

Abstract

Inclusive Jet Measurements in Pb-Pb Collisions at $\sqrt{s_{\text{NN}}} = 5.02 \text{ TeV}$ with ALICE

James Mulligan

2019

Droplets of deconfined quarks and gluons, known as the quark-gluon plasma, are produced experimentally in ultra-relativistic heavy-ion collisions. Studying this deconfined matter may allow insight into a variety of open questions about the high temperature regime of QCD and the emergent behaviors of QCD. One major effort to probe the quark-gluon plasma is the study of high-momentum jets produced in an initial high momentum-transfer scattering of a heavy-ion collision. Measurements have demonstrated that by traversing the dense plasma, jets are modified in several ways, including that jet yields are suppressed in heavy-ion collisions relative to proton-proton collisions. The ALICE detector at the Large Hadron Collider reconstructs jets with high-precision tracking of charged particles combined with particle information from the electromagnetic calorimeter, achieving a unique kinematic range of jets extending to low jet momenta. This thesis describes inclusive jet measurements in Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 5.02 \text{ TeV}$ with ALICE, which constitute the first such full jet measurements at low transverse jet momentum at this collision energy. These measurements are compared to several theoretical predictions, and will help constrain models of jet energy loss.



Inclusive Jet Measurements in Pb-Pb Collisions at
 $\sqrt{s_{\text{NN}}} = 5.02 \text{ TeV}$ with ALICE

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
James Mulligan

Dissertation Director: Professor John W. Harris

May 2019

Copyright © 2019 by James Mulligan
All rights reserved.

Contents

Acknowledgments	vi
1 Introduction	1
1.1 The Standard Model of Particle Physics	1
1.2 Quantum chromodynamics	2
1.2.1 Basics of QCD	2
1.2.2 Chiral symmetry in QCD	5
1.2.3 Scattering and factorization	6
1.3 The quark-gluon plasma	9
1.3.1 A strongly coupled relativistic fluid	12
1.3.2 Physics goals	14
1.4 Heavy-ion physics	15
1.4.1 Evolution of the quark-gluon plasma	15
1.4.2 Experimental variables	19
1.4.3 Signatures of the quark-gluon plasma	21
1.5 Jets in heavy-ion collisions	25
1.5.1 Measuring jets in heavy-ion collisions	27
1.5.2 Theory of jet modification in heavy-ion collisions	30
1.5.3 Phenomenology of jet modification in heavy-ion collisions	32
2 The ALICE Detector	39
2.1 The Large Hadron Collider	42
2.2 ALICE tracking system	45

2.2.1	ITS	45
2.2.2	TPC	46
2.3	ALICE EMCal	47
2.4	ALICE Upgrades	51
2.4.1	TPC Upgrade	51
3	Inclusive jet measurements in Pb–Pb collisions	53
3.1	Experimental data	55
3.1.1	Analysis selections	55
3.1.2	Data quality assurance	62
3.2	Jet reconstruction	68
3.2.1	Jet clustering	68
3.2.2	Background subtraction	69
3.2.3	Background fluctuations	70
3.2.4	Impact of jet cuts on the measured spectra	74
3.2.5	Jet performance	77
3.3	Performance studies	80
3.3.1	Centrality-dependence of the EMCal cluster spectrum	80
3.3.2	EMCal cell thresholds	86
3.4	Unfolding the jet p_T spectrum	89
3.4.1	Constructing the response matrix	90
3.4.2	Performing the unfolding	97
3.4.3	Thermal closure test	100
3.5	Systematics	104
3.5.1	Correlated uncertainties	104
3.5.2	Shape uncertainties	106
3.5.3	Uncertainties on the jet cross-section ratio	109
3.6	Results	111
3.6.1	Jet spectra	111
3.6.2	Jet R_{AA}	118

