologies a large push was made by the heavy ion community to collect a dedicated 5.02 TeV pp reference data set in 2015. Reference quantities could then simply be measured.

Once the 2015 reference data was collected, it was decided to leave this study as a preliminary result in favor of directly measuring the needed reference for a completely data-driven pPb nuclear modification factor. The ATLAS Collaboration, however, did update their study to a fully unfolded fragmentation function with a measured reference [215]. That result is compatible with the results shown here.

Table 5.1: The analysis level track selections used for various data sets.

| year | System | Energy (TeV) | <i>high purity</i> | σ_{p_T}/p_T | d_0 DCA | z DCA | N_{hits} | $\frac{1}{N_{layers}} \frac{\chi^2}{(d.o.f)}$ | E_T/p_T ($p_T^{Track} > 20$ GeV) |
|------|--------|--------------|--------------------|--------------------|------------|------------|------------|---|-------------------------------------|
| 2011 | pp | 7 | yes | <0.1 | $<3\sigma$ | $<3\sigma$ | — | — | — |
| 2013 | pp | 2.76 | yes | <0.1 | $<3\sigma$ | $<3\sigma$ | — | — | — |
| 2013 | pPb | 5.02 | yes | <0.1 | $<3\sigma$ | $<3\sigma$ | — | — | — |
| 2015 | pp | 5.02 | yes | <0.3 | $<3\sigma$ | $<3\sigma$ | — | — | >0.5 |
| 2015 | PbPb | 5.02 | yes | <0.1 | $<3\sigma$ | $<3\sigma$ | >10 | <0.15 | >0.5 |
| 2017 | XeXe | 5.44 | yes | <0.1 | $<3\sigma$ | $<3\sigma$ | >10 | <0.15 | >0.5 |

Table 5.2: Summary of the systematic uncertainties for the 5.02 TeV pPb fragmentation functions. The p_T values are given in units of GeV.

| Symmetric uncertainty source | $60 < p_T^{jet} < 80$ | | | | $140 < p_T^{jet} < 200$ | | | | |
|-------------------------------|---------------------------|--------------------------|--------------------------|---------------------------|--------------------------|--------------------------|---------------------------|--------------------------|--------------------------|
| | $0.5 < p_T^{track} < 5.9$ | $5.9 < p_T^{track} < 45$ | $45 < p_T^{track} < 200$ | $0.5 < p_T^{track} < 5.9$ | $5.9 < p_T^{track} < 45$ | $45 < p_T^{track} < 200$ | $0.5 < p_T^{track} < 5.9$ | $5.9 < p_T^{track} < 45$ | $45 < p_T^{track} < 200$ |
| Symmetric uncertainty source | | | | | | | | | |
| Tracking | 3.9% | 3.9% | 3.9% | 3.9% | 3.9% | 3.9% | 3.9% | 3.9% | 3.9% |
| PDF + gen interpolation | 0.5% | 0.5% | 0.5% | 0.5% | 0.5% | 0.5% | 0.5% | 0.5% | 0.5% |
| JER | 0.1 – 1.9% | 0.1 – 2.1% | 0.1 – 3% | 0.1 – 2.1% | 0.1 – 0.9% | 0.1 – 3% | 0.1 – 0.9% | 0.1 – 0.9% | 0.5 – 1.9% |
| pPb/Pbp difference | 0 – 0.5% | 0.1 – 1.7% | 4.6 – 14% | 0.1 – 1.7% | 0.1 – 2% | 4.6 – 14% | 0.1 – 2% | 0.1 – 1% | 0.1 – 3.9% |
| UE subtraction | 0.8 – 8.9% | 0.1 – 0.3% | 0.5 – 1.5% | 0.1 – 0.3% | 1 – 7.8% | 0.5 – 1.5% | 1 – 7.8% | 0.1 – 0.4% | 0 – 0.4% |
| Asymmetric uncertainty source | | | | | | | | | |
| JES (upwards uncert.) | 0.4 – 3.3% | 5.2 – 26% | 88 – 170% | 5.2 – 26% | 0.1 – 1.6% | 88 – 170% | 0.1 – 1.6% | 2 – 11% | 19 – 250% |
| JES (downwards uncert.) | -0.6 – -2.8% | -4.4 – -21% | -53 – -73% | -4.4 – -21% | -0.1 – -1.4% | -53 – -73% | -0.1 – -1.4% | -2.4 – -9.3% | -16 – -71% |
| Reco/gen difference | -1.8 – 1.4% | 0.4 – 6.8% | 31 – 82% | 0.4 – 6.8% | -0.7 – 3.4% | 31 – 82% | -0.7 – 3.4% | 1.7 – 4.1% | 1.1 – 78% |
| Jet charge fraction cut | -1 – 0.3% | -0.4 – 6.6% | 36 – 92% | -0.4 – 6.6% | -1.5 – -0.8% | 36 – 92% | -1.5 – -0.8% | -1 – 1.5% | 5.6 – 110% |
| Total (upwards uncert.) | | | | | | | | | |
| Total (downwards uncert.) | | | | | | | | | |
| Total (upwards uncert.) | 4.8 – 10% | 6.5 – 28% | 100 – 210% | 6.5 – 28% | 4.6 – 9.1% | 100 – 210% | 4.6 – 9.1% | 5.3 – 12% | 21 – 280% |
| Total (downwards uncert.) | -4.9 – -10% | -5.9 – -22% | -53 – -75% | -5.9 – -22% | -4.5 – -8.9% | -53 – -75% | -4.5 – -8.9% | -4.8 – -10% | -17 – -73% |

Table 5.3: Summary of the systematic uncertainties for the ratio between pPb and pp fragmentation functions. The p_T values are given in units of GeV.

| Symmetric uncertainty source | $60 < p_T^{jet} < 80$ | | | | $140 < p_T^{jet} < 200$ | | | | |
|-------------------------------|---------------------------|--------------------------|--------------------------|---------------------------|--------------------------|--------------------------|---------------------------|--------------------------|--------------------------|
| | $0.5 < p_T^{track} < 5.9$ | $5.9 < p_T^{track} < 45$ | $45 < p_T^{track} < 200$ | $0.5 < p_T^{track} < 5.9$ | $5.9 < p_T^{track} < 45$ | $45 < p_T^{track} < 200$ | $0.5 < p_T^{track} < 5.9$ | $5.9 < p_T^{track} < 45$ | $45 < p_T^{track} < 200$ |
| Tracking | 5.5% | 5.5% | 5.5% | 5.5% | 5.5% | 5.5% | 5.5% | 5.5% | 5.5% |
| PDF + gen interpolation | 0.5% | 0.5% | 0.5% | 0.5% | 0.5% | 0.5% | 0.5% | 0.5% | 0.5% |
| JER | 0.4 – 2% | 0.2 – 2.8% | 0.4 – 4.4% | 0.2 – 0.6% | 0.3 – 1.1% | 2.4 – 22% | 0.4 – 2% | 0.1 – 1% | 2.4 – 22% |
| pPb/Pbp difference | 0 – 0.6% | 0.2 – 1.8% | 4.8 – 15% | 0.1 – 2% | 0.1 – 1% | 0.8 – 3.6% | 0 – 0.6% | 0.1 – 1% | 0.8 – 3.6% |
| UE subtraction | 0.1 – 7.6% | 0.1 – 0.4% | 0.9 – 2.4% | 0.1 – 5.5% | 0 – 0.1% | 0 – 0.4% | 0.1 – 7.6% | 0.1 – 0.4% | 0 – 0.4% |
| Asymmetric uncertainty source | | | | | | | | | |
| JES (upwards uncert.) | 0.4 – 2% | 1.4 – 6.5% | 24 – 36% | 0.3 – 1.4% | 0.8 – 2.4% | 4.7 – 39% | 0.4 – 2% | 1.4 – 6.5% | 24 – 36% |
| JES (downwards uncert.) | -0.4 – -1.7% | -1.5 – -7.1% | -19 – -35% | -0.4 – -1% | -0.8 – -2.4% | -3.9 – -26% | -0.4 – -1.7% | -1.5 – -7.1% | -19 – -35% |
| Reco/gen difference | -4.4 – -1.7% | -1.9 – 0.9% | -33 – -6.2% | -1 – 2% | -0.2 – 1% | -5.3 – 8.4% | -4.4 – -1.7% | -1.9 – 0.9% | -33 – -6.2% |
| Jet charge fraction cut | -0.6 – 2.2% | -0.1 – 0.6% | 5.7 – 20% | -0.4 – 0.5% | -0.1 – 0.2% | 0.7 – 28% | -0.6 – 2.2% | -0.1 – 0.6% | 5.7 – 20% |
| Total (upwards uncert.) | 5.6 – 9.8% | 5.7 – 9.2% | 26 – 44% | 5.6 – 8% | 5.7 – 6.3% | 8.1 – 53% | 5.6 – 9.8% | 5.7 – 9.2% | 26 – 44% |
| Total (downwards uncert.) | -5.9 – -11% | -6 – -9.8% | -22 – -51% | -5.6 – -7.9% | -5.6 – -6.2% | -7.4 – -35% | -5.9 – -11% | -6 – -9.8% | -22 – -51% |




Figure 5-5: The fragmentation functions for 2.76 TeV pp (top row) and 7 TeV pp (bottom row) collisions. The selected jet p_T increase from left to right in each of the panels.

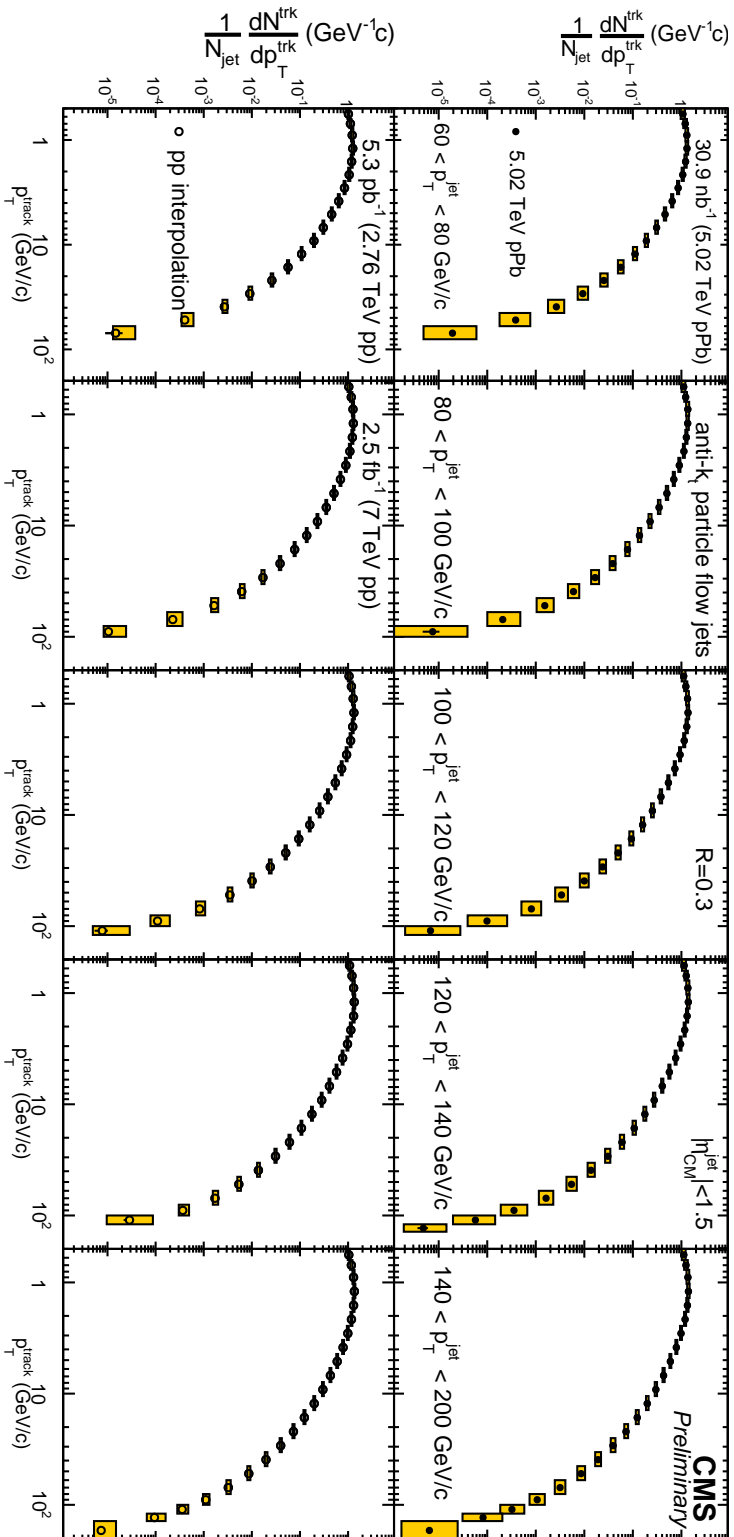


Figure 5-6: The fragmentation functions for 5.02 TeV pp collisions (top row) and an interpolated 5.02 TeV pp reference (bottom row). The selected jet p_T increase from left to right in each of the panels.




Figure 5-7: The ratio of the pPb and pp reference fragmentation functions as a function of track p_T . The selected jet p_T increase from left to right in each of the panels.




Figure 5-8: The ratio of the pPb and pp reference fragmentation functions as a function of ξ . The selected jet p_T increase from left to right in each of the panels.

Chapter 6

pPb and PbPb Nuclear Modification Factors

The unprecedented size of the 5.02 TeV PbPb and pp data sets collected in 2015 opened up an exciting new opportunity to study charged particle production in a p_T range that had not been previously accessible. This chapter concerns an analysis of the charged particle spectra in these collisions [208]. Using this information, the nuclear modification factors in PbPb collisions are measured. Furthermore, the pp reference data allows for an update of the pPb nuclear modification factor. For this update, the pPb spectra from Ref. [126] are used.

6.1 Analysis Procedure

The primary goal of this analysis is to measure the 5.02 TeV PbPb charged particle nuclear modification factor in six centrality intervals. Previous studies indicate the η dependence of R_{AA} is fairly small in the acceptance range of the CMS tracker [145]. Therefore, only the range $|\eta| < 1$ is used to avoid effects related to the sharp increase in the tracker's material budget outside of this region. For PbPb collisions, only tracks above $p_T > 0.7$ GeV are used. The measurement is taken down to $p_T > 0.5$ GeV in pp collisions.

6.1.1 Track Selection

In pp collisions, tracks are required to pass the same track quality selections as were used in the pPb Fragmentation function analysis, with the exception that the cut on the maximum relative p_T uncertainty was relaxed from 0.1 to 0.3. Additionally, tracks having $p_T > 20$ GeV are explicitly required to be matched to a calorimeter deposit using the methods described in Section 3.4.4.

For PbPb collisions, the track selection criteria is much more stringent to reduce the large misreconstruction rate that occurs in central events. Tracks must have a DCA less than 3σ in the xy plane and z direction to increase the probability that they are primary tracks. They must be *highpurity* tracks and their relative p_T uncertainty must also be less than 0.1. Misreconstructed tracks tend to have less hits, so a track is required to have at least 11 hits along its trajectory. A calorimeter match is also required for tracks above 20 GeV. Finally, tracks having a poor fit quality are removed by requiring the χ^2 , divided by the product of the number of degrees of freedom (d.o.f.) and the number of hit layers (N_{layers}), to be less than 0.15 in data. Using the number of hit layers in the denominator of this selection penalizes short tracks and makes it more likely for long tracks having a large number of hits to pass the requirement. A slight disagreement between data and simulation for the distribution of $\chi^2/((N_{layers})(d.o.f))$ exists, most likely because of differences in the real and simulated detector alignment. This discrepancy seems to be larger for high p_T tracks. However, a simple rescaling of the χ^2 distribution in simulation appears to restore the agreement between data and simulation. This can be seen in the left panel of Fig. 6-1 for tracks in the range $4.8 < p_T < 5.6$ GeV. The pink points are the MC distribution, which has a peak that is shifted to the right as compared to the black data points. After applying a scaling along the horizontal axis, the distribution in MC shown by the green points aligns much more closely with data. To prevent the removal of more tracks in simulation than in data, the cut on $\chi^2/((N_{layers})(d.o.f))$ is moved to a slightly larger value depending on how much scaling of the distribution is needed to make MC agree with data. The amount of scaling needed is shown by the




Figure 6-1: (Left) The distribution of $\chi^2/((N_{layers})(d.o.f))$ in 30–100% PbPb collisions for data and MC. (Right) The required amount of scaling of the MC distribution to achieve adequate agreement with data.

right panel of Fig. 6-1, and is calculated as a function of both collision centrality and track p_T . At high track p_T it becomes difficult to extract the values because of the poor statistical accuracy, so these points are fit with a flat function. A summary of the track selection requirements for this analysis are shown in the 4th and 5th rows of Table 5.1.

6.1.2 Tracking Correction

Tracks passing these selection criteria are then simply counted. During this tabulation, tracks are weighted by two factors. The first is a correction that takes into account the effects of tracking efficiency, track misreconstruction, secondary tracks, and the chance of reconstructing a track multiple times. This correction is calculated in a similar way to the correction used in the pPb fragmentation function described in equations 5.2 and 5.3. There are a few important distinctions, however. The parameterization of C^3 does not use the nearest jet p_T , as tracks are not constrained to be inside of a jet cone. Instead, the distance to the nearest anti- k_T jet having $p_T > 40$ GeV is used. The distance is quantified using the ΔR value defined in Equation 5.1. Additionally, an extra correction factor $C^4(centrality)$ is added for PbPb collisions.




Figure 6-2: The tracking efficiency for PbPb collisions as a function of track p_T and collision centrality. The values are calculated with PYTHIA+HYDJET.

This factor depends on the collision centrality and is crucial for describing the strong centrality dependence of the track correction in central events. The value of C^4 is calculated during the same iterative procedure that calculates C^1 , C^2 , and C^3 , and is determined between the iterations that consider the tracking closure as a function p_T and (η, ϕ) . The pp tracking efficiency in PYTHIA is shown in the left panel of Fig. 3-9. It is around 90%, with a slight dip around 20 GeV due to the calorimeter matching requirement. The pp tracking misreconstruction rate is relatively small and is shown in the left panel of Fig. 3-11. The secondary rate is less than 3%. The PbPb tracking efficiency, as evaluated in PYTHIA+HYDJET, is shown in Fig. 6-2 as a function of track p_T and event centrality. In general the efficiency is between 70 and 80%. However, it shrinks to around 50% for low p_T tracks. A further degradation of the efficiency is seen starting around 30% because of the high detector occupancies observed in central events. The misreconstruction rate is less than 1.5% except for tracks under 1 GeV in central events, where it grows to 10%.

The track p_T resolution is $< 5\%$ for the entire kinematic range studied here. The reported distributions are not unfolded to account for this resolution effect, as it is estimated to cause less than a 1% difference in the final charged particle spectra.

6.1.3 Trigger Combination

The second weight applied to each track is used to combine data coming from different triggers. This weight is an event-level quantity; all tracks in the event are given the same weight. The weights are calculated using a data-driven procedure that is similar to the one used in Ref. [126].

The total cross section of MB collisions can be split into different categories based on the p_T of an object that is to be triggered on. The maximum p_T of any track in the event is a natural choice to use because this measurement is interested in charged particle p_T spectra. However, every high p_T charged particle is contained within a high p_T jet, so a jet trigger can also be used. In this analysis, jet triggers are used for the nominal result because they are simpler and tend to be more robust against small changes in the detector conditions. The analysis is repeated using a similar set of data collected with track triggers, and the two results are compared as a cross check. The p_T thresholds used to define the different categories are chosen such that each category corresponds to the region in which a trigger is fully efficient. These thresholds are given in Table 6.1 for 5.02 TeV pp collisions. For each category events from only one trigger are considered. In the lowest p_T category this is usually a MB trigger, while the highest p_T category is usually a trigger that was never prescaled. Without any extra weighting this procedure will give a spectrum with sharp discontinuities when each new trigger threshold is used, because each trigger has a different prescale. The goal is to then calculate an effective weight to give to the events in each category that will smooth out the spectrum by correcting for the prescale effects. For pp collisions, the goal is to measure a differential cross section and not a differential per-event yield. Therefore, the highest p_T unprescaled trigger has its weights set to unity, because the total yield for this trigger can be converted to a cross section by simply dividing by the integrated luminosity. All other categories have their weights calculated relative to this category. The relative weight for the next lowest category is then calculated by taking a ratio of the number of events in a kinematic range where *both* triggers are fully efficient. This range is usually just taken to be the same as

Table 6.1: Event classes for the measurement of charged particle spectra in pp collisions.

| Leading object (p_T [GeV]) | | Trigger |
|-------------------------------|------------------|--------------|
| Jet | $p_T < 40$ | Minimum bias |
| | $60 < p_T < 80$ | Jet40 |
| | $80 < p_T < 100$ | Jet60 |
| | $100 < p_T$ | Jet80 |
| Track | $p_T < 20$ | Minimum bias |
| | $20 < p_T < 26$ | FullTrack18 |
| | $26 < p_T < 36$ | FullTrack24 |
| | $36 < p_T < 47$ | FullTrack34 |
| | $47 < p_T < 55$ | FullTrack45 |
| | $55 < p_T$ | FullTrack53 |

the category for the higher-threshold trigger. The resulting value can be considered an effective ‘average prescale.’ The second highest threshold trigger can now be used as a reference to calibrate the weight of the third highest threshold trigger. Thus, triggers can be chained together until all triggers have a weight assigned to them.

The result of this procedure is shown in Fig. 6-3 for 5.02 TeV pp collisions. The figure displays the distribution of maximum jet p_T after combining the contributions from MB, Jet40, Jet60, and Jet80 triggers. The total spectrum is shown by the black line, and has been scaled by 100 for clarity on the plot. The contributions of each trigger are represented by the colored data points. As the triggers become fully efficient, they converge to lie on top of each other. The dashed lines show the regions which use different triggers. For example, region I uses the MB trigger, region II uses Jet40, region III uses Jet60, and region IV uses Jet80. A similar procedure was applied to data collected with a set of track triggers and using the maximum track p_T . Figure 6-4 shows the ratios of the yields of each trigger with respect to the trigger one threshold lower. Once the triggers are fully efficient, they all converge to unity, indicating that the relative weights of each trigger are properly calculated. Similar results were obtained when using the set of track triggers. This method of combining triggers together is very robust. Because the triggers are copies of each other, only with different thresholds, ratios between pairs of them are insensitive to any trigger bias effects. The kinematic selection of the trigger object is also unimportant.




Figure 6-3: Distributions of the leading jets for pp collisions after each trigger has been weighted by an effective prescale. Refer to the main text for more details.

After the weights have been calculated, the total track spectrum can be formed by adding the contributions from each trigger together. This is shown in Fig. 6-5 for pp collisions. The black points in the top left panel shows the total track spectrum as calculated with jet triggers, before correcting for tracking efficiency. The contributions from each trigger are shown by the stacked colored histograms. The contribution from MB events is dominant at low track p_T . At high p_T , all tracks come from the Jet80 trigger. In between, the total yield is a complicated mixture of various triggers. The top right plot shows a similar spectrum that was calculated with track triggers. Because this trigger object corresponds more closely to the observable examined, the contribution from each trigger is more localized to a p_T range that the trigger is looking for. Nonetheless, the difference between these two spectra was found to be less than 1.5%. The bottom left and bottom right panels display the relative contribution of each trigger to the total spectrum for the jet and track trigger cases, respectively.

A similar procedure is used for combining PbPb triggers together. There are two main differences in the procedure, however. The first difference is that events are




Figure 6-4: Ratios of the distributions of the leading jets for pp collisions after each trigger has been weighted by an effective prescale.

first classified by their centrality into one of three bins: 0–30%, 30–50%, and 50–100%. The 30–100% and 50–100% bins include additional centrality gated triggers into the trigger combination strategy to boost the statistical reach of the sample. In the centrality region where they are fully efficient, these gated triggers are direct copies of the non-gated triggers. Therefore their counts can be simply added to the non-gated counts before calculating any relative overlaps between triggers.

The second difference is that a per-event differential yield is desired. This means that the weight for a MB event should be set to unity (rather than the weight of an unrescaled trigger event). Therefore, the ‘chaining’ procedure for calculating weights proceeds in the opposite direction. The second lowest trigger threshold is compared to MB, then the third highest is compared to the second highest, and so on.