3.6.3	Jet cross-section ratio	120
3.7	Comparison to theory	122
3.7.1	Summary of theoretical models	122
3.8	Discussion and outlook	127
Appendix		130
A TPC upgrade R&D		130
A.1	2-GEM+MMG	130
A.2	Multi-layer extended gating grid	132
A.2.1	Simulation Configurations	132
A.2.2	Simulation Results – ALICE TPC	135
A.2.3	Simulation Results – STAR TPC	140
A.2.4	Discussion	142
B Inclusive jet analysis		144
B.1	Analysis code	144
B.2	Combining $p_{T,\text{hard}}$ bins	145
B.2.1	Scaling $p_{T,\text{hard}}$ bins	145
B.2.2	Removing outliers	145
B.3	EMCal cell threshold studies, additional details	147
B.4	$R = 0.4$ unfolding details	153
References		155

Acknowledgments

I am deeply grateful to my advisor, Professor John Harris, and my “co-advisor”, Professor Helen Caines, for the opportunity to join their research group. John has consistently supported and advocated for me throughout my four-and-a-half years as his student, especially during the final stretch of this thesis. And Helen is one of the brightest physicists I have had the pleasure to work with; in virtually every physics conversation I’ve had with her, I’ve come away learning something new.

I would like to thank Professors Karsten Heeger and Thomas Appelquist for serving on my thesis committee and providing valuable input in the course of this important tradition.

To my colleagues in ALICE. The inclusive jet analysis involved crucial input from Dr. Leticia Cunqueiro Mendez, Dr. Peter Jacobs, and Dr. Marco van Leeuwen. To Raymond and Salvatore, from whom I learned nearly all of my software skills, and whose work ethic and kindness I was able to enjoy over many office chats and social outings.

To my classmates at Yale. To Aleks, Will, Zack, Ben: The insight, curiosity, and humor I have witnessed in all of you is incomparable thus far in my life. To Dili, for being above all else a valuable friend. And to FZ, who is like a brother to me, and who inspires me to set my sights above all others.

To my family. To my mother, for her unparalleled love and kindness and support. To my dad, for his down-to-earth resourcefulness and belief that life is *good*. And to my brother, who has been teaching me since day one.

To Eliane. You are singlehandedly responsible for my biggest growth in these last years, not only as a physicist but as a human. I am excited to see where my trajectory will go, and I am certain it will be somewhere that would not have been possible without you.

And to the thinkers who have shaped me, from Feynman to Hitchens. I am profoundly lucky for the opportunity to participate firsthand in what Feynman called “the great adventure of our time.”

Chapter 1

Introduction

Quantum chromodynamics (QCD) is the strongest yet most poorly understood of the known forces of nature. It is described in the Standard Model of Particle Physics, and while the QCD Lagrangian is known, the behaviors that emerge from it are largely ill-understood. Among these is the nature of deconfined QCD matter known as the quark-gluon plasma. In this chapter, I present an introduction to QCD and the quark-gluon plasma, and the use of jets to learn about them.

1.1 The Standard Model of Particle Physics

The Standard Model of Particle Physics is a quantum field theory describing the strong, weak, and electromagnetic forces of nature. It posits that the universe is populated with a selection of fundamental fields that obey an assortment of symmetries, including the non-Abelian local gauge symmetry $SU(3)_{\text{Strong}} \times SU(2)_{\text{Weak}} \times U(1)_{\text{EM}}$. Elementary particles are excitations of these fields, and can be classified into groups: quarks, leptons, gauge bosons, and the Higgs scalar boson. The quarks and leptons are fermions, while the gauge bosons and Higgs are, unsurprisingly, bosons. Quantum field theories naturally arise from the need to simultaneously describe quantum mechanics and special relativity, accommodating the possibility that a temporary energy fluctuation ΔE (from the Heisenberg uncertainty principle of quantum mechanics) can be converted to a new particle (by the mass-energy duality of special relativity). The Standard Model successfully describes all observed elementary

particles and their interactions up to energies $\mathcal{O}(\approx \text{TeV})$, as well as their resulting bound states. It is the most precise theory in the history of human science, with the theoretical prediction of the magnetic moment of the electron in agreement with experimental measurements to at least 10 significant digits [1]. The free parameters of the Standard Model are the quark masses, the lepton masses, two gauge boson masses, the Higgs mass, the Higgs vacuum expectation value, three couplings, and the strong and weak mixing angles and CP-violating phases.

While the Standard Model is extremely successful as a theory, certain theoretical and observational facts demand that the Standard Model is only a “low-energy” effective model, ignorant to the physics at scales larger than $\mathcal{O}(\approx \text{TeV})$. The presence of dark matter and dark energy are not satisfactorily postulated by the Standard Model. The relatively small mass of the Higgs boson does not naturally arise. It is not clear how the observed matter-antimatter asymmetry originated. There is no clear reason why the strong CP-phase essentially vanishes. Experiments at the Large Hadron Collider (LHC) are pushing the boundary of this regime to higher and higher energies, while other experiments such as in neutrinos, dark matter, and various precision searches are looking for hints of new physics elsewhere. For the remainder of this thesis, however, we will focus not on general questions about the Standard Model, but rather on QCD and its emergent complexities.

1.2 Quantum chromodynamics

1.2.1 Basics of QCD

QCD is a quantum field theory with an $SU(3)$ gauge symmetry, populated by six elementary quark flavors (up, down, strange, charm, bottom, top) and gauge bosons known as gluons. The quarks and gluons carry “color” charge, with the gluon gauge fields acting as “force carriers” that can rotate a quark’s color. Gluons carry a larger color charge than quarks, $\frac{C_g}{C_q} = \frac{9}{4}$, and so gluons propagate less freely. QCD exhibits an approximate flavor isospin symmetry, in which the quark flavors can be grouped into $SU(2)$ doublets, e.g. $\begin{pmatrix} u \\ d \end{pmatrix}$, where they behave as two isospin states of the same particle. There are a wide variety of references describing QCD in great detail [2–4].

The QCD Lagrangian is:

$$\mathcal{L}_{QCD} = -\frac{1}{4g^2} F_{\mu\nu}^a F^{a\mu\nu} + \sum_{j=1}^6 \bar{q}_j (i\gamma^\mu D_\mu - m_j) q_j,$$

where q_j is the quark field of flavor j , m_j is the quark mass, g is the strong coupling constant, $D_\mu = \partial_\mu - iA_\mu$ is the covariant gauge derivative, A_μ is the gluon gauge field, $F_{\mu\nu}^a$ is the gluon field strength tensor, and $a \in \{1, 2, \dots, 8\}$ indexes the $SU(3)$ gauge group.¹ The fundamental parameters of the theory are the dimensionless coupling g and the quark masses m_j .

The coupling g depends on the energy scale of an interaction, due to screening and anti-screening from loop diagrams that increasingly appear with higher resolution scale. The running of the coupling $\alpha_s(\mu) \equiv g(\mu)^2/4\pi$ at leading order in perturbative QCD (pQCD) for a given momentum transfer Q^2 and renormalization scale μ turns out to be given by:

$$\alpha_s(Q^2) = \frac{4\pi}{(11 - \frac{2}{3}n_f) \log(Q^2/\mu^2)},$$

where $n_f = 6$ is the number of fermions. Note that despite the appearance of μ here, physical observables cannot depend on μ . Measured values of α_s are shown in Fig. 1.1. From these, we can see that $\alpha_s \rightarrow 0$ as $Q^2 \rightarrow \infty$. This property is known as asymptotic freedom: the coupling tends to 0 as the momentum transfer Q^2 grows.

The perturbative calculation of the running of the coupling is only valid, however, in the regime where $g \ll 1$. When this is the case, one can perform theoretical calculations in QCD using perturbation theory, and in fact perturbative calculations have been shown to agree with measurements over a wide kinematic range. The elementary Feynman vertices of pQCD are a quark-quark-gluon vertex, a 3-gluon vertex, and a 4-gluon vertex; the color charge is conserved at each vertex. Note that since gluons carry color charge, they interact among themselves, unlike photons in QED.

When g becomes $\mathcal{O}(1)$, however, the perturbative approach breaks down, and in fact

1. There is also a CP-violating term that in principle should be included in the Lagrangian, characterized by a strong CP angle θ . However, θ has been measured to be essentially vanishing, and the term is neglected here. It is not known why the strong force exhibits little-to-no CP violation – this is known as the strong CP problem.