The triggers and p_T categories used for the PbPb trigger combination are shown in Table. 6.2. The overlaps between jet triggers in the 0–30% centrality interval are shown in the left panel of Fig. 6-6. The trigger ratios all converge to unity once the triggers are fully efficient. The right panel displays the track spectrum for 0–5% events calculated from the jet trigger combination, along with the relative contribution of




Figure 6-5: (Top) The uncorrected total track invariant yields for pp collisions, as calculated with jet (left) and track (right) triggers. (Bottom) The relative contribution of each trigger to the spectra shown in the top row.

each trigger. Fig. 6-7 shows the same quantities for the track trigger combination. Both the jet and track trigger methods were found to be consistent with each other within 1%.

6.1.4 Particle Composition Correction

The probability of successfully reconstructing a track in CMS depends on the type of charged particle that creates it. Fig. 6-8 shows the efficiency for pions, kaons, protons, and anti-protons in 5.02 TeV PbPb collisions. The efficiency for pions and kaons is nearly the same. Protons have a slightly lower efficiency at low momenta, but are similar to pions and kaons at high p_T . The efficiency for anti-protons is lower than protons because of the potential of annihilating with the protons in the detector material. The right panel shows the tracking efficiency for four types of strange baryons: Σ^+ , Σ^- , Ξ^- , and Ω^- (note that Σ^+ and Σ^- are not anti-particles of each

Table 6.2: Event classes for the measurement of charged particle spectra in PbPb collisions.

| Leading object (p_T [GeV]) | | Trigger |
|-------------------------------|-------------------|--------------|
| Jet | $p_T < 40$ | Minimum bias |
| | $60 < p_T < 80$ | Jet40 |
| | $80 < p_T < 100$ | Jet60 |
| | $100 < p_T < 120$ | Jet80 |
| | $120 < p_T$ | Jet100 |
| Track | $p_T < 14$ | Minimum bias |
| | $14 < p_T < 20$ | FullTrack12 |
| | $20 < p_T < 35$ | FullTrack18 |
| | $35 < p_T < 50$ | FullTrack24 |
| | $50 < p_T$ | FullTrack34 |

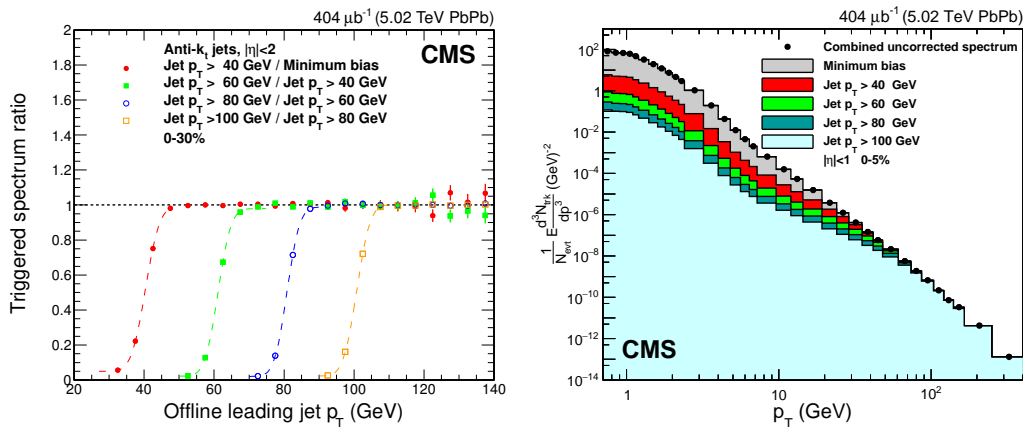


Figure 6-6: (Left) Ratios of the counts from various jet triggers after weighting each one by its ‘relative prescale’ in 0–30% PbPb collisions. (Right) The resulting uncorrected track spectrum in the 0–5% centrality range.

other). These are not stable particles, and they tend to decay into pions, protons, or neutrons after traveling a few cm in the detector. Because they rarely travel far enough to leave a significant number of hits in the tracker, they all have very low efficiencies. As their p_T increases, they are more likely to leave a significant number of hits in the tracker because of time dilation effects, and the efficiency therefore slowly increases. These strange baryons have lifetimes that are just slightly below 1 cm in the rest frame and therefore fall under the definition of a primary particle.

When deriving the tracking efficiency correction, it was assumed that data has the same composition of particles as the MC used to derive the correction. In the




Figure 6-7: (Left) Ratios of the counts from various track triggers after weighting each one by its ‘relative prescale’ in 0–30% PbPb collisions. (Right) The resulting uncorrected track spectrum in the 0–5% centrality range.




Figure 6-8: (Left) The tracking efficiency as a function of p_T for the most common particles produced in hadronic collisions. (Right) The tracking efficiency for various strange baryons.

case of this study, the relevant particle composition is that of the PYTHIA generator. However, it has been observed that the relative amount of strange particles is enhanced in PbPb, pPb, and even high-multiplicity pp collisions [66, 68, 67]. This implies that if a MC generator not modeling this enhancement is used to evaluate the tracking efficiency, the efficiency will be systematically overestimated. Comparisons of PYTHIA to strange particle yields in pp collisions indicate that this is the case [216].

Other generators produce different particle compositions depending on the physics that is implemented. A comparison of the relative fraction of each particle species in PbPb events, as a function of p_T , for embedded PYTHIA, MB HYDJET, and MB

EPOS can be seen in Fig. 6-9. In general, both EPOS and HYDJET have pion fractions that are significantly different than PYTHIA around 3 GeV. This peak at 3 GeV corresponds to the p_T where the effects of hydrodynamic flow have been observed to be maximal. Most of the pions have been replaced by K^\pm , p, and \bar{p} in HYDJET, but the strange baryon production is not significantly changed. Because the tracking efficiency of these particles is similar to that of pions, this means the efficiency computed in HYDJET will be similar to that of PYTHIA. The missing pions in EPOS, however, are replaced with less K^\pm and more strange baryons. Measurements by the ALICE Collaboration indicate that these strange baryons might be slightly overproduced in this model [66]. This makes the EPOS tracking efficiency lower in the range around 3 GeV. At higher p_T the particle fractions converge to each other, and therefore the model-dependence of the tracking efficiency then becomes fairly negligible.

Unfortunately detailed centrality-dependent and p_T -dependent spectra measurements were not available for 5.02 TeV PbPb collisions at the time of this analysis. Additionally the particles that are most important for this effect are the Σ^+ and Σ^- , as they are produced much more often than the Ξ^- and Ω^- . The Σ^\pm tend to decay to a single charged track plus a neutral particle and are therefore extremely difficult to reconstruct in CMS. For these reasons, it was decided to use a conservative model-dependent approach to account for any potential efficiency differences. PYTHIA is expected to underestimate the amount of strange baryons, while EPOS seems to overestimate it. Therefore a second correction is calculated that will place the tracking efficiency halfway between the nominal PYTHIA values and those that would be calculated in EPOS. A large uncertainty is associated with this procedure. A similar correction is applied to pp collisions, using PYTHIA and EPOS. However, in pp collisions preliminary results for the Ξ^- and Ω^- yield were available. Therefore in the model, the relative yields of these particles were replaced by the measured yields from data. The relative pion yields are slightly adjusted after this procedure to ensure that the relative fraction of all charged particles still sums to unity. Fig. 6-10 shows the size of the correction for PbPb and pp collisions in the left and right panels, respectively. It is given as a relative percentage that the charged particle spectra

are increased by. In both cases, the correction peaks strongly at 3 GeV, where there is the largest uncertainty in the amount of strange baryons. The effect is larger in PbPb collisions than in pp collisions. At higher p_T most particles are produced from jets, and therefore PYTHIA does a reasonable job of capturing the relevant physics for describing the particle fractions produced in data.

When revisiting the nuclear modification factor in pPb collisions, the spectra for pPb collisions are taken from Ref. [126]. To ensure that the pPb and pp data are treated in a similar fashion, a particle composition correction was also derived for pPb collisions using PYTHIA and EPOS. This correction was applied retroactively to the pPb data points, and the corresponding systematic uncertainties were updated to reflect this new correction. The only other modification to the pPb spectrum is the removal of a correction that accounts for the finite size of the p_T bins. This correction is not applied for the measured pp reference, so it should also not be applied in the pPb data for a consistent treatment of the data. Any effects caused by the removal of this correction are expected to nearly fully cancel out in a ratio between pPb and pp data.

6.1.5 Normalization

Once different triggers have been combined and the appropriate corrections applied, the spectra are transformed into differential invariant yields [3]:

$$E \frac{d^3 N}{dp^3} = \frac{1}{p_T} \frac{d^3 N}{d\phi d\eta dp_T}. \quad (6.1)$$

This is done by dividing the spectra by factors of $d\phi = 2\pi$, $d\eta = 2$, and the width dp_T of the relevant p_T bin. The center of the p_T bin is also divided out to make the quantity Lorentz-invariant.

Finally, for PbPb collisions the result is divided by the total number of MB events analyzed for the given centrality bin. Around 250 million inclusive MB events are used. In pp collisions, the desired measurement is that of a per-event differential cross section. Because the trigger combination was done with respect to an unrescaled

Table 6.3: Systematic uncertainties associated with the measurement of the charged-particle spectra and R_{AA} using $\sqrt{s_{NN}} = 5.02$ TeV pp and PbPb collision data. The ranges quoted cover both the p_T and the centrality dependence of the uncertainties. The combined uncertainty in R_{AA} does not include the integrated luminosity and the T_{AA} uncertainties.

| Sources | Uncertainty [%] | | |
|--|-----------------|----------|----------|
| | pp | PbPb | R_{AA} |
| Particle species composition | 1–8 | 1.0–13.5 | 1.5–15.5 |
| MC/data tracking efficiency difference | 4 | 4–5 | 2.0–6.5 |
| Tracking correction procedure | 1 | 1–4 | 1.5–4.0 |
| PbPb track selection | – | 4 | 4 |
| Pileup | 3 | <1 | 3 |
| Fraction of misreconstructed tracks | <3 | <1.5 | <3 |
| Trigger combination | <1 | 1 | 1 |
| Momentum resolution | 1 | 1 | 1 |
| Event selection correction | <1 | – | <1 |
| Combined uncertainty | 7–10 | 7–15 | 7.0–17.5 |
| Glauber model uncertainty (T_{AA}) | – | – | 1.7–16.0 |
| Integrated luminosity | 2.3 | – | 2.3 |

trigger, this can be calculated by simply dividing by the total sampled luminosity of 27.4 pb^{-1} .

Nuclear modification factors are built from the measured spectra using Equation 1.9 and values of T_{AA} given in Table 4.2.

6.2 Systematic Uncertainties

A summary of the systematic uncertainties for this analysis can be seen in Table 6.3.

The largest systematic uncertainty originates from the particle species correction, which is largely dependent on the input MC models used in this analysis. The full value of this correction is assigned as an uncertainty, which means that the nominal result is consistent with results that would be produced using only PYTHIA or EPOS. This is a very conservative choice. As can be seen in Fig. 6-10, the size of the uncertainty therefore also peaks around 3 GeV, where it is 13.5% in PbPb collisions and 8% in pp collisions. At $p_T > 10$ GeV, a flat systematic of 1% is used to account for any remaining differences in this kinematic range. This systematic uncertainty is

not assumed to cancel out in any ratios.

The second largest uncertainty results from any remaining differences between data and MC that are not related to different fractions of particles in each sample. In pp collisions this uncertainty is 4%, using a similar data-driven study of D^{*+} decays that was used for the pPb fragmentation function analysis [211]. In PbPb collisions this uncertainty is increased to 5% to account for some small differences observed in the track quality distributions (including the reweighted χ^2 distributions) between data and MC. This uncertainty is assumed to partially cancel in a ratio of spectra in peripheral events. However, it does not cancel for central events.

The iterative tracking correction procedure leads to some small residual nonclosures that are accounted for with a systematic uncertainty. There is also a statistical component related to the finite size of the MC sample used for the correction. This uncertainty is around 1% for most of the p_T and centrality range studied but can be as large as 4% for a few p_T bins in central events.

A looser variant of the PbPb track selection was used to test the robustness of the tracking procedure. This looser variation removes the χ^2 and N_{hits} cuts. Differences of up to 4% compared to the nominal result were observed when using these looser selections. This can be seen for two centrality selection in Fig. 6-11.

The pileup in PbPb collisions is negligible for this analysis. However, the effect of pileup in pp collisions may not be. The spectra calculated with events having precisely 1 primary vertex was compared to the nominal result to calculate a systematic uncertainty for this effect. The differences are mostly less than 1%, but a few p_T bins have up to a 3% difference. An flat uncertainty of 3% is conservatively assigned.

The fraction of misreconstructed tracks in both samples is $< 3\%$ and $< 1.5\%$ in pp and PbPb collisions, respectively, because of the stringent track selections. The full misreconstruction rate is assigned as a systematic uncertainty to cover any differences in the data and MC. This difference between data and MC is not probed by the data-driven studies used to evaluate previously-mentioned differences in the tracking efficiency and is therefore considered as a separate uncertainty.

The results using the trigger combination of either jet or track triggers agree within

1% of each other. A p_T dependent uncertainty of around 1% is used to account for any remaining uncertainties in this procedure.

The momentum resolution of CMS is excellent, so this result is not corrected for the effects of finite p_T resolution. A systematic uncertainty of 1% is used to account for any residual differences from the shifting of tracks between p_T bins because of the resolution.

The effects of event selection are determined to be negligible for this analysis in pp collisions. The event selection is fully efficient at selecting events having at least one track above $p_T = 0.5$ GeV.

For building the nuclear modification factors, the uncertainties for the Glauber model quantities are calculated by propagating the uncertainties on the input parameters into the final results. The uncertainties are displayed in Table 4.2. In general, they are negligible in central events but get large in peripheral events.

The pp luminosity is used to normalize the pp spectra into a per-event differential cross section. The uncertainty on this quantity was determined using a Van der Meer Scan procedure and is 2.3% [206]. This uncertainty, as well as the Glauber uncertainty, is fully correlated across all bins in p_T .

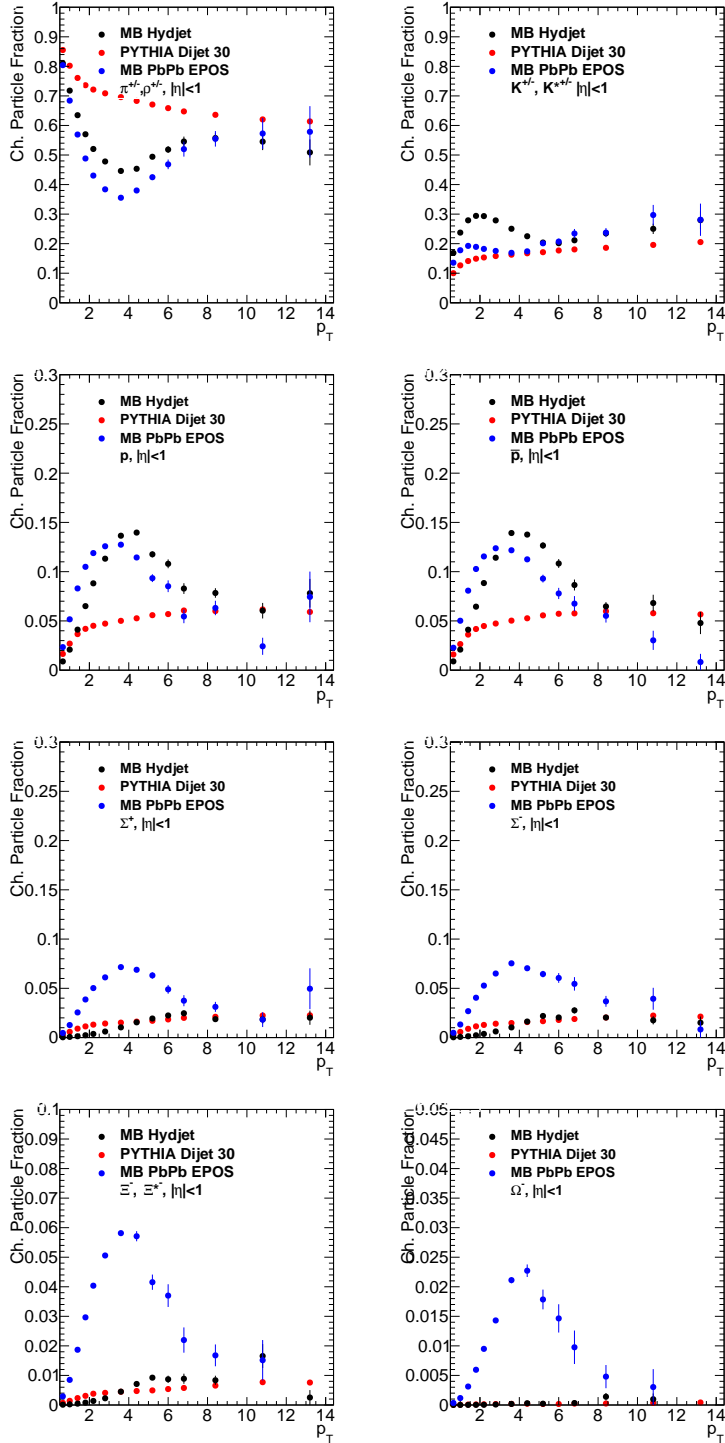


Figure 6-9: The relative fraction of charged particles coming from various particle species, as a function of p_T , for three different MC generators. The particles are, from the top left panel to the bottom right panel, π^\pm , K^\pm , p , \bar{p} , Σ^+ , Σ^- , Ξ^- and Ω^- .




Figure 6-10: The particle species residual correction for 5.02 TeV PbPb collisions (left) and pp collisions (right).




Figure 6-11: A comparison of the PbPb spectra in the 10–30% and 50–100% centrality bins, after varying the track selections to looser criteria.

6.3 Results

6.3.1 PbPb Nuclear Modification Factor

The charged-particle spectra are shown in Fig. 6-12 for six different centrality selections. The data for the 0–5% and 5–10% selections has been scaled by constant factors of 10 and 3, respectively, for visual clarity. The pp measurement is also shown on the same figure. A constant factor of 70 mb is used to convert the measured pp differential cross section into a per-event differential yield so that it can be directly compared to the PbPb data. The statistical error bars are smaller than markers for most of the data shown. In all cases, the reported yields are the average yield across the entire width of the bin; no ‘bin-shift’ correction [126] is applied. At high p_T the spectra have a roughly power law dependence. This seems to be violated for low p_T charged particles. The shape of the distributions is qualitatively similar to a Tsallis function, which has been known to describe charged particle spectra at the LHC [217]. Representative systematic uncertainties are given in the bottom panels as a relative percent uncertainty, for 0–5% and 70–90% PbPb events and pp collisions. The uncertainties tends to peak around $p_T = 4$ GeV, and plateau above around 10 GeV.

The corresponding nuclear modification factors in PbPb collisions are shown in Fig. 6-13. The 5.02 TeV data is represented by the black data points. The black error bars are the statistical uncertainties, while the yellow boxes are the systematic uncertainties. The light blue box around unity represents the relative Glauber model uncertainty, while the light gray box is the pp luminosity uncertainty. A characteristic suppression is observed for most of the p_T range studied. Additionally, a local maximum is seen around 2 GeV, and a corresponding local minimum is present at 7 GeV for all centrality selections. The strength of the suppression is much larger in central collisions. This suppression is presumably the result of an admixture of many effects, including but not limited to nuclear PDF contributions [78], elliptic flow [48], the Cronin effect [83, 84], and parton energy loss [218]. The maximal suppression is seen in 0–5% events, where it is a factor of 7–8 around 6–9 GeV. At larger p_T the




Figure 6-12: The charged particle spectra for 5.02 TeV PbPb and pp collisions. The pp data has been scaled by a constant factor of 70 mb so it can be directly compared to the PbPb measurement.

effects of radial flow are expected to be relatively small and the shape of the R_{AA} should be governed by the effects of parton energy loss. A continual rising slope is observed, until the data points become compatible with unity around 200 GeV. As the centrality selection becomes more peripheral, the R_{AA} line shape becomes less suppressed and flatter. In the 70–90% centrality bin, R_{AA} is relatively flat around a value of approximately 0.75. It should be noted that a large normalization uncertainty is related to the Glauber model in this centrality bin. The flat shape is relatively consistent with the expectation of a rather small suppression in peripheral events. Nonetheless, the suppression of the magnitude below unity is not expected to

be caused by strong quenching effects here, as jet quenching signals in other studies of this centrality range have been rather small [28, 29, 27]. Potential causes for this peripheral suppression will be explored further in Section 7.4.

In Fig. 6-13 previous measurements at 2.76 TeV by the CMS [146], ATLAS [145], and ALICE [142] Collaborations are shown by the red, blue, and green points, respectively. The CMS 2.76 TeV result has a 6% pp luminosity uncertainty that is not shown in the figure. In the range of $1 < p_T < 5$ GeV, the 5.02 TeV data tend to be slightly higher than the 2.76 TeV data (although this discrepancy is not significant). This is most likely due to the inclusion of the particle species correction in this analysis, which moves the data points in this kinematic range slightly upward. Such a correction was not considered in previous analyses. In general, the measured R_{AA} values are in quantitative agreement at high p_T , despite the difference in center of mass energy. This seems to be similar to the situation that arose when comparing $\sqrt{s_{NN}} = 200$ GeV RHIC data to $\sqrt{s_{NN}} = 2.76$ TeV LHC data, where the center of mass energy did not seem to strongly affect the values of R_{AA} . Presumably the effects of a hardening of the particle spectra largely offset any potential increases in parton energy loss. Indeed, recent studies that use a simple BDMPS model (see Section 1.4.1) to account for the hardening spectrum indicate that the relative energy loss in 5.02 TeV collisions is 10–20% larger than at 2.76 TeV [219].