Figure 1.1: Running of the strong coupling α_s as a function of the resolution scale $Q \equiv \sqrt{Q^2}$. The data points show values extracted from experimental measurements, and the band shows a particular theoretical calculation [5].

there is no known analytical way to compute the running of the coupling or any other QCD dynamics in this case. Empirically, however, it has been established that the force between two partons (i.e. quarks or gluons) gets stronger at low momentum-transfer, or large distance – until it becomes favorable to create new partons rather than further separate the partons. This behavior is arguably hinted at by the perturbative calculation shown in Fig. 1.1, but in principle the coupling in the small- Q regime could have been anything; the analytical form of the coupling in the non-perturbative regime remains unknown. A scale $\Lambda_{QCD} \approx 217$ MeV is defined as the scale at which α_s becomes $\mathcal{O}(1)$.² This non-perturbative regime gives rise to the fact that the elementary constituents of the strong interaction have never been observed in isolation – a property known as color confinement. Rather, they exist in color-neutral bound states known as hadrons. The strong force exhibited in nuclei is therefore short distance, since it is only the “van der Waals” force of the strong interaction. The detailed mechanism of confinement remains unknown, and is arguably the biggest open

². This is known as dimensional transmutation, since a dimensionful scale Λ_{QCD} emerges from a dimensionless scale g .

question in QCD.

Therefore, despite knowing the exact Lagrangian of QCD, it is not known how to generally solve the resulting equations. Many basic characteristics of QCD consequently remain poorly understood. The mechanism by which confinement occurs is not known. It is not known how to analytically calculate the hadron masses. It is not known why the spin of the proton is $\frac{1}{2}$. It is not known why certain observed hadronic states exist (the XYZ states), or why certain allowable QCD bound states have not been observed (“glueballs”, qqg , etc.). It is not known if there is a critical point in the QCD phase diagram. Most of these open questions are about the nature of emergent QCD characteristics, which cannot be easily determined from the QCD Lagrangian.

At sufficiently high temperature T , one expects that confined hadronic matter can be “melted” into a deconfined state. And if T is high enough, that the coupling between partons will become small due to asymptotic freedom. The study of this deconfinement transition will be discussed extensively later in this chapter.

1.2.2 Chiral symmetry in QCD

Given a massless quark isospin doublet ψ , the quark part of the QCD Lagrangian can be written:

$$\mathcal{L}_{QCD} = \bar{\psi}_L (i\gamma^\mu D_\mu) \psi_L + \bar{\psi}_R (i\gamma^\mu D_\mu) \psi_R,$$

where the projection operators $P_{L/R} = \frac{1}{2}(1 \mp \gamma^5)$ have been used to decompose $\psi_{L/R} = P_{L/R}\psi$. This Lagrangian is invariant under the $SU(2)_L \times SU(2)_R$ chiral symmetry $\psi_{L/R} \rightarrow e^{-i\theta_{L/R}^a t^a} \psi_{L/R}$, where θ^a is a continuous parameter and t^a is a generator of the $SU(2)$ group. Note that the chiral symmetry transformation acts independently on the left-handed and right-handed quarks. That is, the Lagrangian is unchanged when the left-handed and right-handed components independently transform.

This symmetry, however, is broken in three ways. First, chiral symmetry is explicitly broken by the fact that the quark masses are non-zero. Second, chiral symmetry in QCD is spontaneously broken – while the Lagrangian exhibits the symmetry, the ground state of the theory does not. That is, the “chiral condensate” has nonzero vacuum expectation value:

$\langle 0|\bar{\psi}\psi|0\rangle \neq 0$. Specifically, in the two-quark model, the $SU(2)_L \times SU(2)_R$ chiral symmetry, which can be decomposed into $SU(2)_V \times SU(2)_A$, a vector part (which treats L, R the same) and an axial part (which treats L, R differently), is broken to the vector subgroup $SU(2)_V$. The three associated Goldstone bosons are the charged and neutral pions, which obtain mass due to the explicit quark masses.³ This is reflected also in the hadronic mass spectrum, both because the pions are by far the lightest hadrons, and because there are large mass splittings between pseudo-scalar mesons and their chiral vector partners, such as the ρ and a_1 , unlike the proton and neutron. The cause of spontaneous chiral symmetry breaking in QCD is not well-understood, although lattice QCD predicts a crossover phase transition at low- T , similar to the confinement transition. The chiral symmetry restoration transition and the deconfinement transition are distinct, although they are expected to be related. The exact relationship between them, including at what T the chiral symmetry restoration transition occurs, is unknown.