The measurements in the 0–10% and 30–50% centrality ranges are compared to the theoretical predictions coming from Refs. [220, 94, 113, 105, 117, 221] in Fig. 6-14. The details of these models have been discussed in Section 1.4.1. The SCET_G model is shown by the shaded red band and does an excellent job of predicting the data in both centrality bins. It is slightly below the data for the highest p_T points in the 0–10% range. The Hybrid model is given by the dashed green band and predicts less suppression than other models. It successfully predicts the trend of the data within the sizeable uncertainty of the model. The model of Bianchi et al. is shown by the dashed blue band and underpredicts R_{AA} around $p_T = 10$ GeV. However, it also has the least amount of suppression at high p_T . Although the model manages to be compatible with the data points above approximately 20 GeV, it appears to predict

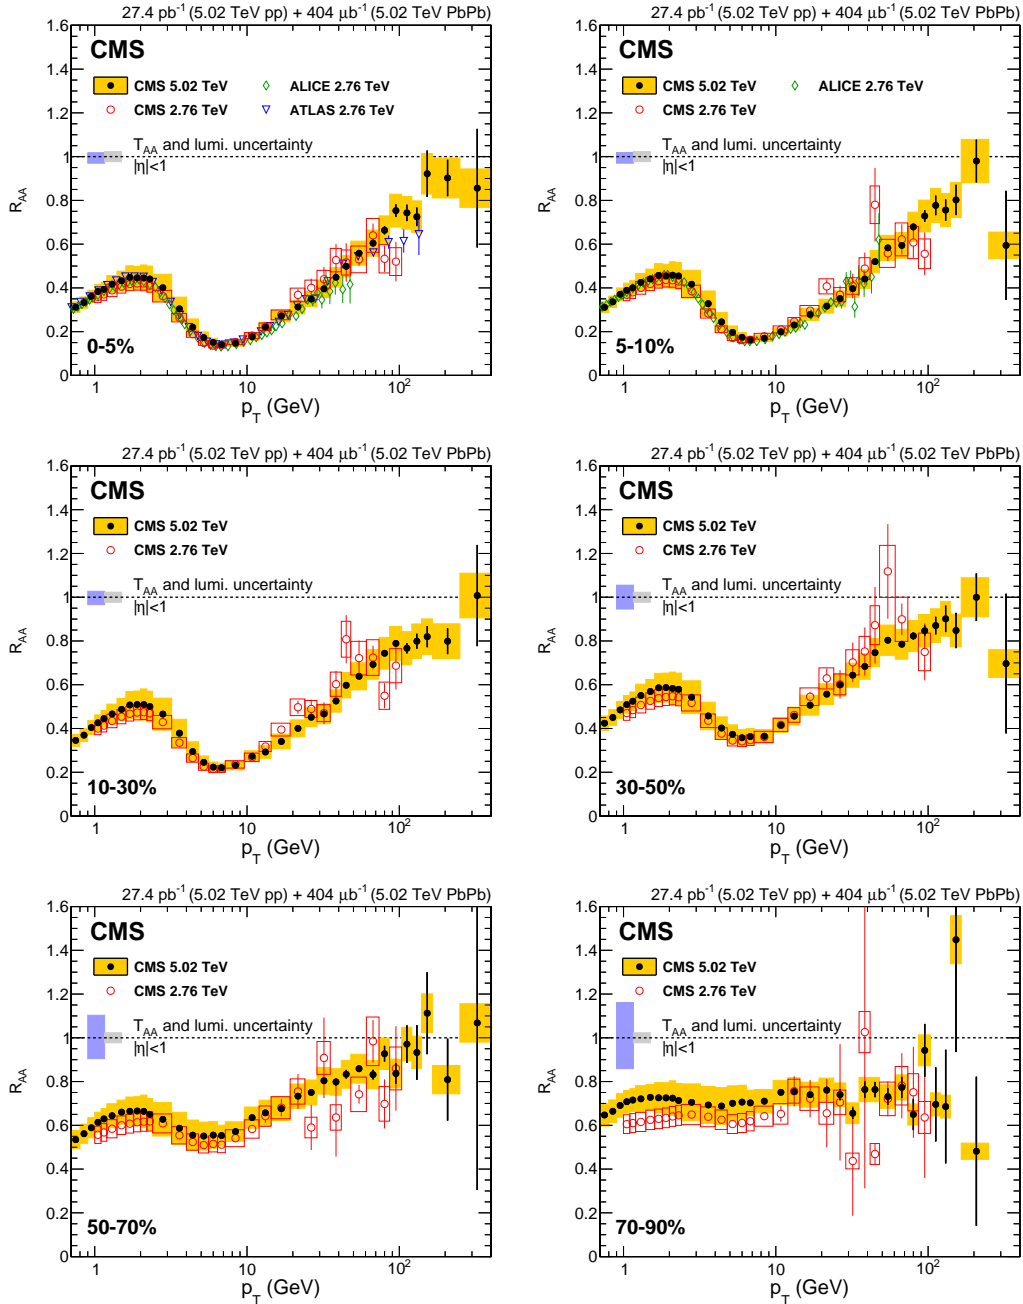


Figure 6-13: The nuclear modification factor in 5.02 TeV PbPb collisions, for six different centrality selections. Previous results at 2.76 TeV are shown by the red, blue, and green data points.

a rising slope that is slightly steeper than what is seen in the data. This model was provided after the release of a preliminary version of these results and therefore is not guaranteed to be a true prediction. The CUJET 3.0 model is shown by the dark blue

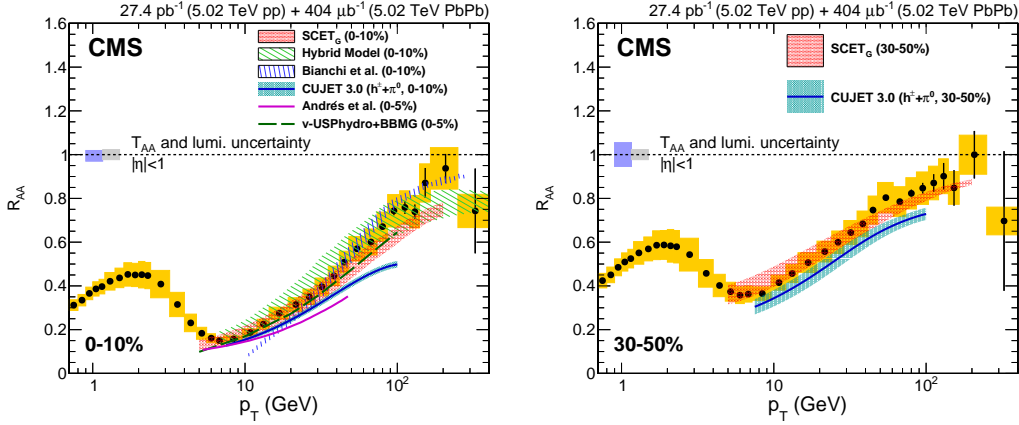


Figure 6-14: Comparison of the nuclear modification factor to six theoretical predictions in the 0–10% and 30–50% centrality bins. The predictions were provided by the authors of Refs. [220, 94, 113, 105, 117, 221].

line with a surrounding light blue band and is calculated as the R_{AA} for both charged hadrons and neutral pions. This model seems to predict a suppression that is stronger than what is seen in data, particularly in the high p_T region for 0–10% events. The prediction lies on the lower edge of the data for the 30–50% centrality range. It should also be noted that a bug was found in the results from this model after the release of the preliminary form of these results. Thus, this model can also not be guaranteed to be a true prediction. The model of Andrés et al. also seems to underpredict the data points, as shown by the purple line. This could partially be because this prediction is for the 0–5% centrality range, and not the 0–10% range. However, comparison of the data points for the 0–5% and 0–10% selections indicates that this is not a large effect. This model uses a quenching weight that is calibrated using 2.76 TeV data, so it is possible that the quenching weight would need to be recalculated using input data at a higher center of mass energy. Finally, the v-USPhydro+BBMG model is shown for 0–5% collisions by the dashed dark green line. The prediction generally does a good job of qualitatively describing the data, although it lies slightly below the data points for high p_T values. This small discrepancy would also be slightly alleviated by comparing to 0–5% events. In general a majority of the models seem to describe the region of $10 < p_T < 40$ GeV reasonably. This could potentially be because a




Figure 6-15: A summary of R_{AA} measurements and predictions for central collisions. The data and models come from References [208, 145, 220, 94, 113, 105, 117, 221, 129, 222, 128, 137, 132, 146, 142].

large amount of 2.76 TeV data already existed in this kinematic range, allowing for tuning of the models. As the probed p_T increases the separation between the models increases. Models that favor a slower increase of R_{AA} with p_T seem to be disfavored by this measurement. Put another way, this measurement indicates that the medium is rather transparent at high p_T , as compared to some predictions.

Figure. 6-15 is a compilation of R_{AA} results in central collisions for a variety of collision systems and beam energies. Various normalization uncertainties for the data presented here are not shown on the plot. The purple and pink points show SPS data

from WA98 [129, 222] and NA49 [128], respectively. These are $\sqrt{s_{NN}} = 17.3$ GeV collisions of PbPb. The WA98 result is for 0–7% centrality events and the NA49 result is for 0–5% collisions. The data agree reasonably and already show a fairly significant suppression at low p_T . The charged hadron data from STAR [137] for AuAu collisions at $\sqrt{s_{NN}} = 200$ GeV reveals an indication of a maximum and minimum structure, although it seems to be less suppressed than the LHC data. The data from PHENIX [132] for neutral pions was gathered for the energy and collision system as the STAR measurement. This measurement probes a higher p_T region than the STAR data, and starts to reveal the increasing slope that occurs after the R_{AA} minimum. There appears to be some tension between the STAR and PHENIX data around 5 GeV. The LHC data for 2.76 TeV PbPb collisions has already been discussed before. The models previously examined are also shown as the bands and lines in the region $p_T > 10$ GeV. It can be seen that this measurement continues a long tradition of R_{AA} measurements and provides an unprecedented view of the physics of the QGP across a huge range of particle p_T . It is also striking how most of the data above 10 GeV aligns, despite the differences in center of mass energy.

The inclusive charged particle R_{AA} is also useful for studies of the energy loss of heavy flavor quarks. Because the inclusive charged particle spectrum is dominated by the dynamics of light quarks and gluons, it can be used as a reference for comparing to heavy flavor particles. Figure 6-16 shows the current nuclear modification results from the CMS heavy flavor program, as compared to the results measured here for light quarks and gluons. Measurements of fully reconstructed D^0 mesons [223] are shown in green. Fully reconstructed B^\pm [224] mesons are shown in blue. A measurement of D^0 mesons originating from B decays is shown by the purple points [225]. Finally, a measurement of J/Ψ mesons coming from B decays is shown by the pink points [226]. It seems as if charm quarks have a modification that is similar to that of light quarks and gluons. However, beauty quarks are slightly less suppressed. Clearly the large p_T range of charged particles examined here is crucial to facilitate comparisons with the unprecedented kinematic range in which these heavy flavor measurements are performed.




Figure 6-16: A selection of heavy flavor measurements from CMS [223, 224, 225, 226], compared to the measurement of the nuclear modification of inclusive charged particles described here. The charged particle measurement is a crucial reference for interpreting these results.

6.3.2 pPb Nuclear Modification Factor

With a measured pp reference at 5.02 TeV, it is straightforward to update the R_{pA} result of Ref. [126] to not rely upon an interpolation or extrapolation procedure. The existing pPb spectrum, which was validated as part of the pPb fragmentation function analysis has a particle composition correction applied to it and a correction for the finite p_T bin width removed. This is done so that the data is treated in a similar fashion as the pp reference. A ratio is then taken and the relevant systematic uncertainties updated. The result can be seen as the black points with green boxes in Fig 6-17. The light green box around unity is the Glauber uncertainty for pPb collisions. The significant ‘bump’ structure in the p_T range 30–100 GeV is now largely gone. The data points seem to plateau at a value of around 1.2. At $p_T = 65$ GeV,




Figure 6-17: The updated result of R_{pA} using a measured 5.02 TeV pp reference, shown in by the points surrounded by green boxes. The inclusive R_{AA} for PbPb collisions is shown by the points with yellow boxes.

where the bump structure was largest when using an interpolated reference, $R_{pA} = 1.19 \pm 0.02$ (stat) $^{+0.13}_{-0.11}$ (syst). For comparison, the old result was $R_{pA}^* = 1.41 \pm 0.01$ (stat) $^{+0.20}_{-0.19}$ (syst). Thus, the replacement of the reference resulted in a change in the data points that was just over one standard deviation of the total uncertainty of the measurement. It should be noted that the uncertainty attributed to the pp reference was slightly smaller than the total uncertainty. This means that the change in the reference was closer to 1.5 standard deviations. The remaining enhancement is suggestive of nuclear anti-shadowing [78] and appears to be consistent with the measured jet R_{pA} seen in Fig. 1-17. The result is still a bit higher than the prediction of EPS09 [151] that is shown in Fig. 1-19. At p_T around 3 GeV the observed R_{pA} is similar to the previous result and is consistent with a slight enhancement caused

by Cronin enhancement and radial flow that has been observed in other dA and pA systems [84, 227]. For $p_T < 1$ GeV the values of R_{pA} are suppressed. Because low- p_T particles are not exclusively produced by hard processes, their production does not scale with N_{coll} . Therefore such a suppression indicates the breakdown of the inherent N_{coll} scaling hypothesis that is used when normalizing the nuclear modification factor.

The inclusive R_{AA} for PbPb collisions at the same energy has also been placed on Fig. 6-17. The suppression is still very strong for inclusive events, highlighting the effect that QGP production has on the yields of high- p_T hadrons. Indeed, it seems that even an R_{AA} of unity would still indicate a suppression with respect to the expectation from the initial state effects probed by the R_{pA} measurement. Furthermore, the enhancements seen for R_{pA} might be even greater in PbPb collisions, because of the presence of two lead ions. Clearly initial-state effects must be taken into account when attempting to fully interpret the results of PbPb collisions.

The results presented here have been largely confirmed by other experimental collaborations. Both the ATLAS [228] and ALICE [229] Collaborations have measured the PbPb R_{AA} at 5.02 TeV and obtained values that are compatible with this measurement. It should be noted, however, that the ATLAS data for central events tends to lie on the lower edge of the uncertainty band of this measurement at $p_T > 100$ GeV. Additional measurements of the ATLAS *jet* R_{AA} seem to plateau at values of around 0.6 for jets above approximately 300 GeV [35]. Thus, it appears that the charged particle R_{AA} is larger than the jet R_{AA} in these kinematic regions. This could imply that very high- p_T particles are only produced in a subset of jets which undergo less energy loss as compared to the inclusive jet population. There may be other high- p_T jets that lose a larger amount of energy and only have relatively soft constituent p_T . This interpretation is supported by a recent measurement of the jet fragmentation function in PbPb collisions at 5.02 TeV by ATLAS [34], where an *enhancement* is observed for jets containing a hadron that carries a majority of the jet's momentum. Data from the CMS Collaboration on the jet fragmentation function in 2.76 TeV PbPb collisions for the range $150 < p_T^{jet} < 300$ GeV also slightly hints at an enhancement [31]. However, these data would also be consistent with an unmodified fragmentation because

of large systematic uncertainties. With this interpretation, there would then seem to be some minor tension between existing jet and charged particle R_{AA} measurements at the LHC. Future studies with larger data sets and smaller systematic uncertainties may be able to better clarify the relationship between these two observables in this extremely high p_T region.

ATLAS has also updated their measurement of the charged particle R_{pA} , observing an enhancement around 1.2 that is similar to the data reported here [230].

Chapter 7

XeXe Nuclear Modification Factor

The delivery of a small amount of collisions of Xe¹²⁹ at 5.44 TeV in 2017 allows for a unique test of the system size dependence charged particle production. The radii of xenon and lead nuclei are ≈ 5.4 fm and ≈ 6.6 fm, respectively [200]. A simple Glauber model assuming that parton energy loss is linearly (quadratically) related to only the length of the path the parton traverse inside the QGP would conclude that the average energy loss in central XeXe collisions is 17% (31%) smaller than in PbPb. Although the actual dynamics of a heavy ion collision complicate this picture, XeXe collisions in principle allow a second method of varying the geometry of the collision (with the first method being a centrality selection). This was also the first time collisions at the LHC have involved a particle other than a proton or a lead ion and is therefore also an important proof-of-principle that other collision systems can be studied at LHC energies. The analysis detailed here follows Ref. [173] and focuses on a measurement of the nuclear modification factor of XeXe collisions, as well as comparisons to the same quantities in PbPb collisions. Because comparison with PbPb collisions is desired, much of the analysis procedure was kept the same as the PbPb analysis in order to facilitate cancellations of systematic uncertainties.

To allow comparisons between XeXe and PbPb collisions, a new ratio is defined:

$$R_{Pb}^{Xe}(p_T) = \frac{dN^{XeXe}/dp_T}{dN^{PbPb}/dp_T} \frac{T_{PbPb}}{T_{XeXe}}. \quad (7.1)$$

The PbPb data were gathered at $\sqrt{s_{NN}} = 5.02$ TeV, so these two collision systems have different center-of-mass energies. Deviations of R_{Pb}^{Xe} from expected values, after taking this energy difference into account, would indicate a different spectral modification between PbPb and XeXe events.

7.1 Analysis Procedure

The analysis procedure for XeXe collisions is very similar to the strategy for PbPb collisions described in Section 6.1. The XeXe data were the first heavy ion collisions to be collected with the CMS Phase 1 pixel upgrade. The tracking algorithms were therefore significantly updated to include quadruplet and Cellular Automaton seeding. However, because the strip detector was the same as it was in 2015, the properties of the final output tracks are fairly similar to the tracks in the 2015 PbPb data. Therefore, the same track selection requirements as the PbPb R_{AA} analysis are used. These selections are detailed in the last row of Table 5.1. Rescaling the χ^2 distributions was not needed for the XeXe data. Additionally, the XeXe analysis is taken down to 0.5 GeV instead of the 0.7 GeV threshold used for the PbPb analysis.

Unlike previous analyses, the tracking efficiency is evaluated as a function of p_T and centrality, while integrating over η and ϕ . This was done because it was possible to generate enough MC to full fill a two dimensional table without having to resort to using the previously used iterative method. This is mostly because it takes less time to generate the heavy ion background for XeXe collisions, as compared to PbPb collisions, and also because the use of PYTHIA 8 instead of PYTHIA 6 enables the generation of a more uniform distribution of the statistical power of the sample in the high p_T kinematic region. The tracking efficiency is similar to that of PbPb collisions, with a plateau value of around 70% for tracks having $p_T > 1$ GeV. The tracking efficiency as a function of p_T and centrality can be seen in Fig. 3-10. The misreconstruction rate is less than 1% for most of the p_T range studied, but spikes to up to 34% for tracks with p_T below 0.7 GeV in central events. This can be seen in Fig. 3-12 for the six centrality selections used here.

A similar strategy was used for the particle composition correction as in the PbPb R_{AA} analysis. The nominal tracking efficiency was calculated with PYTHIA signal embedded into MB HYDJET background. A sample of EPOS was then used as a cross check sample having a higher strangeness content. The data is corrected to lie halfway between the results given by both of these generators, keeping the same assumptions that were used in the PbPb R_{AA} analysis. The shape of the correction is similar to those shown in Fig. 6-10, although the total magnitude reaches a maximum of around 8% in central events. Above 10 GeV, the uncertainty on this correction is less than 3%.

The XeXe data was all gathered with a single unrescaled MB trigger; every collision passing this trigger was stored. This simplifies the analysis procedure immensely, as no trigger combination is needed. The normalization is done in a manner similar to the PbPb R_{AA} analysis, with the XeXe spectra being normalized by the total number of minimum bias events in a given centrality range. The T_{AA} numbers used for normalization are given in Table 4.2.

7.2 Extrapolation Procedure

Unfortunately, no pp reference data exists at $\sqrt{s} = 5.44$ TeV. Thus, a reference must be constructed from existing measurements. Here we choose to use a MC-based extrapolation procedure, using a previously measured 5.02 TeV pp reference as a baseline. This method has the advantage of being easily understandable and replaceable if an appropriate reference spectrum is measured in the future. It is expected to work better than the previous extrapolation attempts that were conducted to predict the pp spectrum at 5.02 TeV because of the much narrower energy range used. The extrapolated pp reference is calculated using the formula

$$\left(\frac{d\sigma_{5.44}^{pp}}{dp_T}\right)_{\text{Extrap.}} = \left(\frac{d\sigma_{5.44}^{pp}}{dp_T} \bigg/ \frac{d\sigma_{5.02}^{pp}}{dp_T}\right)_{\text{MC}} \left(\frac{d\sigma_{5.02}^{pp}}{dp_T}\right)_{\text{Data}}. \quad (7.2)$$

The data input is given by the pp points in Fig. 6-12. The MC-based extrapolation factor is calculated using PYTHIA 8 tune CUETP8M1. This MC generator was found to reproduce the measured charged particle spectra within experimental uncertainties at both 5.02 and 7 TeV [231] for most of the p_T range examined in this study. Some small discrepancies between the data and MC are observed for the ranges $p_T < 1$ GeV and around $p_T = 10$ GeV, but these are similar at both center of mass energies. Thus, these discrepancies are expected to largely cancel when a ratio of spectra is taken. The extrapolation factor calculated by PYTHIA is shown by the black points in Fig. 7-1. A polynomial fit of the form $a_0 + a_1x + a_2x^2 + a_3x^3 + a_4x^4$, where $x = \ln(p_T/1\text{GeV})$, is shown by the red line. The fit parameters are $a_0 = 1.04$, $a_1 = 2.56 \times 10^{-2}$, $a_2 = 1.27 \times 10^{-2}$, $a_3 = -4.72 \times 10^{-3}$ and $a_4 = 4.80 \times 10^{-4}$. This functional form gives a good phenomenological description of the data in the kinematic region of interest. The statistical uncertainty of the PYTHIA sample used to evaluate the fit is smaller than the markers for most of the data points. The extrapolation factor was checked using different MC models. EPOS tune LHC was used to validate the low p_T region, but can not produce high p_T points because it does not support a way to only generate hard scattering events. Instead, HERWIG++ tune EE5C was used to confirm the data points above $p_T = 20$ GeV. A linear fit to these points differs from the nominal fit by no more than 2%. Different fit functions, including sigmoid functions and ratios of Tsallis distributions [217] also agree with the nominal fit to within 1%.

The extrapolation procedure was also repeated using data at 7 TeV as a starting point. This extrapolation was found to be consistent with the nominal extrapolation. Other procedures, such as “relative placement” and x_T interpolation procedures [126] give results similar to the nominal reference values because they are tightly constrained by the input 5.02 TeV spectra.

Because the results for the nuclear modification factor do depend somewhat on this MC extrapolation procedure, we denote the measured quantity R_{AA}^* . Here the asterisk indicates the usage of this extrapolation procedure.




Figure 7-1: The extrapolation factor needed to calculate a pp reference at 5.44 TeV for comparison with XeXe data.

7.3 Systematic Uncertainties

A summary of the contributions from various systematic uncertainties is shown in Table. 7.1. The measurement of XeXe collisions includes large uncertainties on the total normalization, which are fully correlated across all p_T bins. These uncertainties are kept separate from uncertainties that are not fully correlated, so that the potential variation of the shape of the measured distributions can be evaluated regardless of the uncertainty in normalization.

For most of the kinematic range examined, the track misreconstruction rate is less than 1%. In the regions where this is true, the full misreconstruction rate is conservatively assigned as a systematic uncertainty. However, this rate is very large for tracks with $p_T < 0.7$ GeV in 0–10% events. For these events, the uncertainty related to the misreconstruction rate is determined using a data-driven study of the difference of track d_0 DCA significance between data and MC. A side band region that is between 25 and 30 standard deviations is first examined. Tracks in this region tend

Table 7.1: The systematic uncertainties related to the XeXe measurements reported here. The values quoted cover the centrality and p_T dependence of each uncertainty. They are separated into normalization uncertainties and all other systematic uncertainties.

| Sources | Uncertainty [%] | | |
|--|-----------------|------------|---------------|
| | XeXe Spectra | R_{AA}^* | R_{Pb}^{Xe} |
| Fraction of misreconstructed tracks | 0.1–16 | 0.1–16 | 0.1–5 |
| Particle species composition | 0.5–8 | 0.5–8 | 1–8 |
| Track selection | 3–6 | 3–6 | 5–7 |
| MC/data tracking efficiency difference | 5 | 2.0–6.4 | – |
| Tracking corrections | 0.5–2 | 0.5–2 | 1–5 |
| p_T resolution | 0.5 | 0.5 | – |
| Extrapolated pp reference | – | 4–9 | – |
| Trigger combination | – | – | 1 |
| Combined uncertainty | 7–18 | 6–18 | 6–11 |
| XeXe event selection efficiency | 0.3–26 | 0.3–26 | 0.3–26 |
| Glauber model uncertainty (T_{AA}) | – | 5–14 | 6–21 |
| pp reference luminosity | – | 2.3 | – |
| Combined normalization uncertainty | 0.3–26 | 6–30 | 6–33 |

to result from random combinations of hits and are therefore mostly misreconstructed. In general the fraction of tracks in this region is higher in data than in MC. This can be seen in the left panel of Fig. 7-2 for tracks in 0–5% events having p_T in the range 0.5–0.6 GeV. The relative fraction of misreconstructed tracks in the MC is then increased so that the total fraction of tracks in this side band region matches data. During this procedure, the shape of the distribution of misreconstructed and real tracks is assumed to stay the same in MC. Thus, the relative contribution of the blue line is increased with respect to the red line in Fig. 7-2. After this scaling procedure is done, the distributions are shown in the right panel of the same figure. The relative misreconstruction fraction is then calculated in the signal region requiring a $DCA < 3\sigma$ (shown by the green dotted lines). Comparison of the misreconstruction fraction in the unscaled and scaled MC gives a data-driven handle on how much the misreconstruction fraction can vary. The uncertainty in the misreconstruction rate for tracks under 0.7 GeV is calculated to be at most 16%. This uncertainty is not assumed to cancel in any ratios but is relatively small for the measurement of R_{Pb}^{Xe} because tracks below 0.7 GeV are only used for the charged particle spectra and R_{AA} .




Figure 7-2: A comparison of the d_0 DCA significance for XeXe data (black points) and MC (black line). The red and blue lines show the fraction of read and fake tracks making up the MC, respectively. The unscaled data is shown in the left panel, while MC that has had the fake fraction scale up to match data in a side band region is shown on the right panel.

The uncertainty related to the particle species correction is once again taken to be the full value of the correction, in a fashion similar to what was done for the PbPb R_{AA} analysis. It is a maximum of 8% around $p_T = 3$ GeV. Because this correction is evaluated using the same MC generators, this uncertainty is partially correlated between the XeXe and PbPb spectra measurements. This means it can be partially canceled in the ratio of R_{Pb}^{Xe} .

The track quality selections were varied twice, and the analysis was repeated with each variation. The first variation, referred to as selection ‘D’, required only 9 hits in the tracker, a normalized χ^2 less than 0.18, and an $E_T/p_T > 0.4$. Thus, it is a looser track selection. The second variation, selection ‘E’, required at least 11 hits, a normalized χ^2 less than 0.12, and an $E_T/p_T > 0.6$. The ratio of these different variations with respect to the nominal result is shown in the left and right panels of Fig. 7-3 for central and peripheral events, respectively. The results for selection ‘E’ tend to deviate more than those for selection ‘D’ and are greatest at lower p_T values. A systematic uncertainty of 6% is used to account for the observed discrepancies below $p_T = 20$ GeV. Above this threshold, the systematic uncertainty is only 3%.




Figure 7-3: The relative difference in the XeXe particle spectra after varying the track selections to tighter and looser cuts. The left panel shows 0–10% events, while the right panel shows 50–70% events. Refer to the text for explanations about selections D and E.

This systematic uncertainty is assumed to not cancel in any ratios.

The systematic uncertainty related to differences between the tracking efficiency in data and MC is evaluated in a way similar to the uncertainty in PbPb collisions. The logic is described in Section 6.2. The uncertainty for the XeXe spectra is 5%. Partial cancellation is assumed for measurements of R_{AA} , as previously described for the PbPb analysis. Because the detector occupancies for heavy ion collisions at the similar centralities are closer than compared to pp collisions, this uncertainty is assumed to cancel in measurements of R_{Pb}^{Xe} .

An uncertainty is related to the statistical uncertainty of the MC samples used to calculate the tracking efficiencies. This uncertainty is between 0.5 and 2% depending on the centrality and p_T range considered. The statistical uncertainties of the tracking efficiency can be seen by the light bands in Fig. 3-10. This does not cancel in any ratios. Measurements of R_{Pb}^{Xe} also include a contribution from the tracking correction of PbPb events as well, as described in Section 6.2.

The track p_T resolution is excellent in CMS, and a systematic uncertainty of 0.5% is applied to account for any bin to bin smearing related to the resolution. This is assumed to cancel for the measurement of R_{Pb}^{Xe} . The particle spectra are not unfolded to take into account resolution effects.

The uncertainty on the extrapolated pp reference is calculated by propagating the

uncertainty on the input 5.02 TeV reference spectra. To avoid double counting, some of the contributions to this uncertainty have been removed because they are included in other items. For example, this does not include the data vs MC efficiency difference uncertainty, because that contribution has been partially canceled out when accounting for the XeXe contribution. The remaining uncertainty from this contribution is between 4 and 9%.

No trigger combination is done in the XeXe analysis, but a 1% uncertainty is propagated to measurements of R_{Pb}^{Xe} because a trigger combination was needed to measure the PbPb spectra.

One challenge in analyzing the XeXe data is the relatively low event selection efficiency ($95 \pm 3\%$) of the MB trigger. This inefficiency is not localized to the most peripheral bins, and extends from 100% events to around 85%. Because of this inefficiency, it was decided to not analyze events in the 80–90% centrality region. Furthermore, the large uncertainty on this efficiency affects the centrality calibration. If the selection is actually more efficient, then events are systematically classified as being slightly more central than they really are. Likewise, if the selection is actually less efficient, then events are classified as being too peripheral. This shifting of centrality bins becomes more pronounced for peripheral centrality selections. Unlike the PbPb R_{AA} analysis, this shift is handled by creating three different centrality tables and repeating the analysis after varying the centrality calibration. In this way, the effects of this uncertainty are propagated to the final result as a systematic uncertainty. The net result is an uncertainty that is very small (0.3%) for 0–5% events, but quite large (26%) in the 70–80% centrality region. This uncertainty is correlated across all p_T bins.

The uncertainties on values calculated using the Glauber Model are given in Table 4.2. Uncertainties on Glauber model parameters are assumed to not cancel in ratios because they originate from different measurements of the nuclear structure of lead and xenon nuclei. In R_{Pb}^{Xe} the Glauber uncertainty can be as large as 21% in the most peripheral centrality selections. Finally, a correlated uncertainty of 2.3% on the total 5.02 TeV pp luminosity [206] applies for the measurement of the XeXe R_{AA} .

Graphical representations of the relative systematic uncertainty for the measurement of the XeXe spectra, XeXe R_{AA} , and R_{Pb}^{Xe} are shown in Fig. 7-4.

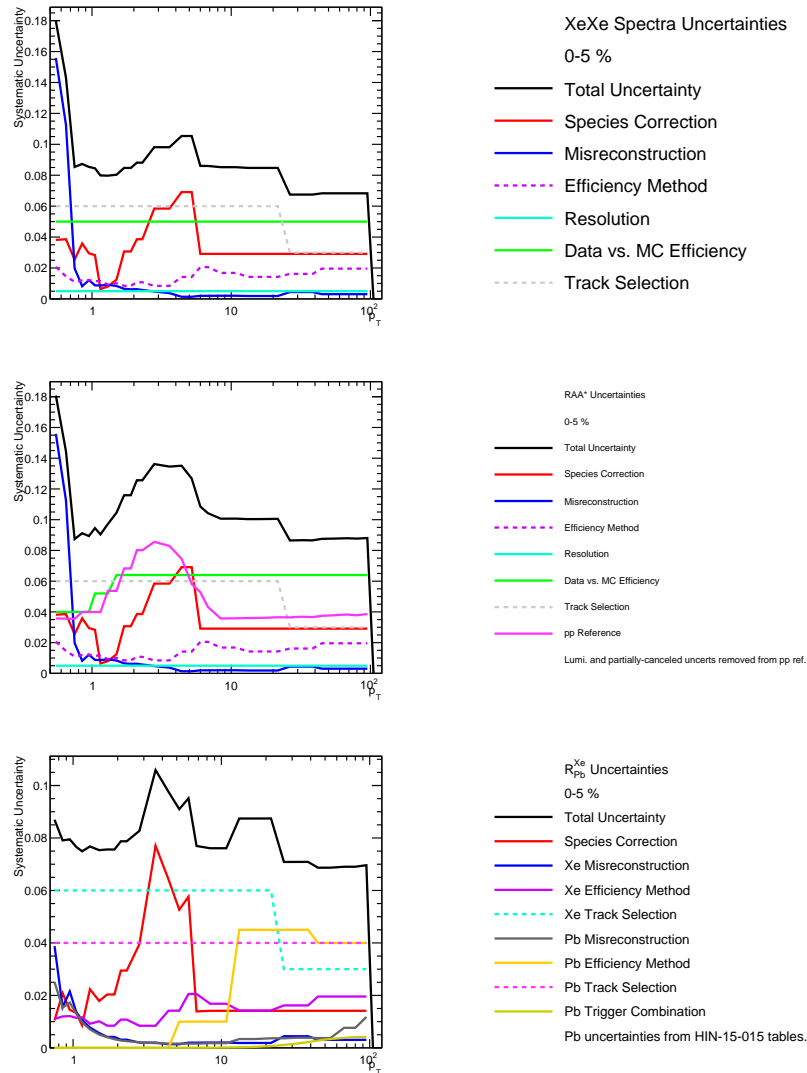


Figure 7-4: A graphical summary of the systematic uncertainties related to measurements of the 0–5% XeXe spectra (top), 0–5% XeXe R_{AA} (middle), and 0–5% R_{Pb}^{Xe} (bottom).

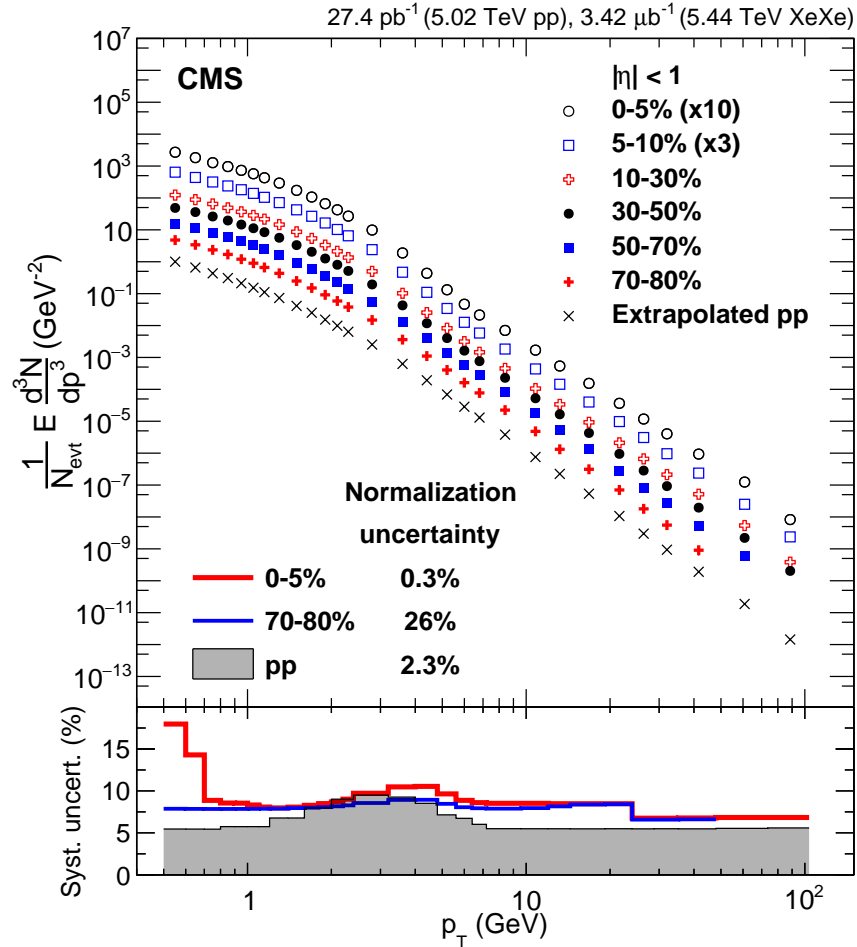


Figure 7-5: The charged particle spectra for 5.44 TeV XeXe and pp collisions. The pp data has been scaled by a constant factor of 70 mb so it can be directly compared to the XeXe measurement.

7.4 Results

The measured XeXe spectra at 5.44 TeV are shown for six different centrality selections in Fig. 7-5. The 0–5% and 5–10% data have been scaled by constant factors of 10 and 3, respectively, for clarity on the figure. The extrapolated 5.44 TeV pp reference is also shown on the same figure. To convert the pp data points from a differential invariant cross section to a differential invariant yield, a constant factor of 70 mb is used. The data points extend to around $p_T = 100$ GeV, although the peripheral centralities terminate earlier because of the limited statistical power of the gathered sample. Representative systematic uncertainties are shown in the bottom

panel for 0–5% and 70–80% XeXe collisions, as well as the pp reference. The shape of the uncertainties is similar to that observed for PbPb collisions, with a peak around 4 GeV and a plateau value at higher p_T . In central events the sharp increase at low p_T is related to the uncertainty on the misreconstruction rate. Values for the normalization uncertainties are also shown on the figure, as they are not included in the uncertainties shown in the bottom panel. The spectra all follow a roughly power-law behavior at high p_T , similar to what was observed for PbPb collisions.

The resulting values of R_{AA}^* in 5.44 TeV XeXe events are shown in Fig. 7-6. The black points shown the nuclear modification factor, and the black bars display the statistical uncertainties. The pink boxes represent the systematic uncertainties that are not fully correlated across of p_T bins. The red box around unity shows the fully correlated relative systematic uncertainties. The previously discussed 5.02 TeV PbPb R_{AA} measurement is shown by the open blue points and blue boxes. The dark blue box around unity also shows the related normalization uncertainties for this measurement. In general the normalization uncertainties for the XeXe data are larger because of the larger uncertainty in the event selection efficiency, and also because the input parameters to the Glauber model are less constrained. The XeXe data show a very strong suppression in 0–5% events. This suppression is most pronounced in the range $5 < p_T < 30$ GeV. The line shape is very similar to that seen in PbPb collisions, indicating that a hot medium is also created in XeXe collisions at the same centrality. In general, the XeXe and PbPb data show excellent agreement with each other in the range $p_T < 4$ GeV. In this region, the dynamics that govern particle production are more sensitive to phenomena like hydrodynamic flow [48], strangeness enhancement [66, 232], and the Cronin effect [83, 84] rather than parton energy loss. The similarity of the suppression between the two systems in this region indicates that the dynamics of these processes are not strongly affected by the difference in system size. This conclusion is supported by recent measurements of both radial flow [233, 234, 235] and the charged particle multiplicity [236, 237] for XeXe collisions. As the p_T range of interest increases, the XeXe data points are more likely to be on the upper edge of the PbPb data. In the range of $6 < p_T < 50$ GeV and for events

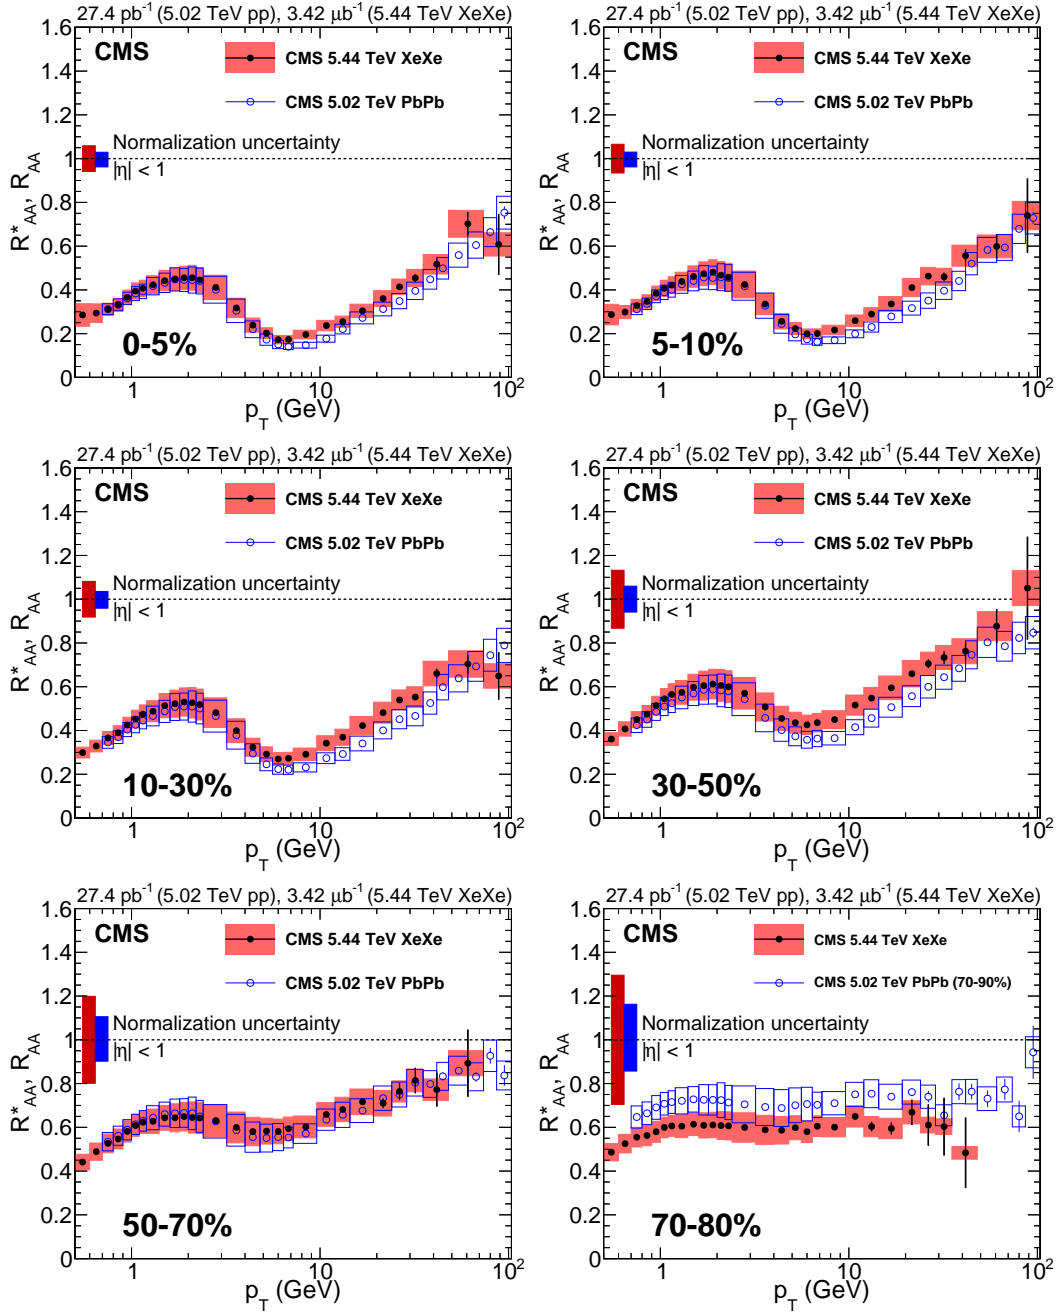


Figure 7-6: The nuclear modification factor in 5.44 TeV XeXe collisions, shown by the black points with pink boxes for six different centrality selections. The nuclear modification factor in 5.02 TeV PbPb collisions is shown by the blue points.

from 0–50%, the systematic uncertainties of the XeXe R_{AA}^* just barely touch those of the PbPb measurement. It should be noted that these measurements have a large correlated contribution to their uncertainties from the pp reference, since the 5.44 TeV

extrapolation is based on the 5.02 TeV measurement. When this correlation is taken into account, the separation between the two sets of data points would become more significant. This difference can be understood by the onset of parton energy loss as a dominant contributor to the modification of charged particle yields in this kinematic range. Because the strength of parton energy loss depends on the size of the QGP produced, and at a given centrality XeXe collisions produce a smaller medium, less suppression is expected. Interestingly, the slope of this suppression seems similar between XeXe and PbPb collisions; only the suppression magnitude seems strongly affected. This is consistent with the observation in PbPb collisions that the R_{AA} seems to continue increasing towards values that are close to unity, rather than flattening out.

As the collisions become more peripheral, the strength and oscillatory nature of the suppression decreases. In 50–70% events, the data for XeXe collisions lies directly on top of the PbPb data points. Remarkably, in 70–80% collisions this trend of the XeXe data crossing over the PbPb data continues. The XeXe data points are relatively flat at a value of around 0.6. It should be noted that there is a sizeable normalization uncertainty for this centrality selection, but the data still seems to be significantly below unity. As mentioned in Section 6.3.1, a suppression of the magnitude below unity is not expected to be caused by strong quenching effects in peripheral collisions. This is supported by recent results on dijet asymmetry in XeXe collisions [235]. It is now believed that this effect is the result of an experimental bias resulting from both the effects of underlying event fluctuations and the correlation of the yields of high p_T hadrons with forward activity. Because forward activity is used to calibrate the collision centrality, this correlation means that the centrality calibration can be biased.

This effect is illustrated in Fig. 7-7. The left panel shows the per-event differential yield of charged particles after scaling by N_{coll} in 5.44 TeV XeXe HYDJET collisions as a function of N_{coll} . Low p_T particles in the range $0.1 < p_T < 0.5$ GeV are shown by the red lines, while particles in the range from $4 < p_T < 10$ GeV are shown by black lines. The high p_T selection has been scaled by 2000 for clarity on the figure. The

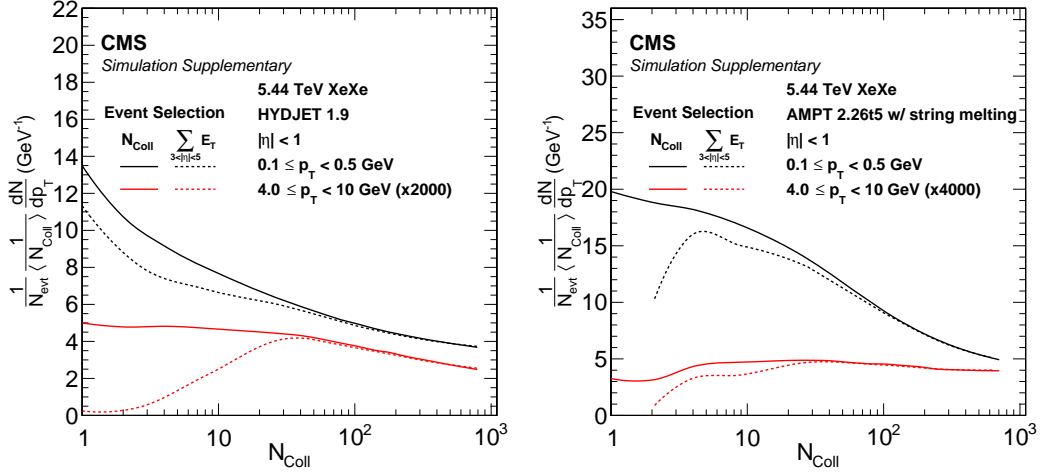


Figure 7-7: The N_{coll} scaled differential charged particle yields as a function of N_{coll} in two different MC generators for two different p_T selections. Two sets of lines are shown, one that uses the generator-level N_{coll} values (solid lines), and a second that uses the $\langle N_{coll} \rangle$ of a centrality bin calibrated in a fashion similar to data. The two curves differ for low values of N_{coll} , which correspond to peripheral events.

solid variations of the lines show the result calculated using the generator level N_{coll} information. The dashed lines show the results of a work flow in which the events are classified based on their centrality and then each centrality selection is divided by the $\langle N_{coll} \rangle$ calculated for the bin. The centrality is calculated using the same procedure as in data, with the HF energy deposits in the range $3 < |\eta| < 5$, and also includes detector effects. Thus, the dashed lines are equivalent to what is done for data. For both p_T ranges studied, the dashed and solid lines converge for large values of N_{coll} , indicating that the analysis procedure in data correctly reproduces the generator level result for central events. However, a large depletion in the dashed line relative to the solid line is observed for events with low values of N_{coll} . This corresponds to peripheral events; an N_{coll} value of 10 roughly corresponds to a 70% event. This indicates that the yield of observed charged hadrons in peripheral bins is affected by the centrality calibration procedure itself. It is well known that in pp collisions the production of high p_T jets, and therefore high p_T particles as well, is associated with an increase in the soft particle production [238, 239, 240, 241, 242]. Therefore heavy ion events with a high p_T particle may also have a systematically larger forward

energy for a given impact parameter, as compared to MB events. This means that the events will migrate into more central selections and peripheral selections will have a depletion of high p_T particles. Unfortunately this effect is difficult to correct for using a MC generator, because different models predict wildly different forward activity and different correlation between the forward and midrapidity production. This can be seen by comparing the left panel with right panel in Fig. 7-7. The left panel uses HYDJET, while the right panel uses AMPT 2.265 with string melting turned on (see Section 4.1 for a description of the physics of these generators). The depletion of high p_T particles in AMPT is much less pronounced than in HYDJET, but the relative number of soft particles produced per hard particle is much larger.

Figure 7-8 shows a quantification of the potential bias to R_{AA}^* that could result from this effect. The lines on the figure are calculated by taking a ratio of the dashed and solid lines shown in Fig 7-7 and displaying the result as a function of p_T rather than N_{coll} . The solid lines are the result from HYDJET and the dashed lines are from AMPT with string melting. For both generators, the potential bias is very small for central events, shown by the blue lines. Similarly, 30–50% events (shown in red) seem relatively unaffected by any potential bias. However, 70–80% collisions are suppressed by up to 20% and 50% in AMPT and HYDJET respectively. This suppression becomes larger as the p_T examined becomes larger. Fully understanding the exact magnitude of this bias will require a detailed quantification of the correlations between charged particle production in midrapidity and the forward region. Until this is done, physics conclusions based on ratios of central and peripheral collisions (so-called R_{CP} measurements) should be carefully scrutinized; there is large potential for a bias coming from the peripheral part of the measurement.

Recent studies in MC have used a simple model of overlaying PYTHIA events to create an ‘effective’ heavy ion event that contains no extra physics compared to a MB pp baseline. This model is known as HG-PYTHIA [243]. Examination of the nuclear modification factor for hard processes in this sample indicates a similar conclusion, i.e. a potential geometric bias exists when performing a centrality selection on the sample. This study observes a suppression factor of nearly 2 for 90% events even though




Figure 7-8: An estimation of the potential bias to R_{AA} caused by the effects of the centrality calibration procedure. Central events are largely unbiased, but peripheral events can be strongly affected. The magnitude of the bias depends on the MC model used in the calculation.

the model does not contain any physics process that is unique to heavy ion collisions. Recent measurements in 5.02 TeV PbPb collisions by the ALICE Collaboration for this peripheral centrality range have shown that this suppression qualitatively follows the prediction of this model [244]. One more interesting observation from this model is that a bias of up to 20% can be seen in very peripheral events before any detector effects are considered. This is because fluctuations in the event activity at a given impact parameter can cause significant centrality bin migration effects. This is a known issue when trying to calculate a centrality calibration for pA and dA collisions [245, 246, 247], but it is also clearly relevant for AA collisions which have a low number of participating nucleons.