The third way chiral symmetry in QCD is broken is by a quantum anomaly – while the Lagrangian exhibits the symmetry, the path integral measure does not. This is known as the “chiral anomaly”. It is an open question whether the anomalous breaking of chiral symmetry is restored at high T [6].

1.2.3 Scattering and factorization

As discussed in Section 1.2.1, the one regime in which QCD scattering amplitudes can be reliably calculated is for large Q^2 . However, experimentally we cannot collide two calibrated parton beams together – instead, we collide hadrons, which are non-perturbative systems. Fortunately, QCD Factorization allows us in certain cases to separate the non-perturbative initial state, described by the Parton Distribution Function (PDF), from the perturbative high- Q^2 (“hard”) scattering, and the non-perturbative final state, described by the Fragmentation Function (FF).

3. When the quark masses are taken to 0, the pion masses also go to 0 – but not for any other hadrons (in the 2-quark model).

The cross-section for the process $pp \rightarrow \text{jet} + X$, for example, has the form:

$$\sigma^{pp \rightarrow \text{jet} + X} = \sum_{a,b} \int_0^1 dx_1 \int_0^1 dx_2 f_{a/p}(x_1, \mu) f_{b/p}(x_2, \mu) \sigma_{ab \rightarrow c+X}(x_1 p_1, x_2 p_2, \mu),$$

where we consider partons a, b from protons 1,2 scattering to an outgoing jet c . We denote the fraction of the hadron momentum carried by a parton by $x \equiv \frac{p_{\text{parton}}}{p_{\text{hadron}}}$. That is, we integrate the likelihood to have a parton with a given x , and then hard-scatter with another parton. If instead the final state involved a hadron, we would also integrate over the likelihood of the outgoing parton to fragment into that given final state.

Accordingly, we must describe the abundance of partons in the initial hadrons, as well as the likelihood of the outgoing partons to fragment into a given set of hadrons – that is, the PDFs and FFs.

Parton Distribution Functions

In general, a hadron consists of valence quarks as well as gluons and “sea” quarks (generated inside the hadron by $g \rightarrow \bar{q}q$). The PDF $f_i(x, Q^2)$ is the probability of a parton i having a fraction x of the hadron momentum, given that the square of the exchanged four-momentum is Q^2 . Note that the PDF depends on Q^2 since with a higher resolution scale, one sees more loop diagrams to scatter from.

The momentum transfer Q^2 sets the resolution scale with which we probe the proton. At low Q^2 , i.e. coarse resolution scale, the PDF is dominated by the three valence quarks, each carrying $\approx \frac{1}{3}$ of the proton momentum. As Q^2 increases, and we probe the short-range structure of the proton, the softer contributions of sea quarks and gluons increases. At very low x , the PDF is dominated by soft gluons. This can be understood since at a coarse resolution scale, we may see only a propagating quark, but as we zoom in, we will see many soft gluon radiations from the quark, which carry part of the quark’s momentum (and therefore a strong enhancement at low- x). Note that as $\sqrt{s}$ increases, the typical Q^2 increases, and the proton increasingly consists of low- x gluons. Accordingly the pp cross-section increases with $\sqrt{s}$, as the PDFs contain more gluons, and are therefore more likely to interact.




Figure 1.2: Proton PDF as a function of x at $Q^2 = 10 \text{ GeV}^2$ and $Q^2 = 10^4 \text{ GeV}^2$, from the MSTW2008 parameterization. Note that at higher Q^2 , the gluon and sea quark PDFs grow dramatically [7].