It is also possible that this suppression is caused by a bias in the calculation of

the Glauber parameters for peripheral events. This could be the case if the spatial distribution of hard partons within the nucleus is narrower than expected [248]. Measurements of the yields of electroweak probes, which scale with N_{coll} and are not expected to be modified by the medium, seem to currently disfavor this interpretation for central collisions [86, 87, 88, 89, 90, 91, 92]. However, because the limited statistical power of available samples in the range from 70–100%, it is difficult to conclusively rule out this interpretation for peripheral events.

The difference in charged particle suppression is quantified using R_{Pb}^{Xe} for five centrality selections in Fig. 7-9. Once again, the black lines represent statistical uncertainties and the pink boxes show systematic uncertainties. The red box around unity shows the total normalization uncertainty of the measurement. Because the xenon and lead data are gathered at different center of mass energies, this ratio is not expected to be unity even in the absence of hot nuclear matter effects. The blue line shows the expected ratio of 5.44 and 5.02 TeV pp collisions, as calculated in PYTHIA. This is the same as the extrapolation factor shown in Fig. 7-1. Because no pp reference is used in this measurement, the systematic uncertainties are smaller than those for the measurement of R_{AA} . At low p_T , R_{Pb}^{Xe} is consistent with the PYTHIA expectation, indicating a similar amount of suppression in both systems. Starting around $p_T = 4$ GeV, a sharp rising trend is observed in the data, up to values of around 1.4. This value is $2-3\sigma$ above the PYTHIA expectation, indicating a significant difference in the amount of suppression between the two systems in this region. At even higher p_T the ratio tends to flatten out, while the PYTHIA expectation continues to rise. These observations are consistent with the interpretation that parton energy loss becomes more important at high p_T .

Comparing PbPb and XeXe collisions at the same centrality varies the amount of QGP produced while keeping the *shape* of the initial overlapping region between the two nuclei roughly constant. For example in 0% collisions, both overlapping regions will be roughly circular but the XeXe collision will have a smaller radius. It is also interesting to compare the systems when they have a similar overall size, as measured by N_{part} , while letting the shape of the overlapping region vary more.

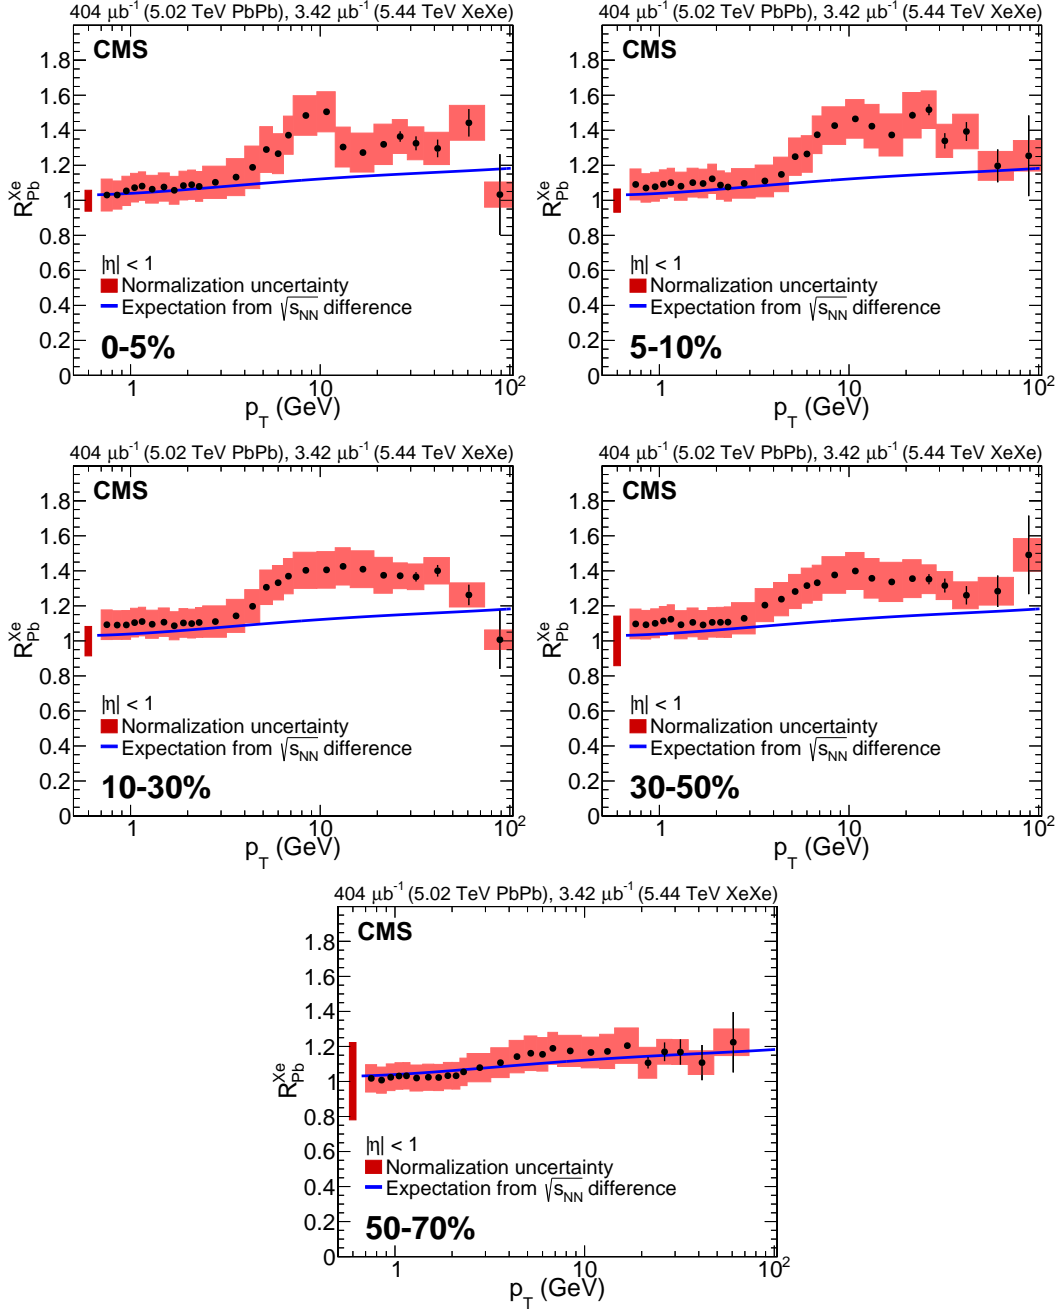


Figure 7-9: A measurement of R_{Pb}^{Xe} for five centrality selections. The blue line shows the unmodified expectation from the difference in center of mass energies between the numerator and denominator.

Such a comparison is shown in Fig. 7-10. Here, the values of both R_{AA} and R_{AA}^* in a given p_T range have been plotted versus the average N_{part} of each centrality selection. The left panel shows the R_{AA} values for charged particles in the range

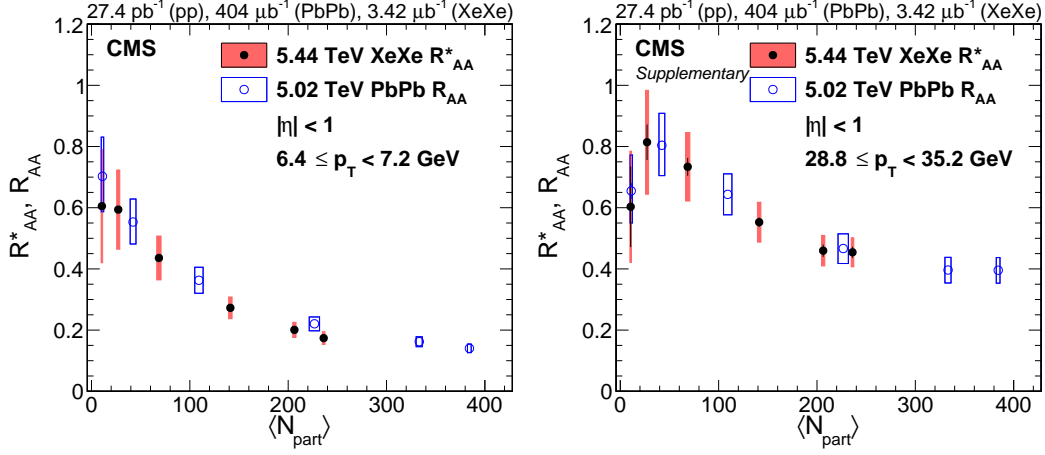


Figure 7-10: The values of R_{AA} and R_{AA}^* for PbPb and XeXe collisions as a function of N_{part} for two p_T selections. The left panel shows charged particles at the local minimum in R_{AA} , while the right panel shows particles that are in the middle of the rising high p_T trend.

$6.4 < p_T < 7.2$ GeV and the right panel shows the same quantity for particles having $28.8 < p_T < 35.2$ GeV. The first kinematic selection corresponds to particles at the minimum of the R_{AA} line shape and the second corresponds to particles that are in the middle of the ‘rising’ slope at high p_T . In both cases, the suppression of XeXe and PbPb collisions seems to roughly align with each other. In particular, the most central XeXe collisions have a suppression that is similar to 10–30% PbPb events with a nearly identical number of participating nucleons. It is also interesting to note that a significant suppression signal is expected for relatively small values of N_{part} . This means that it might be possible to study QGP properties using smaller ions that can be focused into more intense beams. Recent proposals for the future of the LHC heavy ion physics program are seriously investigating this possibility [249].

Two XeXe centrality selections have very similar N_{part} values to centrality bins in the PbPb measurement. The 0–5% XeXe collisions are close to 10–30% PbPb collisions, and 70–80% XeXe collisions are close to 70–90% PbPb collisions. Figure 7-11 shows the full R_{Pb}^{Xe} when taking ratios of these different centrality bins. Unlike the previously shown values of R_{Pb}^{Xe} (which compared the same centrality bins), no strong rising structure is observed around 5 GeV. Instead, the ratio is slightly below

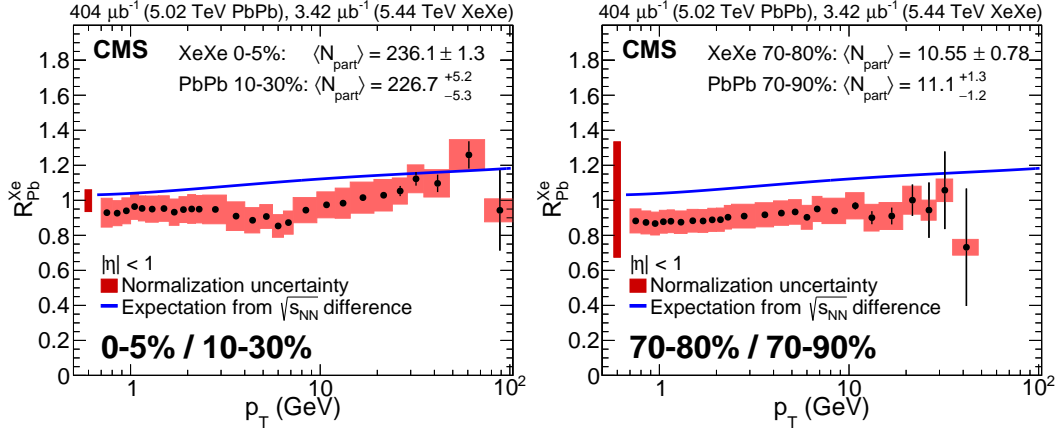


Figure 7-11: Measurements of R_{Pb}^{Xe} comparing different centrality bins for XeXe and PbPb collisions. The relevant centralities were selected to have similar N_{part} values. The left panel shows the result for central XeXe collisions, while the right panel shows the measurement in peripheral collisions.

the PYTHIA expectation. For central XeXe collisions, this deviation is only significant where R_{AA} has a minimum, around 4–10 GeV. For peripheral events the difference is not significant because of the very large normalization uncertainty. The trends seen when comparing collisions having similar N_{part} values further support the interpretation that differences in suppression found when comparing similar centrality selections originate from variations in the size of the QGP that is created in the collision.

The measured R_{AA}^* for 5.44 TeV XeXe collisions is compared to various models in Fig. 7-12. The left panel displays 0–10% collisions, while the right panel shows 30–50% collisions. The bottom inset shows a ratio of each model with the data. Systematic uncertainties are shown by the black boxes in the top part of the panel, but are shown as a gray band around unity in the bottom portion. The green lines show a prediction from the linear Boltzmann transport (LBT) model of jet quenching [112, 250]. This prediction uses the CLV_{isc} hydrodynamics package for medium evolution [252], with a freeze out temperature $T_c = 154$ MeV. It predicts a quadratic path-length dependence of energy loss in a static medium. The model lies on the upper edge of the measured R_{AA}^* in the range $20 < p_T < 60$ GeV for 0–10% events, but otherwise agrees with the data well. The orange band displays a model by Djordjevic that uses a dynamic energy loss formalism [103, 104]. The medium undergoes Bjorken expansion [253], and

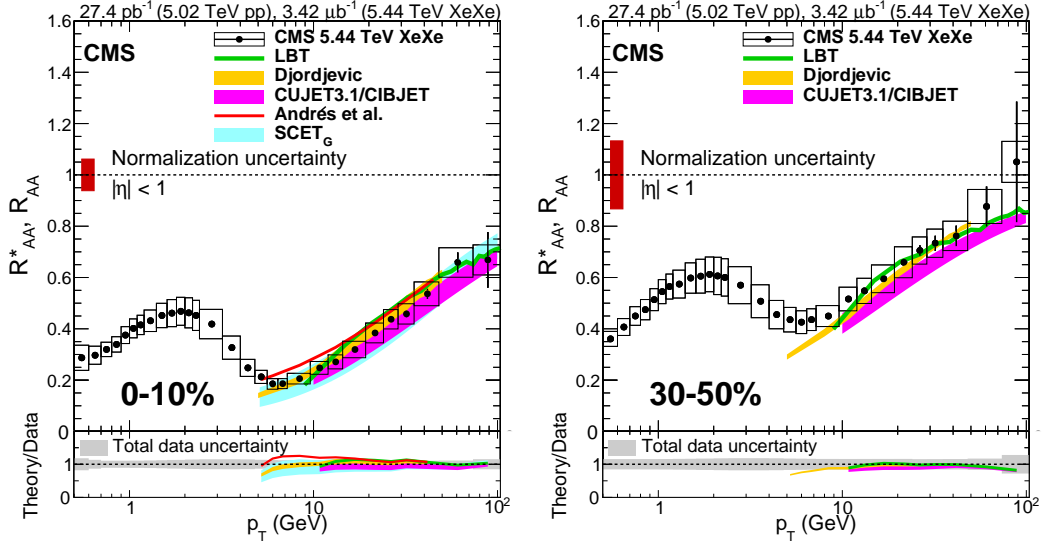


Figure 7-12: A comparison of R_{AA}^* for both 0–10% (left) and 30–50% (right) XeXe collisions with various theoretical models. The models are described in Refs. [112, 250, 103, 104, 105, 251, 117, 109, 220].

the path-length dependence is between linear and quadratic. The freeze out temperature is assumed to be $T_c = 160$ MeV. In general this model agrees with the data very nicely, although it undershoots the data in the region $p_T < 10$. This could be because of the increasing need to consider other effects such as a hydrodynamic flow in this p_T range. The magenta band shows the prediction of the CUJET3.1/CIBJET model. This model has two components. The first is a jet quenching model (CUJET 3.1) that includes suppression of the quark and gluon degrees of freedom and the emergence of chromo-magnetic monopole degrees of freedom [105]. The second component calculates the correlation between soft and hard flow harmonics on an event by event basis. This portion is done by the CIBJET framework [251]. The freeze out temperature is $T_c = 160$ MeV. The expected path-length dependence is roughly $L^{3/2}$, although the exact value of the exponent depends on the medium temperature, the amount of medium traversed, and the parton energy [105]. The model is able to describe the 0–10% data well, but lies on the lower edge of the uncertainty band for the 30–50% range. It is interesting to note that this model agrees with the XeXe measurement, while the CUJET 3.0 model seemed to systematically underpredict the PbPb R_{AA} . It

is possible that the model was tuned using the PbPb data in order to improve its performance for higher energy collisions. The model of Andrés et al. is shown by the red line in the 0–10% range. This uses a ‘quenching weights’ formalism for estimating the behavior of the medium transport coefficient, \hat{q} [117]. Medium evolution is performed using the EKRT event-by-event hydrodynamics package [254, 255]. The prediction lies on the upper edge of the data. This is interesting because a similar model prediction underestimated the R_{AA} in PbPb collisions. The effective quenching weight was likely refit to new data at a higher center of mass energy. The final prediction shown is from the SCET_G model [109, 220]. This model uses the IEBC hydrodynamics package to generate the background medium [256]. The freeze out temperature is assumed to be $T_c = 165$ MeV. The energy loss scales roughly as $N_{part}^{0.6}$ [257], which corresponds to a path length dependence of $L^{1.8}$ under the assumption $N_{part}^{1/3} \propto L$. The model is in good agreement with all the data in the range of its prediction. This model was also able to successfully predict the behavior of the PbPb R_{AA} . However, this model also has the largest uncertainty of all the models tested.

All of the models tested seem to be able to predict R_{AA}^* fairly well. This might be expected because of the close proximity in center of mass energy to readily available results for PbPb collisions. More stringent constraints on the models can be obtained if they are asked to reproduce the correct charged particles yields in both PbPb and XeXe collisions at the same time without tuning parameters for each system. Figure 7-13 compares the models to measurements of R_{Pb}^{Xe} in the 5–10% and 30–50% centrality selections. Some of the models only provided prediction for the 0–10% range, but they are still shown here because the 0–5% and 5–10% results for R_{Pb}^{Xe} in data are very similar. It should be noted that the experimental uncertainties for R_{Pb}^{Xe} are smaller because of cancellations in the ratio. Uncertainties in the theoretical models also seem to be smaller. This is because the uncertainty on the total normalization caused by the pp spectrum is removed. Thus, using this observable is more constraining than simply looking at R_{AA} .

The LBT model is unable to correctly predict the behavior of R_{Pb}^{Xe} , despite reproducing the XeXe result for R_{AA}^* . In both centrality ranges, it shows a large excess

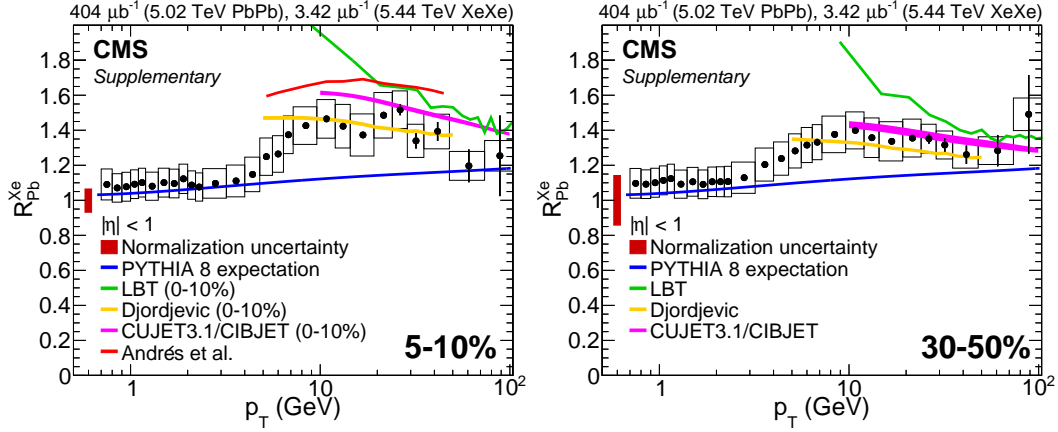


Figure 7-13: A comparison of four theoretical models to measurements of R_{Pb}^{Xe} . The 5–10% centrality range is shown on the left, while the 30–50% range is shown on the right. The models come from Refs. [112, 250, 103, 104, 105, 251, 117]

in the range $10 < p_T < 30$. At higher p_T the result seems to converge towards the data. This implies that for this model the predicted PbPb R_{AA} would be too low, and that the path-length dependence of this model is too large. The model from Djordjevic successfully predicts both the XeXe R_{AA}^* and R_{Pb}^{Xe} . The CUJET 3.1 model slightly over predicts the ratio in the 5–10% centrality range. This is consistent with the observation that it can predict the XeXe R_{AA}^* while under predicting the PbPb R_{AA} . This could indicate that the path length dependence of the model (of around $L^{1.5}$) is also too large. In the 30–50% events, the model predicts the R_{Pb}^{Xe} well, but this is most likely because it under estimates both the XeXe R_{AA}^* and the PbPb R_{AA} , causing a cancellation effect. Finally the prediction of the Andrés et al. model is shown for 5–10% collisions. It is slightly above the data, consistent with the slight over prediction of the XeXe data and slight under estimation of the PbPb R_{AA} . This model may improve if it is retuned using the new data to fit an updated value of the effective quenching weight.

7.5 Discussion

In summary, the LBT and CUJET models, which have energy loss proportional to L^2 and $\approx L^{1.5}$ seem to be unable to predict the R_{AA} of XeXe and PbPb collisions at the same time. The Djordjevic model, which has a dynamic energy loss that scales between L and L^2 seems to be able to predict the data well. Finally, the SCET_G model, which has a length-dependence of $\approx L^{1.8}$ seems to be able to predict R_{AA} for both XeXe and PbPb collisions, but has a large uncertainty and has not provided a direction comparison to R_{Pb}^{Xe} . The models shown here are not the only studies that have used this data, however. Recent phenomenological studies of the published CMS XeXe and PbPb data [98] using a BDMPS formalism have attempted to extract the average energy loss of a parton ($\langle\bar{\epsilon}\rangle$) for 2.76 TeV and 5.02 TeV PbPb collisions, as well as 5.44 TeV XeXe collisions. Their result is shown in Fig. 7-14. They find the average energy loss in XeXe collisions to be around 20% lower than that of 5.02 TeV PbPb collisions. This corresponds to a path-length dependence of roughly $L^{1.3\pm 0.5}$. Although this calculation makes many simplifying assumptions and has a large uncertainty, it illustrates that global analysis of data from many different center of mass energies and collision systems is now a viable option for attempting to extract the properties of the QGP. Based on the available information, it appears that the path-length dependence of energy loss should lie somewhere between a linear and quadratic relationship. This is consistent with parton energy loss being a mixture of both radiative and collisional processes.

Theoretical models have become increasingly complex over the years. Each model here uses a different energy loss kernel, a different type of hydrodynamic background evolution, and a different freeze out temperature. As can be seen by the comparisons of models to the XeXe R_{AA}^* and R_{Pb}^{Xe} , it appears that comparisons to multiple data sets are now needed to extract useful information about the QGP from theoretical models, as most models have enough free parameters that can be tuned to fit any one distribution reasonably well. Charged particle observables are unique in that they can be measured with high precision by essentially every heavy ion experiment.




Figure 7-14: The average energy loss of a fast parton, as calculated in a phenomenological study using the data from this measurement. The figure is reproduced from Ref. [98].

Therefore a huge amount of charged particle data is available for examination. A concerted effort to do a global fit of all the data with theoretical models that have similar hydrodynamic profiles, freeze out temperatures, etc. is needed to fully extract the wealth of information in these measurements. Such an effort is currently underway in the form of the JETSCAPE framework [258]. This software package intends to have powerful statistical analysis tools that will allow comparisons of theoretical models to a host of charged particle and jet observables at the same time. Preliminary results using only 2.76 TeV data, as shown in Fig 7-15, seem promising [259]. In the figure, three energy loss models having the same hydrodynamic background and hadronization code are compared to 2.76 TeV CMS results. The inclusion of the data measured in this thesis should be possible with the release of version 2 of this framework and will hopefully allow greater separation power between theoretical models.

Similar measurements for the XeXe R_{AA} were performed by the ALICE [260] and ATLAS [235] Collaborations at the same time as this analysis. In general the re-

sults from these two measurements, as well as the one presented here, agree quite impressively. Similar conclusions concerning the difference in nuclear modification factors between XeXe and PbPb collisions were reached independently by all three collaborations. Perhaps the only major point of tension that remains is the measurements in peripheral bins. For the 70–80% range, the CMS data lies between 1.5 and 2 standard deviations away from the other two results. However, as already mentioned this is likely due to experimental biases which can differ from detector to detector. Analyzing these peripheral events reliably remains a challenge to be addressed.

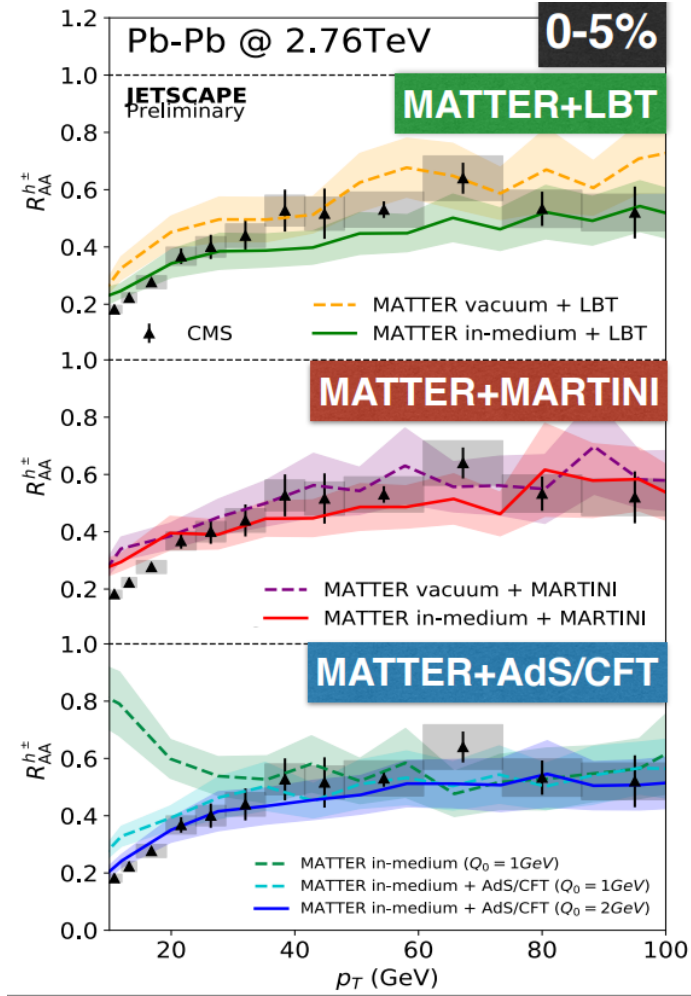


Figure 7-15: An example of the tools being developed by the JETSCAPE Collaboration to compare charged hadron spectra to theoretical models that have similar parameters. Here three energy loss models having the same hydrodynamic background and hadronization code are compared to 2.76 TeV CMS results. Figure taken from Ref. [259].

Chapter 8

Conclusions

The CMS detector provides an excellent basis for performing charged particle measurements; it has a very fine granularity tracking system and a large magnetic field that results in precise track p_T reconstruction. Track reconstruction in the heavy ion environment is challenging, and a large variety of tracking algorithms are employed to deal with the large detector occupancies observed. These tracking procedures were updated to accommodate upgrades to the CMS pixel tracker. Detailed studies in both data and simulation were performed to understand the performance of the tracking as a function of p_T and collision centrality.

A set of charged particle spectra measurements spanning a wide range of system sizes has been presented. The first measurement examines the fragmentation function of jets in 5.02 TeV pPb collisions and observes no significant modification with respect to an interpolated pp reference. This measurement confirmed concerns about an already-existing measurement of the charged particle nuclear modification factor in pPb collisions (R_{pA}) which implied a modified fragmentation function. It was found that different assumptions about the pp reference can drastically affect the conclusions about potential cold nuclear matter effects. This study helped motivate the need for the LHC to deliver a dedicated reference run of 5.02 TeV pp data in 2015. These data were analyzed and the measurement of the R_{pA} was updated. After this update, R_{pA} has a much flatter shape and plateaus at a value of around 1.2 in the high p_T region. This is indicative of nuclear anti-shadowing effects but is still slightly higher

than some nPDF predictions.

Data from 5.02 TeV PbPb collisions were also analyzed in six bins of collision centrality across an unprecedentedly large range of charged particle p_T . The nuclear modification factor (R_{AA}) was constructed and found to be strongly suppressed, especially in the region around $p_T = 7$ GeV. This suppression was found to be very similar to that of 2.76 TeV PbPb collisions. However, because the charged particle spectra are much flatter at 5.02 TeV, a similar R_{AA} value implies a larger amount of energy loss at the higher collision energy. The values of R_{AA} were also found to continue increasing above $p_T > 80$, up to values that are around 0.8 at $p_T = 200$ GeV. This rules out some previous speculation that R_{AA} would plateau at a value of around 0.6. This measurement also serves as a useful reference for studies of the energy loss of heavy flavor particles.

Finally, a measurement of the nuclear modification factor with extrapolated pp reference (R_{AA}^*) was performed in 5.44 TeV XeXe collisions. Care was taken to perform the analysis in a manner that would allow the result to be compared to the previous PbPb results. The XeXe R_{AA}^* was found to be less suppressed than the PbPb R_{AA} in the kinematic region where parton energy loss is a dominant effect. This implies that the difference in the two results originates from the different system size when comparing XeXe and PbPb collisions at the same centrality. Comparisons between different centrality selections that have a similar number of participating nucleons further confirm this conclusion.

Experimental difficulties are still found in very peripheral heavy ion collisions. A significant suppression of both the PbPb and XeXe R_{AA} is observed but is not expected to be caused by large parton energy loss. This suppression results most likely from a geometric bias related to the correlation of particle yields in the signal region and the forward region used for the centrality calibration. It is difficult to correct this effect with a Monte Carlo generator, because the activity in the forward region can vary wildly between different models. Until this correlation can be explored in more detail, it is advisable to refrain from using peripheral events as a reference, as is commonly done in central-peripheral ratios (R^{CP}).

Comparisons of the R_{AA} results with theoretical models were performed. Models are able to reliably predict the PbPb and XeXe R_{AA} but have more trouble describing both collision systems simultaneously. A ratio of the two systems, R_{Pb}^{Xe} , is found to be more constraining because of smaller systematic uncertainties on both experimental and theoretical fronts. It is found that models having a fractional parton energy loss that is between linear and quadratic in the path-length dependence describe the data best. These data are valuable inputs to theoretical efforts (such as the ongoing JETSCAPE project [258]) that perform global fits to constrain the properties of the QGP.

These measurements continue a long tradition of using charged particle spectra as one of the first ‘bellwether’ observations in new collision systems and energies. Although modern analysis methods are becoming increasingly more complicated and using multiple combinations of different physics objects, simple spectra measurements still have their place as a touchstone that can be used to calibrate theoretical predictions. Furthermore, charged particle measurements are some of the only measurements that can be reliably reproduced by more than two experiments and serve as a useful guidepost for validating different experimental approaches to tracking in the heavy ion environment.

The LHC Run 2 era has given rise to a host of novel energies, collision systems, and unprecedented luminosities. Future heavy ion data collection periods at the LHC will most likely not involve significant progression on either the energy or luminosity frontiers because of machine limitations. Nonetheless, exciting new opportunities still exist with planned detector upgrades and more potential collision systems. Significant progress can still be made by reducing systematic uncertainties, improving analysis techniques, and improving theoretical methods. Although much remains to be clarified about the properties of QCD in heavy ion events, it is impressive how rich the dynamics of the QGP have already proven to be.

Bibliography

- [1] D. Dominguez, “Standard Model. Le modèle standard”,. General Photo.
- [2] M. K. Gaillard, P. D. Grannis, and F. J. Sciulli, “The Standard model of particle physics”, *Rev. Mod. Phys.* **71** (1999) S96–S111, doi:10.1103/RevModPhys.71.S96, arXiv:hep-ph/9812285.
- [3] Particle Data Group Collaboration, “Review of Particle Physics”, *Chin. Phys.* **C38** (2014) 090001, doi:10.1088/1674-1137/38/9/090001.
- [4] H. D. Politzer, “Reliable Perturbative Results for Strong Interactions?”, *Phys. Rev. Lett.* **30** (1973) 1346–1349, doi:10.1103/PhysRevLett.30.1346. [,274(1973)].
- [5] D. J. Gross and F. Wilczek, “Ultraviolet Behavior of Nonabelian Gauge Theories”, *Phys. Rev. Lett.* **30** (1973) 1343–1346, doi:10.1103/PhysRevLett.30.1343. [,271(1973)].
- [6] S. Bethke, “The 2009 World Average of $\alpha(s)$ ”, *Eur. Phys. J.* **C64** (2009) 689–703, doi:10.1140/epjc/s10052-009-1173-1, arXiv:0908.1135. [,111(2009)].
- [7] Y. Koma and M. Koma, “Spin-dependent potentials from lattice qcd”, *Nuclear Physics B* **769** (2007), no. 1, 79 – 107, doi:https://doi.org/10.1016/j.nuclphysb.2007.01.033.
- [8] NNPDF Collaboration, “Parton distributions for the LHC Run II”, *JHEP* **04** (2015) 040, doi:10.1007/JHEP04(2015)040, arXiv:1410.8849.
- [9] S. Dulat et al., “New parton distribution functions from a global analysis of quantum chromodynamics”, *Phys. Rev.* **D93** (2016), no. 3, 033006, doi:10.1103/PhysRevD.93.033006, arXiv:1506.07443.
- [10] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt, “Parton distributions for the LHC”, *Eur. Phys. J.* **C63** (2009) 189–285, doi:10.1140/epjc/s10052-009-1072-5, arXiv:0901.0002.
- [11] G. Altarelli and G. Parisi, “Asymptotic Freedom in Parton Language”, *Nucl. Phys.* **B126** (1977) 298–318, doi:10.1016/0550-3213(77)90384-4.

- [12] M. E. Peskin and D. V. Schroeder, “An Introduction to quantum field theory”. Addison-Wesley, Reading, USA, 1995. ISBN 9780201503975, 0201503972,
- [13] ALEPH Collaboration, “Quark and gluon jet properties in symmetric three jet events”, *Phys. Lett.* **B384** (1996) 353–364, doi:10.1016/0370-2693(96)00849-0.
- [14] Y. L. Dokshitzer, “Calculation of the Structure Functions for Deep Inelastic Scattering and e+ e- Annihilation by Perturbation Theory in Quantum Chromodynamics.”, *Sov. Phys. JETP* **46** (1977) 641–653. [Zh. Eksp. Teor. Fiz.73,1216(1977)].
- [15] V. N. Gribov and L. N. Lipatov, “e+ e- pair annihilation and deep inelastic e p scattering in perturbation theory”, *Sov. J. Nucl. Phys.* **15** (1972) 675–684. [Yad. Fiz.15,1218(1972)].
- [16] “A new state of matter created at CERN”. <https://newstate-matter.web.cern.ch/newstate-matter/index.html>. Accessed: 2019-03-22.
- [17] M. Gyulassy, “The QGP discovered at RHIC”, in *Structure and dynamics of elementary matter. Proceedings, NATO Advanced Study Institute, Camyuva-Kemer, Turkey, September 22-October 2, 2003*, pp. 159–182. 2004. arXiv:nucl-th/0403032.
- [18] A. Bazavov et al., “Equation of state and QCD transition at finite temperature”, *Phys. Rev.* **D80** (2009) 014504, doi:10.1103/PhysRevD.80.014504, arXiv:0903.4379.
- [19] “The Frontiers of Nuclear Science, A Long Range Plan”, arXiv:0809.3137.
- [20] C. Shen, “The standard model for relativistic heavy-ion collisions and electromagnetic tomography”. PhD thesis, Ohio State U., 2014-07-25.
- [21] D. d’Enterria, “Jet quenching”, *Landolt-Bornstein* **23** (2010) 471, doi:10.1007/978-3-642-01539-7_16, arXiv:0902.2011.
- [22] Y. Mehtar-Tani, C. A. Salgado, and K. Tywoniuk, “Jets in QCD Media: From Color Coherence to Decoherence”, *Phys. Lett.* **B707** (2012) 156–159, doi:10.1016/j.physletb.2011.12.042, arXiv:1102.4317.
- [23] J. D. Bjorken, “Energy Loss of Energetic Partons in Quark - Gluon Plasma: Possible Extinction of High p(t) Jets in Hadron - Hadron Collisions”,.
- [24] STAR Collaboration, “Disappearance of back-to-back high p_T hadron correlations in central Au+Au collisions at $\sqrt{s_{NN}} = 200\text{-GeV}$ ”, *Phys. Rev. Lett.* **90** (2003) 082302, doi:10.1103/PhysRevLett.90.082302, arXiv:nucl-ex/0210033.

- [25] STAR Collaboration, “Measurements of jet quenching with semi-inclusive hadron+jet distributions in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV”, *Phys. Rev. C* **C96** (2017), no. 2, 024905, doi:10.1103/PhysRevC.96.024905, arXiv:1702.01108.
- [26] PHENIX Collaboration, “Trends in Yield and Azimuthal Shape Modification in Dihadron Correlations in Relativistic Heavy Ion Collisions”, *Phys. Rev. Lett.* **104** (2010) 252301, doi:10.1103/PhysRevLett.104.252301, arXiv:1002.1077.
- [27] ALICE Collaboration, “Particle-yield modification in jet-like azimuthal di-hadron correlations in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Phys. Rev. Lett.* **108** (2012) 092301, doi:10.1103/PhysRevLett.108.092301, arXiv:1110.0121.
- [28] CMS Collaboration, “Observation and studies of jet quenching in PbPb collisions at nucleon-nucleon center-of-mass energy = 2.76 TeV”, *Phys. Rev. C* **C84** (2011) 024906, doi:10.1103/PhysRevC.84.024906, arXiv:1102.1957.
- [29] ATLAS Collaboration, “Observation of a Centrality-Dependent Dijet Asymmetry in Lead-Lead Collisions at $\sqrt{s_{NN}} = 2.77$ TeV with the ATLAS Detector at the LHC”, *Phys. Rev. Lett.* **105** (2010) 252303, doi:10.1103/PhysRevLett.105.252303, arXiv:1011.6182.
- [30] ATLAS Collaboration, “Measurement of jet p_T correlations in Pb+Pb and pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV with the ATLAS detector”, *Phys. Lett. B* **774** (2017) 379–402, doi:10.1016/j.physletb.2017.09.078, arXiv:1706.09363.
- [31] CMS Collaboration, “Measurement of jet fragmentation in PbPb and pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Phys. Rev. C* **C90** (2014), no. 2, 024908, doi:10.1103/PhysRevC.90.024908, arXiv:1406.0932.
- [32] CMS Collaboration, “Measurement of inclusive jet nuclear modification factor in pPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV with CMS”,.
- [33] CMS Collaboration, “Modification of jet shapes in PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Phys. Lett. B* **730** (2014) 243–263, doi:10.1016/j.physletb.2014.01.042, arXiv:1310.0878.
- [34] ATLAS Collaboration, “Measurement of jet fragmentation in Pb+Pb and pp collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the ATLAS detector”, *Phys. Rev. C* **C98** (2018), no. 2, 024908, doi:10.1103/PhysRevC.98.024908, arXiv:1805.05424.
- [35] ATLAS Collaboration, “Measurement of the nuclear modification factor for inclusive jets in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the ATLAS detector”, *Phys. Lett. B* **790** (2019) 108–128, doi:10.1016/j.physletb.2018.10.076, arXiv:1805.05635.

- [36] CMS Collaboration, “Jet shapes of isolated photon-tagged jets in PbPb and pp collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Submitted to: Phys. Rev. Lett.* (2018) arXiv:1809.08602.
- [37] CMS Collaboration, “Study of jet quenching with isolated-photon+jet correlations in PbPb and pp collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Phys. Lett. B* **785** (2018) 14–39, doi:10.1016/j.physletb.2018.07.061, arXiv:1711.09738.
- [38] CMS Collaboration, “Study of Jet Quenching with Z + jet Correlations in Pb-Pb and pp Collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Phys. Rev. Lett.* **119** (2017), no. 8, 082301, doi:10.1103/PhysRevLett.119.082301, arXiv:1702.01060.
- [39] CMS Collaboration, “Observation of Medium-Induced Modifications of Jet Fragmentation in Pb-Pb Collisions at $\sqrt{s_{NN}} = 5.02$ TeV Using Isolated Photon-Tagged Jets”, *Phys. Rev. Lett.* **121** (2018), no. 24, 242301, doi:10.1103/PhysRevLett.121.242301, arXiv:1801.04895.
- [40] ATLAS Collaboration, “Comparison of fragmentation functions for light-quark- and gluon-dominated jets from pp and Pb+Pb collisions in ATLAS”, arXiv:1902.10007.
- [41] ATLAS Collaboration, “Measurement of photon–jet transverse momentum correlations in 5.02 TeV Pb + Pb and pp collisions with ATLAS”, *Phys. Lett. B* **789** (2019) 167–190, doi:10.1016/j.physletb.2018.12.023, arXiv:1809.07280.
- [42] CMS Collaboration, “Studies of dijet transverse momentum balance and pseudorapidity distributions in pPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Eur. Phys. J. C* **74** (2014), no. 7, 2951, doi:10.1140/epjc/s10052-014-2951-y, arXiv:1401.4433.
- [43] ATLAS Collaboration, “Centrality and rapidity dependence of inclusive jet production in $\sqrt{s_{NN}} = 5.02$ TeV proton-lead collisions with the ATLAS detector”, *Phys. Lett. B* **748** (2015) 392–413, doi:10.1016/j.physletb.2015.07.023, arXiv:1412.4092.
- [44] CMS Collaboration, “Measurement of inclusive jet production and nuclear modifications in pPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Eur. Phys. J. C* **76** (2016), no. 7, 372, doi:10.1140/epjc/s10052-016-4205-7, arXiv:1601.02001.
- [45] STAR Collaboration, “Azimuthal anisotropy in Au+Au collisions at $s(NN)^{(1/2)} = 200$ -GeV”, *Phys. Rev. C* **72** (2005) 014904, doi:10.1103/PhysRevC.72.014904, arXiv:nucl-ex/0409033.
- [46] STAR Collaboration, “Azimuthal anisotropy in Cu+Au collisions at $\sqrt{s_{NN}} = 200$ GeV”, *Phys. Rev. C* **98** (2018), no. 1, 014915, doi:10.1103/PhysRevC.98.014915, arXiv:1712.01332.

- [47] PHENIX Collaboration, “Flow measurements via two particle azimuthal correlations in Au+Au collisions at $\sqrt{s(NN)}^{1/2} = 130\text{-GeV}$ ”, *Phys. Rev. Lett.* **89** (2002) 212301, doi:10.1103/PhysRevLett.89.212301, arXiv:nucl-ex/0204005.
- [48] CMS Collaboration, “Measurement of the elliptic anisotropy of charged particles produced in PbPb collisions at $\sqrt{s_{NN}}=2.76\text{ TeV}$ ”, *Phys. Rev.* **C87** (2013), no. 1, 014902, doi:10.1103/PhysRevC.87.014902, arXiv:1204.1409.
- [49] ALICE Collaboration, “Elliptic flow of charged particles in Pb-Pb collisions at 2.76 TeV”, *Phys. Rev. Lett.* **105** (2010) 252302, doi:10.1103/PhysRevLett.105.252302, arXiv:1011.3914.
- [50] ATLAS Collaboration, “Measurement of the azimuthal anisotropy for charged particle production in $\sqrt{s_{NN}} = 2.76\text{ TeV}$ lead-lead collisions with the ATLAS detector”, *Phys. Rev.* **C86** (2012) 014907, doi:10.1103/PhysRevC.86.014907, arXiv:1203.3087.
- [51] ATLAS Collaboration, “Measurement of the azimuthal anisotropy of charged particles produced in $\sqrt{s_{NN}} = 5.02\text{ TeV}$ Pb+Pb collisions with the ATLAS detector”, *Eur. Phys. J.* **C78** (2018), no. 12, 997, doi:10.1140/epjc/s10052-018-6468-7, arXiv:1808.03951.
- [52] J.-Y. Ollitrault, “Anisotropy as a signature of transverse collective flow”, *Phys. Rev.* **D46** (1992) 229–245, doi:10.1103/PhysRevD.46.229.
- [53] CMS Collaboration, “Observation of long-range near-side angular correlations in proton-lead collisions at the LHC”, *Phys. Lett.* **B718** (2013) 795–814, doi:10.1016/j.physletb.2012.11.025, arXiv:1210.5482.
- [54] CMS Collaboration, “Observation of Long-Range Near-Side Angular Correlations in Proton-Proton Collisions at the LHC”, *JHEP* **09** (2010) 091, doi:10.1007/JHEP09(2010)091, arXiv:1009.4122.
- [55] ATLAS Collaboration, “Observation of Associated Near-Side and Away-Side Long-Range Correlations in $\sqrt{s_{NN}} = 5.02\text{ TeV}$ Proton-Lead Collisions with the ATLAS Detector”, *Phys. Rev. Lett.* **110** (2013), no. 18, 182302, doi:10.1103/PhysRevLett.110.182302, arXiv:1212.5198.
- [56] ALICE Collaboration, “Long-range angular correlations on the near and away side in $p\text{-Pb}$ collisions at $\sqrt{s_{NN}} = 5.02\text{ TeV}$ ”, *Phys. Lett.* **B719** (2013) 29–41, doi:10.1016/j.physletb.2013.01.012, arXiv:1212.2001.
- [57] PHENIX Collaboration, “Creating small circular, elliptical, and triangular droplets of quark-gluon plasma”, *Nature Phys.* **15** (2019), no. 3, 214–220, doi:10.1038/s41567-018-0360-0, arXiv:1805.02973.
- [58] F. Gelis, “Color Glass Condensate and Glasma”, *Int. J. Mod. Phys.* **A28** (2013) 1330001, doi:10.1142/S0217751X13300019, arXiv:1211.3327.

- [59] PHENIX Collaboration, “Enhanced production of direct photons in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV and implications for the initial temperature”, *Phys. Rev. Lett.* **104** (2010) 132301, doi:10.1103/PhysRevLett.104.132301, arXiv:0804.4168.
- [60] ALICE Collaboration, “Direct photon production in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Phys. Lett.* **B754** (2016) 235–248, doi:10.1016/j.physletb.2016.01.020, arXiv:1509.07324.
- [61] C. Shen, U. W. Heinz, J.-F. Paquet, and C. Gale, “Thermal photons as a quark-gluon plasma thermometer reexamined”, *Phys. Rev.* **C89** (2014), no. 4, 044910, doi:10.1103/PhysRevC.89.044910, arXiv:1308.2440.
- [62] P. Koch, B. Müller, and J. Rafelski, “From strangeness enhancement to quark-gluon plasma discovery”, *Int. J. Mod. Phys.* **A32** (2017), no. 31, 1730024, doi:10.1142/S0217751X17300241, arXiv:1708.08115.
- [63] P. Seyboth, “Strange particle production in ultrarelativistic heavy ion collisions at the CERN SPS”, *J. Phys.* **G23** (1997) 1787–1801, doi:10.1088/0954-3899/23/12/002.
- [64] NA49 Collaboration, “Production of Phi mesons in p + p, p + Pb and central Pb + Pb collisions at E(beam) = 158-A-GeV”, *Phys. Lett.* **B491** (2000) 59–66, doi:10.1016/S0370-2693(00)01023-6.
- [65] STAR Collaboration, “Enhanced strange baryon production in Au + Au collisions compared to p + p at $\sqrt{s_{NN}} = 200$ GeV”, *Phys. Rev.* **C77** (2008) 044908, doi:10.1103/PhysRevC.77.044908, arXiv:0705.2511.
- [66] ALICE Collaboration, “Multi-strange baryon production at mid-rapidity in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Phys. Lett.* **B728** (2014) 216–227, doi:10.1016/j.physletb.2014.05.052, 10.1016/j.physletb.2013.11.048, arXiv:1307.5543. [Erratum: *Phys. Lett.*B734,409(2014)].
- [67] ALICE Collaboration, “Enhanced production of multi-strange hadrons in high-multiplicity proton-proton collisions”, *Nature Phys.* **13** (2017) 535–539, doi:10.1038/nphys4111, arXiv:1606.07424.
- [68] ALICE Collaboration, “Multi-strange baryon production in p-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Phys. Lett.* **B758** (2016) 389–401, doi:10.1016/j.physletb.2016.05.027, arXiv:1512.07227.
- [69] A. Andronic et al., “Heavy-flavour and quarkonium production in the LHC era: from proton-proton to heavy-ion collisions”, *Eur. Phys. J.* **C76** (2016), no. 3, 107, doi:10.1140/epjc/s10052-015-3819-5, arXiv:1506.03981.
- [70] D. Kharzeev, C. Lourenco, M. Nardi, and H. Satz, “A Quantitative analysis of charmonium suppression in nuclear collisions”, *Z. Phys.* **C74** (1997) 307–318, doi:10.1007/s002880050392, arXiv:hep-ph/9612217.

- [71] NA50 Collaboration, “Anomalous J/ψ suppression in Pb - Pb interactions at 158 GeV/c per nucleon”, *Phys. Lett.* **B410** (1997) 337–343, doi:10.1016/S0370-2693(97)00915-5.
- [72] STAR Collaboration, “ J/ψ production at high transverse momenta in $p + p$ and Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV”, *Phys. Lett.* **B722** (2013) 55–62, doi:10.1016/j.physletb.2013.04.010, arXiv:1208.2736.
- [73] PHENIX Collaboration, “ J/ψ suppression at forward rapidity in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV”, *Phys. Rev.* **C84** (2011) 054912, doi:10.1103/PhysRevC.84.054912, arXiv:1103.6269.
- [74] CMS Collaboration, “Observation of sequential Upsilon suppression in PbPb collisions”, *Phys. Rev. Lett.* **109** (2012) 222301, doi:10.1103/PhysRevLett.109.222301, 10.1103/PhysRevLett.120.199903, arXiv:1208.2826. [Erratum: Phys. Rev. Lett.120,no.19,199903(2018)].
- [75] CMS Collaboration, “Measurement of Prompt $\psi(2S) \rightarrow J/\psi$ Yield Ratios in Pb-Pb and $p - p$ Collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Phys. Rev. Lett.* **113** (2014), no. 26, 262301, doi:10.1103/PhysRevLett.113.262301, arXiv:1410.1804.
- [76] ALICE Collaboration, “Centrality, rapidity and transverse momentum dependence of J/ψ suppression in Pb-Pb collisions at $\sqrt{s_{NN}}=2.76$ TeV”, *Phys. Lett.* **B734** (2014) 314–327, doi:10.1016/j.physletb.2014.05.064, arXiv:1311.0214.
- [77] CMS Collaboration, “Suppression of Excited Υ States Relative to the Ground State in Pb-Pb Collisions at $\sqrt{s_{NN}}=5.02$ TeV”, *Phys. Rev. Lett.* **120** (2018), no. 14, 142301, doi:10.1103/PhysRevLett.120.142301, arXiv:1706.05984.
- [78] M. Arneodo, “Nuclear effects in structure functions”, *Physics Reports* **240** (1994), no. 5, 301 – 393, doi:https://doi.org/10.1016/0370-1573(94)90048-5.
- [79] K. J. Eskola, H. Paukkunen, and C. A. Salgado, “EPS09: A New Generation of NLO and LO Nuclear Parton Distribution Functions”, *JHEP* **04** (2009) 065, doi:10.1088/1126-6708/2009/04/065, arXiv:0902.4154.
- [80] K. Kovarik et al., “nCTEQ15 - Global analysis of nuclear parton distributions with uncertainties in the CTEQ framework”, *Phys. Rev.* **D93** (2016), no. 8, 085037, doi:10.1103/PhysRevD.93.085037, arXiv:1509.00792.
- [81] K. J. Eskola, P. Paakkinen, H. Paukkunen, and C. A. Salgado, “EPPS16: Nuclear parton distributions with LHC data”, *Eur. Phys. J.* **C77** (2017), no. 3, 163, doi:10.1140/epjc/s10052-017-4725-9, arXiv:1612.05741.

- [82] D. de Florian, R. Sassot, P. Zurita, and M. Stratmann, “Global Analysis of Nuclear Parton Distributions”, *Phys. Rev.* **D85** (2012) 074028, doi:10.1103/PhysRevD.85.074028, arXiv:1112.6324.
- [83] J. W. Cronin et al., “Production of hadrons at large transverse momentum at 200, 300, and 400 gev”, *Phys. Rev. D* **11** (Jun, 1975) 3105–3123, doi:10.1103/PhysRevD.11.3105.
- [84] PHENIX Collaboration, “Spectra and ratios of identified particles in Au+Au and d+Au collisions at $\sqrt{s_{NN}} = 200$ GeV”, *Phys. Rev.* **C88** (2013), no. 2, 024906, doi:10.1103/PhysRevC.88.024906, arXiv:1304.3410.
- [85] B. Z. Kopeliovich, J. Nemchik, A. Schafer, and A. V. Tarasov, “Cronin effect in hadron production off nuclei”, *Phys. Rev. Lett.* **88** (2002) 232303, doi:10.1103/PhysRevLett.88.232303, arXiv:hep-ph/0201010.
- [86] CMS Collaboration, “Study of Z production in PbPb and pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV in the dimuon and dielectron decay channels”, *JHEP* **03** (2015) 022, doi:10.1007/JHEP03(2015)022, arXiv:1410.4825.
- [87] CMS Collaboration, “Study of W boson production in PbPb and pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Phys. Lett.* **B715** (2012) 66–87, doi:10.1016/j.physletb.2012.07.025, arXiv:1205.6334.
- [88] CMS Collaboration, “Measurement of isolated photon production in pp and PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Phys. Lett.* **B710** (2012) 256–277, doi:10.1016/j.physletb.2012.02.077, arXiv:1201.3093.
- [89] ATLAS Collaboration, “Measurement of Z boson Production in Pb+Pb Collisions at $\sqrt{s_{NN}} = 2.76$ TeV with the ATLAS Detector”, *Phys. Rev. Lett.* **110** (2013), no. 2, 022301, doi:10.1103/PhysRevLett.110.022301, arXiv:1210.6486.
- [90] ATLAS Collaboration, “Measurement of the production and lepton charge asymmetry of W bosons in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV with the ATLAS detector”, *Eur. Phys. J.* **C75** (2015), no. 1, 23, doi:10.1140/epjc/s10052-014-3231-6, arXiv:1408.4674.
- [91] ATLAS Collaboration, “Centrality, rapidity and transverse momentum dependence of isolated prompt photon production in lead-lead collisions at $\sqrt{s_{NN}} = 2.76$ TeV measured with the ATLAS detector”, *Phys. Rev.* **C93** (2016), no. 3, 034914, doi:10.1103/PhysRevC.93.034914, arXiv:1506.08552.
- [92] ATLAS Collaboration Collaboration, “Z boson production in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the ATLAS detector at the LHC”, Technical Report ATLAS-CONF-2017-010, CERN, Geneva, Feb, 2017.

- [93] G.-Y. Qin and X.-N. Wang, “Jet quenching in high-energy heavy-ion collisions”, *Int. J. Mod. Phys.* **E24** (2015), no. 11, 1530014, doi:10.1142/S0218301315300143, 10.1142/9789814663717_0007, arXiv:1511.00790. [,309(2016)].
- [94] J. Casalderrey-Solana et al., “A Hybrid Strong/Weak Coupling Approach to Jet Quenching”, *JHEP* **10** (2014) 019, doi:10.1007/JHEP09(2015)175, 10.1007/JHEP10(2014)019, arXiv:1405.3864. [Erratum: JHEP09,175(2015)].
- [95] J. M. Maldacena, “The Large N limit of superconformal field theories and supergravity”, *Int. J. Theor. Phys.* **38** (1999) 1113–1133, doi:10.1023/A:1026654312961, 10.4310/ATMP.1998.v2.n2.a1, arXiv:hep-th/9711200. [Adv. Theor. Math. Phys.2,231(1998)].
- [96] R. Baier et al., “Radiative energy loss of high-energy quarks and gluons in a finite volume quark - gluon plasma”, *Nucl. Phys.* **B483** (1997) 291–320, doi:10.1016/S0550-3213(96)00553-6, arXiv:hep-ph/9607355.
- [97] R. Baier, Y. L. Dokshitzer, A. H. Mueller, and D. Schiff, “Quenching of hadron spectra in media”, *JHEP* **09** (2001) 033, doi:10.1088/1126-6708/2001/09/033, arXiv:hep-ph/0106347.
- [98] F. Arleo and G. Falmagne, “Quenching of hadron spectra in XeXe and PbPb collisions at the LHC”, *PoS HardProbes2018* (2018) 075, arXiv:1902.05032.
- [99] F. Arleo, “Tomography of cold and hot QCD matter: Tools and diagnosis”, *JHEP* **11** (2002) 044, doi:10.1088/1126-6708/2002/11/044, arXiv:hep-ph/0210104.
- [100] M. Gyulassy, P. Levai, and I. Vitev, “Jet quenching in thin quark gluon plasmas. 1. Formalism”, *Nucl. Phys.* **B571** (2000) 197–233, doi:10.1016/S0550-3213(99)00713-0, arXiv:hep-ph/9907461.
- [101] M. Djordjevic and U. W. Heinz, “Radiative energy loss in a finite dynamical QCD medium”, *Phys. Rev. Lett.* **101** (2008) 022302, doi:10.1103/PhysRevLett.101.022302, arXiv:0802.1230.
- [102] M. Djordjevic, “Collisional energy loss in a finite size qcd matter”, *Phys. Rev. C* **74** (Dec, 2006) 064907, doi:10.1103/PhysRevC.74.064907.
- [103] M. Djordjevic, “Theoretical formalism of radiative jet energy loss in a finite size dynamical QCD medium”, *Phys. Rev.* **C80** (2009) 064909, doi:10.1103/PhysRevC.80.064909, arXiv:0903.4591.
- [104] M. Djordjevic and M. Djordjevic, “LHC jet suppression of light and heavy flavor observables”, *Phys. Lett.* **B734** (2014) 286–289, doi:10.1016/j.physletb.2014.05.053, arXiv:1307.4098.

- [105] J. Xu, J. Liao, and M. Gyulassy, “Bridging Soft-Hard Transport Properties of Quark-Gluon Plasmas with CUJET3.0”, *JHEP* **02** (2016) 169, doi:10.1007/JHEP02(2016)169, arXiv:1508.00552.
- [106] M. H. Thoma and M. Gyulassy, “Quark damping and energy loss in the high temperature qcd”, *Nuclear Physics B* **351** (1991), no. 3, 491 – 506, doi:https://doi.org/10.1016/S0550-3213(05)80031-8.
- [107] C. W. Bauer, D. Pirjol, and I. W. Stewart, “Soft collinear factorization in effective field theory”, *Phys. Rev.* **D65** (2002) 054022, doi:10.1103/PhysRevD.65.054022, arXiv:hep-ph/0109045.
- [108] A. Idilbi and A. Majumder, “Extending Soft-Collinear-Effective-Theory to describe hard jets in dense QCD media”, *Phys. Rev.* **D80** (2009) 054022, doi:10.1103/PhysRevD.80.054022, arXiv:0808.1087.
- [109] Z.-B. Kang et al., “Jet quenching phenomenology from soft-collinear effective theory with Glauber gluons”, *Phys. Rev. Lett.* **114** (2015), no. 9, 092002, doi:10.1103/PhysRevLett.114.092002, arXiv:1405.2612.
- [110] X.-N. Wang and X.-f. Guo, “Multiple parton scattering in nuclei: Parton energy loss”, *Nucl. Phys.* **A696** (2001) 788–832, doi:10.1016/S0375-9474(01)01130-7, arXiv:hep-ph/0102230.
- [111] R. J. Fries, “Higher twist effects in nuclei”, in *Ultrarelativistic heavy-ion collisions. Proceedings, 30th International Workshop on Gross Properties of Nuclei and Nuclear Excitations, Hirschegg, Austria, January 13-19, 2002*, pp. 348–357. 2002. arXiv:hep-ph/0201311.
- [112] Y. He, T. Luo, X.-N. Wang, and Y. Zhu, “Linear Boltzmann Transport for Jet Propagation in the Quark-Gluon Plasma: Elastic Processes and Medium Recoil”, *Phys. Rev.* **C91** (2015) 054908, doi:10.1103/PhysRevC.91.054908, 10.1103/PhysRevC.97.019902, arXiv:1503.03313. [Erratum: Phys. Rev.C97,no.1,019902(2018)].
- [113] E. Bianchi et al., “The x and Q^2 dependence of \hat{q} , quasi-particles and the JET puzzle”, arXiv:1702.00481.
- [114] A. Majumder and C. Shen, “Suppression of the High p_T Charged Hadron R_{AA} at the LHC”, *Phys. Rev. Lett.* **109** (2012) 202301, doi:10.1103/PhysRevLett.109.202301, arXiv:1103.0809.
- [115] B. Betz, M. Gyulassy, and G. Torrieri, “Fourier Harmonics of High- p_T Particles Probing the Fluctuating Initial Condition Geometries in Heavy-Ion Collisions”, *Phys. Rev.* **C84** (2011) 024913, doi:10.1103/PhysRevC.84.024913, arXiv:1102.5416.

- [116] B. Betz and M. Gyulassy, “Constraints on the Path-Length Dependence of Jet Quenching in Nuclear Collisions at RHIC and LHC”, *JHEP* **08** (2014) 090, doi:10.1007/JHEP10(2014)043, 10.1007/JHEP08(2014)090, arXiv:1404.6378. [Erratum: JHEP10,043(2014)].
- [117] C. Andrés et al., “Energy versus centrality dependence of the jet quenching parameter \hat{q} at RHIC and LHC: a new puzzle?”, *Eur. Phys. J.* **C76** (2016), no. 9, 475, doi:10.1140/epjc/s10052-016-4320-5, arXiv:1606.04837.
- [118] BRAHMS Collaboration, “Transverse-momentum spectra in Au + Au and $d + Au$ collisions at $\sqrt{s_{NN}} = 200$ GeV and the pseudorapidity dependence of high- p_T suppression”, *Phys. Rev. Lett.* **91** (Aug, 2003) 072305, doi:10.1103/PhysRevLett.91.072305.
- [119] PHENIX Collaboration, “Absence of suppression in particle production at large transverse momentum in $S(NN)^{1/2} = 200$ GeV $d + Au$ collisions”, *Phys. Rev. Lett.* **91** (2003) 072303, doi:10.1103/PhysRevLett.91.072303, arXiv:nucl-ex/0306021.
- [120] PHENIX Collaboration, “High transverse momentum η meson production in p^+p , $d^+ Au$ and $Au+Au$ collisions at $S(NN)^{(1/2)} = 200$ -GeV”, *Phys. Rev.* **C75** (2007) 024909, doi:10.1103/PhysRevC.75.024909, arXiv:nucl-ex/0611006.
- [121] PHOBOS Collaboration Collaboration, “Centrality dependence of charged-hadron transverse-momentum spectra in $d + Au$ collisions at $\sqrt{s_{NN}} = 200$ GeV”, *Phys. Rev. Lett.* **91** (Aug, 2003) 072302, doi:10.1103/PhysRevLett.91.072302.
- [122] STAR Collaboration, “Evidence from $d + Au$ measurements for final-state suppression of high- p_T hadrons in $Au + Au$ collisions at rhic”, *Phys. Rev. Lett.* **91** (Aug, 2003) 072304, doi:10.1103/PhysRevLett.91.072304.
- [123] ALICE Collaboration, “Transverse momentum dependence of inclusive primary charged-particle production in p-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Eur. Phys. J.* **C74** (2014), no. 9, 3054, doi:10.1140/epjc/s10052-014-3054-5, arXiv:1405.2737.
- [124] ALICE Collaboration, “Neutral pion and η meson production in p-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Eur. Phys. J.* **C78** (2018), no. 8, 624, doi:10.1140/epjc/s10052-018-6013-8, arXiv:1801.07051.
- [125] ATLAS Collaboration, “Transverse momentum, rapidity, and centrality dependence of inclusive charged-particle production in $\sqrt{s_{NN}} = 5.02$ TeV $p + Pb$ collisions measured by the ATLAS experiment”, *Phys. Lett.* **B763** (2016) 313–336, doi:10.1016/j.physletb.2016.10.053, arXiv:1605.06436.

- [126] CMS Collaboration, “Nuclear Effects on the Transverse Momentum Spectra of Charged Particles in pPb Collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Eur. Phys. J.* **C75** (2015), no. 5, 237, doi:10.1140/epjc/s10052-015-3435-4, arXiv:1502.05387.
- [127] A. Angelis et al., “Large transverse momentum π^0 production in $\alpha\alpha$, dd and pp collisions at the cern isr”, *Physics Letters B* **185** (1987), no. 1, 213 – 217, doi:https://doi.org/10.1016/0370-2693(87)91557-7.
- [128] NA49 Collaboration, “High Transverse Momentum Hadron Spectra at $s(NN)^{(1/2)} = 17.3$ -GeV, in Pb+Pb and p+p Collisions, Measured by CERN-NA49”, *Phys. Rev.* **C77** (2008) 034906, doi:10.1103/PhysRevC.77.034906, arXiv:0711.0547.
- [129] WA98 Collaboration, “Transverse mass distributions of neutral pions from Pb-208 induced reactions at 158-A-GeV”, *Eur. Phys. J.* **C23** (2002) 225–236, doi:10.1007/s100520100886, arXiv:nucl-ex/0108006.
- [130] J. Silnova, “Azimuthal Correlations of High-pt Pions in 158 AGeV/c Pb-Au Collisions measured by the CERES/NA45 Experiment”. PhD thesis, Charles University, 2003.
- [131] WA80 Collaboration, “Transverse momentum distributions of neutral pions from nuclear collisions at 200-A/GeV”, *Eur. Phys. J.* **C5** (1998) 255–267, doi:10.1007/s100520050267, arXiv:nucl-ex/9805007.
- [132] PHENIX Collaboration, “Neutral pion production with respect to centrality and reaction plane in Au+Au collisions at $\sqrt{s_{NN}}=200$ GeV”, *Phys. Rev.* **C87** (2013), no. 3, 034911, doi:10.1103/PhysRevC.87.034911, arXiv:1208.2254.
- [133] PHENIX Collaboration, “Suppressed π^0 production at large transverse momentum in central Au+ Au collisions at $S(NN)^{1/2} = 200$ GeV”, *Phys. Rev. Lett.* **91** (2003) 072301, doi:10.1103/PhysRevLett.91.072301, arXiv:nucl-ex/0304022.
- [134] PHENIX Collaboration, “High p_T charged hadron suppression in Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV”, *Phys. Rev.* **C69** (2004) 034910, doi:10.1103/PhysRevC.69.034910, arXiv:nucl-ex/0308006.
- [135] PHENIX Collaboration, “Evolution of π^0 suppression in Au+Au collisions from $\sqrt{s_{NN}} = 39$ to 200 GeV”, *Phys. Rev. Lett.* **109** (2012) 152301, doi:10.1103/PhysRevLett.109.152301, arXiv:1204.1526.
- [136] PHOBOS Collaboration, “Centrality dependence of charged hadron transverse momentum spectra in Au + Au collisions from $s(NN)^{(1/2)} = 62.4$ -GeV to 200-GeV”, *Phys. Rev. Lett.* **94** (2005) 082304, doi:10.1103/PhysRevLett.94.082304, arXiv:nucl-ex/0405003.

- [137] STAR Collaboration, “Transverse momentum and collision energy dependence of high p_T hadron suppression in Au+Au collisions at ultrarelativistic energies”, *Phys. Rev. Lett.* **91** (2003) 172302, doi:10.1103/PhysRevLett.91.172302, arXiv:nucl-ex/0305015.
- [138] BRAHMS Collaboration, “Rapidity and centrality dependence of particle production for identified hadrons in Cu+Cu collisions at $\sqrt{s_{NN}} = 200$ GeV”, *Phys. Rev.* **C94** (2016), no. 1, 014907, doi:10.1103/PhysRevC.94.014907, arXiv:1602.01183.
- [139] PHENIX Collaboration, “Onset of π^0 Suppression Studied in Cu+Cu Collisions at $s_{NN}=22.4, 62.4,$ and 200 GeV”, *Phys. Rev. Lett.* **101** (2008) 162301, doi:10.1103/PhysRevLett.101.162301, arXiv:0801.4555.
- [140] PHOBOS Collaboration, “System size and centrality dependence of charged hadron transverse momentum spectra in Au + Au and Cu + Cu collisions at $s_{NN}^{1/2} = 62.4$ -GeV and 200 -GeV”, *Phys. Rev. Lett.* **96** (2006) 212301, doi:10.1103/PhysRevLett.96.212301, arXiv:nucl-ex/0512016.
- [141] STAR Collaboration, “Spectra of identified high- p_T π^\pm and $p(\bar{p})$ in Cu+Cu collisions at $\sqrt{s_{NN}} = 200$ GeV”, *Phys. Rev.* **C81** (2010) 054907, doi:10.1103/PhysRevC.81.054907, arXiv:0911.3130.
- [142] ALICE Collaboration, “Centrality Dependence of Charged Particle Production at Large Transverse Momentum in Pb–Pb Collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Phys. Lett.* **B720** (2013) 52–62, doi:10.1016/j.physletb.2013.01.051, arXiv:1208.2711.
- [143] ALICE Collaboration, “Neutral pion production at midrapidity in pp and Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Eur. Phys. J.* **C74** (2014), no. 10, 3108, doi:10.1140/epjc/s10052-014-3108-8, arXiv:1405.3794.
- [144] ALICE Collaboration, “Neutral pion and η meson production at mid-rapidity in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Phys. Rev.* **C98** (2018), no. 4, 044901, doi:10.1103/PhysRevC.98.044901, arXiv:1803.05490.
- [145] ATLAS Collaboration, “Measurement of charged-particle spectra in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV with the ATLAS detector at the LHC”, *JHEP* **09** (2015) 050, doi:10.1007/JHEP09(2015)050, arXiv:1504.04337.
- [146] CMS Collaboration, “Study of high- p_T charged particle suppression in PbPb compared to pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Eur. Phys. J.* **C72** (2012) 1945, doi:10.1140/epjc/s10052-012-1945-x, arXiv:1202.2554.
- [147] European Muon Collaboration, “The ratio of the nucleon structure functions F_2^n for iron and deuterium”, *Phys. Lett.* **123B** (1983) 275–278, doi:10.1016/0370-2693(83)90437-9.

- [148] CLAS Collaboration, “Modified structure of protons and neutrons in correlated pairs”, *Nature* **566** (2019), no. 7744, 354–358, doi:10.1038/s41586-019-0925-9.
- [149] BRAHMS Collaboration, “Quark gluon plasma and color glass condensate at RHIC? The Perspective from the BRAHMS experiment”, *Nucl. Phys.* **A757** (2005) 1–27, doi:10.1016/j.nuclphysa.2005.02.130, arXiv:nucl-ex/0410020.
- [150] CMS Collaboration, “Constraining gluon distributions in nuclei using dijets in proton-proton and proton-lead collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Phys. Rev. Lett.* **121** (2018), no. 6, 062002, doi:10.1103/PhysRevLett.121.062002, arXiv:1805.04736.
- [151] H. Paukkunen, “The LHC p+Pb run from the nuclear PDF perspective”, *PoS DIS2014* (2014) 053, doi:10.22323/1.203.0053, arXiv:1408.4657.
- [152] D. G. d’Enterria, “Indications of suppressed high p_T hadron production in nucleus-nucleus collisions at CERN-SPS”, *Phys. Lett. B* **596** (2004), no. nucl-ex/0403055, 32–43. 18 p.
- [153] PHENIX Collaboration, “Formation of dense partonic matter in relativistic nucleus-nucleus collisions at RHIC: Experimental evaluation by the PHENIX collaboration”, *Nucl. Phys.* **A757** (2005) 184–283, doi:10.1016/j.nuclphysa.2005.03.086, arXiv:nucl-ex/0410003.
- [154] L. Evans and P. Bryant, “LHC Machine”, *JINST* **3** (2008) S08001, doi:10.1088/1748-0221/3/08/S08001.
- [155] T. C. Collaboration, “The cms experiment at the cern lhc”, *Journal of Instrumentation* **3** (2008), no. 08, S08004.
- [156] CMS Collaboration, “CMS Tracking Performance Results from early LHC Operation”, *Eur. Phys. J.* **C70** (2010) 1165–1192, doi:10.1140/epjc/s10052-010-1491-3, arXiv:1007.1988.
- [157] CMS Collaboration, “Description and performance of track and primary-vertex reconstruction with the CMS tracker”, *JINST* **9** (2014), no. 10, P10009, doi:10.1088/1748-0221/9/10/P10009, arXiv:1405.6569.
- [158] CMS Collaboration, “CMS Technical Design Report for the Pixel Detector Upgrade”, doi:10.2172/1151650.
- [159] R. L. Gluckstern, “Uncertainties in track momentum and direction, due to multiple scattering and measurement errors”, *Nucl. Instrum. Meth.* **24** (1963) 381–389, doi:10.1016/0029-554X(63)90347-1.

- [160] CMS Collaboration, “Commissioning of the CMS High-Level Trigger with Cosmic Rays”, *JINST* **5** (2010) T03005, doi:10.1088/1748-0221/5/03/T03005, arXiv:0911.4889.
- [161] CMS Collaboration, “The CMS trigger system”, *JINST* **12** (2017), no. 01, P01020, doi:10.1088/1748-0221/12/01/P01020, arXiv:1609.02366.
- [162] C. Collaboration, “CMS Heavy Ion Event Display - 25 November 2015”,. General Photo.
- [163] PHOBOS Collaboration, “Importance of correlations and fluctuations on the initial source eccentricity in high-energy nucleus-nucleus collisions”, *Phys. Rev. C* **77** (2008) 014906, doi:10.1103/PhysRevC.77.014906, arXiv:0711.3724.
- [164] P. Billoir and S. Qian, “Fast vertex fitting with a local parametrization of tracks”, *Nucl. Instrum. Meth.* **A311** (1992) 139–150, doi:10.1016/0168-9002(92)90859-3.
- [165] M. Swartz et al., “A new technique for the reconstruction, validation, and simulation of hits in the CMS pixel detector”, *PoS VERTEX2007* (2007) 035, doi:10.22323/1.057.0035.
- [166] M. Swartz, “CMS pixel simulations”, *Nucl. Instrum. Meth.* **A511** (2003) 88–91, doi:10.1016/S0168-9002(03)01757-1.
- [167] R. Frühwirth et al., “Track finding in silicon trackers with a small number of layers”, *Nucl. Instrum. Meth.* **A732** (2013) 95–98, doi:10.1016/j.nima.2013.06.035.
- [168] R. Fruhwirth, “Application of Kalman filtering to track and vertex fitting”, *Nucl. Instrum. Meth.* **A262** (1987) 444–450, doi:10.1016/0168-9002(87)90887-4.
- [169] P. Billoir, “Progressive track recognition with a Kalman like fitting procedure”, *Comput. Phys. Commun.* **57** (1989) 390–394, doi:10.1016/0010-4655(89)90249-X.
- [170] P. Billoir and S. Qian, “Simultaneous pattern recognition and track fitting by the Kalman filtering method”, *Nucl. Instrum. Meth.* **A294** (1990) 219–228, doi:10.1016/0168-9002(90)91835-Y.
- [171] H. Drucker and C. Cortes, “Boosting decision trees”, in *Proceedings of the 8th International Conference on Neural Information Processing Systems, NIPS’95*, pp. 479–485. MIT Press, Cambridge, MA, USA, 1995.
- [172] A. Hocker et al., “TMVA - Toolkit for Multivariate Data Analysis”, arXiv:physics/0703039.

- [173] CMS Collaboration, “Charged-particle nuclear modification factors in XeXe collisions at $\sqrt{s_{NN}} = 5.44$ TeV”, *JHEP* **10** (2018) 138, doi:10.1007/JHEP10(2018)138, arXiv:1809.00201.
- [174] K. Rose, “Deterministic annealing for clustering, compression, classification, regression, and related optimization problems”, *IEEE Proc.* **86** (1998), no. 11, 2210–2239, doi:10.1109/5.726788.
- [175] R. Frühwirth, W. Waltenberger, and P. Vanlaer, “Adaptive Vertex Fitting”, Technical Report CMS-NOTE-2007-008, CERN, Geneva, Mar, 2007.
- [176] CMS Collaboration, “Particle-flow reconstruction and global event description with the CMS detector”, *JINST* **12** (2017), no. 10, P10003, doi:10.1088/1748-0221/12/10/P10003, arXiv:1706.04965.
- [177] M. Cacciari, G. P. Salam, and G. Soyez, “The anti- k_t jet clustering algorithm”, *JHEP* **04** (2008) 063, doi:10.1088/1126-6708/2008/04/063, arXiv:0802.1189.
- [178] M. Cacciari, G. P. Salam, and G. Soyez, “FastJet User Manual”, *Eur. Phys. J.* **C72** (2012) 1896, doi:10.1140/epjc/s10052-012-1896-2, arXiv:1111.6097.
- [179] M. Cacciari and G. P. Salam, “Pileup subtraction using jet areas”, *Phys. Lett.* **B659** (2008) 119–126, doi:10.1016/j.physletb.2007.09.077, arXiv:0707.1378.
- [180] CMS Collaboration, “Jet momentum dependence of jet quenching in PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Phys. Lett.* **B712** (2012) 176–197, doi:10.1016/j.physletb.2012.04.058, arXiv:1202.5022.
- [181] CMS Collaboration, “Studies of jet quenching using isolated-photon+jet correlations in PbPb and pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Phys. Lett.* **B718** (2013) 773–794, doi:10.1016/j.physletb.2012.11.003, arXiv:1205.0206.
- [182] CMS Collaboration Collaboration, “Jet Energy Corrections determination at 7 TeV”, Technical Report CMS-PAS-JME-10-010, CERN, Geneva, 2010.
- [183] M. L. Miller, K. Reygers, S. J. Sanders, and P. Steinberg, “Glauber modeling in high energy nuclear collisions”, *Ann. Rev. Nucl. Part. Sci.* **57** (2007) 205–243, doi:10.1146/annurev.nucl.57.090506.123020, arXiv:nucl-ex/0701025.
- [184] J. C. Collins, “Sudakov form-factors”, *Adv. Ser. Direct. High Energy Phys.* **5** (1989) 573–614, doi:10.1142/9789814503266_0006, arXiv:hep-ph/0312336.

- [185] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjostrand, “Parton Fragmentation and String Dynamics”, *Phys. Rept.* **97** (1983) 31–145, doi:10.1016/0370-1573(83)90080-7.
- [186] T. Sjostrand, S. Mrenna, and P. Z. Skands, “PYTHIA 6.4 Physics and Manual”, *JHEP* **05** (2006) 026, doi:10.1088/1126-6708/2006/05/026, arXiv:hep-ph/0603175.
- [187] R. Field, “Early LHC Underlying Event Data - Findings and Surprises”, in *Hadron collider physics. Proceedings, 22nd Conference, HCP 2010, Toronto, Canada, August 23-27, 2010*. 2010. arXiv:1010.3558.
- [188] T. Sjöstrand et al., “An Introduction to PYTHIA 8.2”, *Comput. Phys. Commun.* **191** (2015) 159–177, doi:10.1016/j.cpc.2015.01.024, arXiv:1410.3012.
- [189] CMS Collaboration, “Event generator tunes obtained from underlying event and multiparton scattering measurements”, *Eur. Phys. J.* **C76** (2016), no. 3, 155, doi:10.1140/epjc/s10052-016-3988-x, arXiv:1512.00815.
- [190] M. Bahr et al., “Herwig++ Physics and Manual”, *Eur. Phys. J.* **C58** (2008) 639–707, doi:10.1140/epjc/s10052-008-0798-9, arXiv:0803.0883.
- [191] 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** (1994) 307, doi:10.1016/0010-4655(94)90057-4, arXiv:nucl-th/9502021.
- [192] I. P. Lokhtin and A. M. Snigirev, “A Model of jet quenching in ultrarelativistic heavy ion collisions and high-p(T) hadron spectra at RHIC”, *Eur. Phys. J.* **C45** (2006) 211–217, doi:10.1140/epjc/s2005-02426-3, arXiv:hep-ph/0506189.
- [193] K. Werner, F.-M. Liu, and T. Pierog, “Parton ladder splitting and the rapidity dependence of transverse momentum spectra in deuteron-gold collisions at RHIC”, *Phys. Rev.* **C74** (2006) 044902, doi:10.1103/PhysRevC.74.044902, arXiv:hep-ph/0506232.
- [194] T. Pierog et al., “EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider”, *Phys. Rev.* **C92** (2015), no. 3, 034906, doi:10.1103/PhysRevC.92.034906, arXiv:1306.0121.
- [195] Z.-W. Lin et al., “A Multi-phase transport model for relativistic heavy ion collisions”, *Phys. Rev.* **C72** (2005) 064901, doi:10.1103/PhysRevC.72.064901, arXiv:nucl-th/0411110.
- [196] ALICE Collaboration, “Measurement of the Cross Section for Electromagnetic Dissociation with Neutron Emission in Pb-Pb Collisions at $\sqrt{s_{NN}} = 2.76$

- TeV”, *Phys. Rev. Lett.* **109** (2012) 252302, doi:10.1103/PhysRevLett.109.252302, arXiv:1203.2436.
- [197] S. R. Klein et al., “STARlight: A Monte Carlo simulation program for ultra-peripheral collisions of relativistic ions”, *Comput. Phys. Commun.* **212** (2017) 258–268, doi:10.1016/j.cpc.2016.10.016, arXiv:1607.03838.
- [198] GEANT4 Collaboration, “GEANT4: A Simulation toolkit”, *Nucl. Instrum. Meth.* **A506** (2003) 250–303, doi:10.1016/S0168-9002(03)01368-8.
- [199] B. Alver, M. Baker, C. Loizides, and P. Steinberg, “The PHOBOS Glauber Monte Carlo”, arXiv:0805.4411.
- [200] C. Loizides, J. Nagle, and P. Steinberg, “Improved version of the PHOBOS Glauber Monte Carlo”, doi:10.1016/j.softx.2015.05.001, arXiv:1408.2549. [SoftwareX1-2,13(2015)].
- [201] C. Loizides, J. Kamin, and D. d’Enterria, “Improved Monte Carlo Glauber predictions at present and future nuclear colliders”, *Phys. Rev.* **C97** (2018), no. 5, 054910, doi:10.1103/PhysRevC.97.054910, arXiv:1710.07098.
- [202] J. M. Allmond et al., “Double-Magic Nature of Sn132 and Pb208 through Lifetime and Cross-Section Measurements”, *Phys. Rev. Lett.* **112** (2014), no. 17, 172701, doi:10.1103/PhysRevLett.112.172701.
- [203] H. De Vries, C. W. De Jager, and C. De Vries, “Nuclear charge and magnetization density distribution parameters from elastic electron scattering”, *Atom. Data Nucl. Data Tabl.* **36** (1987) 495–536, doi:10.1016/0092-640X(87)90013-1.
- [204] ALICE Collaboration, “Centrality dependence of the pseudorapidity density distribution for charged particles in Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Phys. Lett.* **B772** (2017) 567–577, doi:10.1016/j.physletb.2017.07.017, arXiv:1612.08966.
- [205] NA50 Collaboration, “Scaling of charged particle multiplicity in Pb Pb collisions at SPS energies”, *Phys. Lett.* **B530** (2002) 43–55, doi:10.1016/S0370-2693(02)01353-9.
- [206] CMS Collaboration Collaboration, “CMS Luminosity Calibration for the pp Reference Run at $\sqrt{s} = 5.02$ TeV”, Technical Report CMS-PAS-LUM-16-001, CERN, Geneva, 2016.
- [207] S. van der Meer, “Calibration of the Effective Beam Height in the ISR”,.
- [208] CMS Collaboration, “Charged-particle nuclear modification factors in PbPb and pPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *JHEP* **04** (2017) 039, doi:10.1007/JHEP04(2017)039, arXiv:1611.01664.

- [209] CMS Collaboration Collaboration, “Jet Fragmentation Function in pPb Collisions at $\sqrt{s_{NN}} = 5.02$ TeV and pp Collisions at $\sqrt{s} = 2.76$ and 7 TeV”, Technical Report CMS-PAS-HIN-15-004, CERN, Geneva, 2015.
- [210] CMS Collaboration, “Measurement of transverse momentum relative to dijet systems in PbPb and pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *JHEP* **01** (2016) 006, doi:10.1007/JHEP01(2016)006, arXiv:1509.09029.
- [211] CMS Collaboration Collaboration, “Measurement of Tracking Efficiency”, Technical Report CMS-PAS-TRK-10-002, CERN, Geneva, 2010.
- [212] CMS Collaboration Collaboration, “Jet Energy Resolution in CMS at $\sqrt{s}=7$ TeV”, Technical Report CMS-PAS-JME-10-014, CERN, Geneva, 2011.
- [213] NNPDF Collaboration, “Unbiased global determination of parton distributions and their uncertainties at NNLO and at LO”, *Nucl. Phys.* **B855** (2012) 153–221, doi:10.1016/j.nuclphysb.2011.09.024, arXiv:1107.2652.
- [214] “Jet Fragmentation in p+Pb Collisions”, Technical Report ATLAS-CONF-2015-022, CERN, Geneva, Jun, 2015.
- [215] ATLAS Collaboration, “Measurement of jet fragmentation in 5.02 TeV proton-lead and proton-proton collisions with the ATLAS detector”, *Nucl. Phys.* **A978** (2018) 65, doi:10.1016/j.nuclphysa.2018.07.006, arXiv:1706.02859.
- [216] ALICE Collaboration, “Multi-strange baryon production in *pp* collisions at $\sqrt{s} = 7$ TeV with ALICE”, *Phys. Lett.* **B712** (2012) 309–318, doi:10.1016/j.physletb.2012.05.011, arXiv:1204.0282.
- [217] C.-Y. Wong and G. Wilk, “Tsallis Fits to p_t Spectra for pp Collisions at LHC”, *Acta Phys. Polon.* **B43** (2012) 2047–2054, doi:10.5506/APhysPolB.43.2047, arXiv:1210.3661.
- [218] D. d’Enterria, “Jet quenching”, *Landolt-Bornstein* **23** (2010) 471, doi:10.1007/978-3-642-01539-7_16, arXiv:0902.2011.
- [219] F. Arleo, “Quenching of Hadron Spectra in Heavy Ion Collisions at the LHC”, *Phys. Rev. Lett.* **119** (2017), no. 6, 062302, doi:10.1103/PhysRevLett.119.062302, arXiv:1703.10852.
- [220] Y.-T. Chien et al., “Jet Quenching from QCD Evolution”, *Phys. Rev.* **D93** (2016), no. 7, 074030, doi:10.1103/PhysRevD.93.074030, arXiv:1509.02936.
- [221] J. Noronha-Hostler, B. Betz, J. Noronha, and M. Gyulassy, “Event-by-event hydrodynamics + jet energy loss: A solution to the $R_{AA} \otimes v_2$ puzzle”, *Phys. Rev. Lett.* **116** (2016), no. 25, 252301, doi:10.1103/PhysRevLett.116.252301, arXiv:1602.03788.

- [222] D. G. d’Enterria, “Indications of suppressed high p(T) hadron production in nucleus - nucleus collisions at CERN-SPS”, *Phys. Lett.* **B596** (2004) 32–43, doi:10.1016/j.physletb.2004.06.071, arXiv:nucl-ex/0403055.
- [223] CMS Collaboration, “Nuclear modification factor of D^0 mesons in PbPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Phys. Lett.* **B782** (2018) 474–496, doi:10.1016/j.physletb.2018.05.074, arXiv:1708.04962.
- [224] CMS Collaboration, “Measurement of the B^\pm Meson Nuclear Modification Factor in Pb-Pb Collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Phys. Rev. Lett.* **119** (2017), no. 15, 152301, doi:10.1103/PhysRevLett.119.152301, arXiv:1705.04727.
- [225] CMS Collaboration, “Studies of beauty suppression via nonprompt D^0 mesons in PbPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV”, *Submitted to: Phys. Rev. Lett.* (2018) arXiv:1810.11102.
- [226] CMS Collaboration, “Measurement of prompt and nonprompt charmonium suppression in PbPb collisions at 5.02 TeV”, *Eur. Phys. J.* **C78** (2018), no. 6, 509, doi:10.1140/epjc/s10052-018-5950-6, arXiv:1712.08959.
- [227] CMS Collaboration, “Evidence for Collective Multiparticle Correlations in p-Pb Collisions”, *Phys. Rev. Lett.* **115** (2015), no. 1, 012301, doi:10.1103/PhysRevLett.115.012301, arXiv:1502.05382.
- [228] ATLAS Collaboration, “Measurement of nuclear modification factor R_{AA} in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the ATLAS detector at the LHC”, Technical Report ATLAS-CONF-2017-012, CERN, Geneva, Feb, 2017.
- [229] ALICE Collaboration, “Transverse momentum spectra and nuclear modification factors of charged particles in pp, p-Pb and Pb-Pb collisions at the LHC”, *JHEP* **11** (2018) 013, doi:10.1007/JHEP11(2018)013, arXiv:1802.09145.
- [230] ATLAS Collaboration, “Measurement of charged particle spectra in pp collisions and nuclear modification factor R_{pPb} at $\sqrt{s_{NN}} = 5.02$ TeV with the ATLAS detector at the LHC”,.
- [231] CMS Collaboration, “Charged particle transverse momentum spectra in pp collisions at $\sqrt{s} = 0.9$ and 7 TeV”, *JHEP* **08** (2011) 086, doi:10.1007/JHEP08(2011)086, arXiv:1104.3547.
- [232] ALICE Collaboration, “ K_S^0 and Λ production in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV”, *Phys. Rev. Lett.* **111** (2013) 222301, doi:10.1103/PhysRevLett.111.222301, arXiv:1307.5530.

- [233] ALICE Collaboration, “Anisotropic flow in Xe-Xe collisions at $\sqrt{s_{\text{NN}}} = 5.44$ TeV”, *Phys. Lett.* **B784** (2018) 82–95, doi:10.1016/j.physletb.2018.06.059, arXiv:1805.01832.
- [234] CMS Collaboration, “Charged-particle angular correlations in XeXe collisions at $\sqrt{s_{\text{NN}}} = 5.44$ TeV”, *Submitted to: Phys. Rev.* (2019) arXiv:1901.07997.
- [235] ATLAS Collaboration Collaboration, “Measurement of the azimuthal anisotropy of charged particle production in Xe+Xe collisions at $\sqrt{s_{\text{NN}}}=5.44$ TeV with the ATLAS detector”, Technical Report ATLAS-CONF-2018-011, CERN, Geneva, May, 2018.
- [236] ALICE Collaboration, “Centrality and pseudorapidity dependence of the charged-particle multiplicity density in Xe–Xe collisions at $\sqrt{s_{\text{NN}}} = 5.44$ TeV”, *Phys. Lett.* **B790** (2019) 35–48, doi:10.1016/j.physletb.2018.12.048, arXiv:1805.04432.
- [237] CMS Collaboration, “Pseudorapidity distributions of charged hadrons in xenon-xenon collisions at $\sqrt{s_{\text{NN}}} = 5.44$ TeV”, arXiv:1902.03603.
- [238] ALICE Collaboration, “Underlying Event measurements in pp collisions at $\sqrt{s} = 0.9$ and 7 TeV with the ALICE experiment at the LHC”, *JHEP* **07** (2012) 116, doi:10.1007/JHEP07(2012)116, arXiv:1112.2082.
- [239] CMS Collaboration, “Measurement of the underlying event activity using charged-particle jets in proton-proton collisions at $\sqrt{s} = 2.76$ TeV”, *JHEP* **09** (2015) 137, doi:10.1007/JHEP09(2015)137, arXiv:1507.07229.
- [240] ATLAS Collaboration, “Measurement of charged-particle distributions sensitive to the underlying event in $\sqrt{s} = 13$ TeV proton-proton collisions with the ATLAS detector at the LHC”, *JHEP* **03** (2017) 157, doi:10.1007/JHEP03(2017)157, arXiv:1701.05390.
- [241] CDF Collaboration, “Study of the energy dependence of the underlying event in proton-antiproton collisions”, *Phys. Rev.* **D92** (2015), no. 9, 092009, doi:10.1103/PhysRevD.92.092009, arXiv:1508.05340.
- [242] CDF Collaboration, “Charged Jet Evolution and the Underlying Event in $p\bar{p}$ Collisions at 1.8 TeV”, *Phys. Rev.* **D65** (2002) 092002, doi:10.1103/PhysRevD.65.092002.
- [243] C. Loizides and A. Morsch, “Absence of jet quenching in peripheral nucleus–nucleus collisions”, *Phys. Lett.* **B773** (2017) 408–411, doi:10.1016/j.physletb.2017.09.002, arXiv:1705.08856.
- [244] ALICE Collaboration, “Analysis of the apparent nuclear modification in peripheral Pb-Pb collisions at 5.02 TeV”, arXiv:1805.05212.

- [245] ALICE Collaboration, “Centrality dependence of particle production in p-Pb collisions at $\sqrt{s_{NN}}=5.02$ TeV”, *Phys. Rev.* **C91** (2015), no. 6, 064905, doi:10.1103/PhysRevC.91.064905, arXiv:1412.6828.
- [246] PHENIX Collaboration, “Centrality categorization for $R_{p(d)+A}$ in high-energy collisions”, *Phys. Rev.* **C90** (2014), no. 3, 034902, doi:10.1103/PhysRevC.90.034902, arXiv:1310.4793.
- [247] PHENIX Collaboration, “Centrality-dependent modification of jet-production rates in deuteron-gold collisions at $\sqrt{s_{NN}}=200$ GeV”, *Phys. Rev. Lett.* **116** (2016), no. 12, 122301, doi:10.1103/PhysRevLett.116.122301, arXiv:1509.04657.
- [248] J. Jia, “Influence of the nucleon-nucleon collision geometry on the determination of the nuclear modification factor for nucleon-nucleus and nucleus-nucleus collisions”, *Phys. Lett.* **B681** (2009) 320–325, doi:10.1016/j.physletb.2009.10.044, arXiv:0907.4175.
- [249] Z. Citron et al., “Future physics opportunities for high-density QCD at the LHC with heavy-ion and proton beams”, in *HL/HE-LHC Workshop: Workshop on the Physics of HL-LHC, and Perspectives at HE-LHC Geneva, Switzerland, June 18-20, 2018*. 2018. arXiv:1812.06772.
- [250] S. Cao, T. Luo, G.-Y. Qin, and X.-N. Wang, “Heavy and light flavor jet quenching at RHIC and LHC energies”, *Phys. Lett.* **B777** (2018) 255–259, doi:10.1016/j.physletb.2017.12.023, arXiv:1703.00822.
- [251] S. Shi, J. Liao, and M. Gyulassy, “Probing the Color Structure of the Perfect QCD Fluids via Soft-Hard-Event-by-Event Azimuthal Correlations”, *Chin. Phys.* **C42** (2018), no. 10, 104104, doi:10.1088/1674-1137/42/10/104104, arXiv:1804.01915.
- [252] L.-G. Pang, Y. Hatta, X.-N. Wang, and B.-W. Xiao, “Analytical and numerical Gubser solutions of the second-order hydrodynamics”, *Phys. Rev.* **D91** (2015), no. 7, 074027, doi:10.1103/PhysRevD.91.074027, arXiv:1411.7767.
- [253] J. D. Bjorken, “Highly Relativistic Nucleus-Nucleus Collisions: The Central Rapidity Region”, *Phys. Rev.* **D27** (1983) 140–151, doi:10.1103/PhysRevD.27.140.
- [254] H. Niemi, K. J. Eskola, and R. Paatelainen, “Event-by-event fluctuations in a perturbative QCD + saturation + hydrodynamics model: Determining QCD matter shear viscosity in ultrarelativistic heavy-ion collisions”, *Phys. Rev.* **C93** (2016), no. 2, 024907, doi:10.1103/PhysRevC.93.024907, arXiv:1505.02677.

- [255] K. J. Eskola, K. Kajantie, P. V. Ruuskanen, and K. Tuominen, “Scaling of transverse energies and multiplicities with atomic number and energy in ultrarelativistic nuclear collisions”, *Nucl. Phys.* **B570** (2000) 379–389, doi:10.1016/S0550-3213(99)00720-8, arXiv:hep-ph/9909456.
- [256] C. Shen et al., “The iEBE-VISHNU code package for relativistic heavy-ion collisions”, *Comput. Phys. Commun.* **199** (2016) 61–85, doi:10.1016/j.cpc.2015.08.039, arXiv:1409.8164.
- [257] I. Vitev, “Testing the mechanism of QGP-induced energy loss”, *Phys. Lett.* **B639** (2006) 38–45, doi:10.1016/j.physletb.2006.05.083, arXiv:hep-ph/0603010.
- [258] JETSCAPE Collaboration, “JETSCAPE v1.0 Quickstart Guide”, *Nucl. Phys.* **A982** (2019) 615–618, doi:10.1016/j.nuclphysa.2018.09.033, arXiv:1807.09615.
- [259] JETSCAPE Collaboration, C. Park, “Multi-stage jet evolution through qgp using the jetscape framework: inclusive jets, correlations and leading hadrons”, January, 2019. 2nd JETSCAPE Winter School and Workshop 2019.
- [260] ALICE Collaboration, “Transverse momentum spectra and nuclear modification factors of charged particles in Xe-Xe collisions at $\sqrt{s_{\text{NN}}} = 5.44$ TeV”, *Phys. Lett.* **B788** (2019) 166–179, doi:10.1016/j.physletb.2018.10.052, arXiv:1805.04399.