The PDF is a non-perturbative object, and cannot be calculated directly (although lattice QCD attempts are progressing). However, given a PDF at one scale $f_i(x, Q_0^2)$, it can be evolved to another $f_i(x, Q^2)$ using the perturbative DGLAP equations to a given fixed order. The PDF can be measured by deep inelastic scattering (DIS) experiments in which a high-energy electron is scattered on a hadron as a function of x, Q^2 . Figure 1.2 shows the proton PDF at $Q^2 = 10 \text{ GeV}^2$ and $Q^2 = 10^4 \text{ GeV}^2$.

In a nucleus, the PDF of a nucleon is different than that of a free nucleon, due to a variety of effects. Shadowing causes a suppression of the nuclear PDF (nPDF) at $x < 0.1$. Anti-shadowing causes an enhancement at slightly higher x . The EMC effect gives suppression at intermediate $0.5 < x < 0.8$. Fermi motion causes enhancement at high- x . The cause of shadowing, anti-shadowing, and the EMC effect are not yet clear. It is also believed that at very low x , the gluon density cannot continue to grow without breaking unitarity, and so there must be an onset of gluon saturation, in which gluon splitting $g \rightarrow gg$ is balanced by gluon fusion $gg \rightarrow g$. The observation of gluon saturation is one of the main goals of the proposed Electron Ion Collider [8].

Fragmentation Functions

The FF $D_p^h(z)$ is the probability density that a final state parton p hadronizes into a mean number $D_p^h(z)dz$ of hadrons per dz , where $z = \frac{p_{\text{hadron}}}{p_{\text{parton}}}$. The FF is also a non-perturbative object, and cannot be described from fundamental QCD, but it can be evolved perturbatively from one scale to another as for PDFs. The FF involves a parton shower, which is a perturbative process governed by the QCD splitting function, to fragment hard partons to soft partons, as well as a non-perturbative fragmentation and hadronization mechanism. There are several phenomenological models of hadronization, such as the Lund string model or cluster fragmentation models, which will be discussed later in this chapter.

Fragmentation functions are approximately independent of the parton p_T . FFs for gluons are significantly softer than for quarks, due to their enhanced color factor. Moreover, heavy quark FFs are notably harder than the light quark fragmentation functions, due to the large mass of the heavy quarks.

1.3 The quark-gluon plasma

At typical temperatures in the universe today, QCD matter is typically confined into colorless hadrons. However, at extreme temperature or pressure, hadronic matter ceases to exist in lieu of a deconfined state of QCD matter, known as the quark-gluon plasma. Experimentally, one can achieve sufficiently high energy densities to reach such a state by colliding heavy-ion nuclei at ultra-relativistic energies. This was done in the past decades, and by the mid-2000s a consensus developed that a new state of matter with delocalized partons had been created with temperatures $\mathcal{O}(10^{12} \text{ K})$. The experiments and physical signatures that have driven this consensus will be explained in the next section. First, let us consider in a bit more depth what the quark-gluon plasma is, and why it is interesting.

The QCD phase diagram as a function of the temperature T and the net baryon density is shown in Fig. 1.3. At typical temperatures and baryon densities in the universe, QCD matter is confined to color-neutral hadrons. At high- T or large baryon density, however, a deconfinement transition is expected, as well as a chiral symmetry transition. In nature, a quark-gluon plasma at extremely high- T and low baryon density existed in the early universe




Figure 1.3: Phase diagram of QCD [9]. The curve “RHIC,LHC” denotes that both facilities produce high- T , low- μ_B QCD matter when operating at their top energies (with the LHC creating lower- μ_B than RHIC). Note, however, that the RHIC Beam Energy Scan, in which $\sqrt{s_{NN}}$ is decreased in order to search for the critical point, is able to reach much larger μ_B than the top RHIC energy.

for the vast majority of $\mathcal{O}(10 \mu s)$ after the Big Bang.⁴ Additionally, it is speculated that a quark-gluon plasma at low- T and large baryon density may be present in the cores of neutron stars.

Commonly, the chemical potential μ would be plotted on such a phase diagram. Recall that chemical potential denotes the change in the system’s energy associated with adding a particle to it, and describes a system’s tendency to create or destroy particles. Low μ denotes freedom to easily create more particles (e.g. photons in QED, for which the energy does not depend on number of photons), while high- μ denotes a strong aversion of the system to create more particles. In QCD, net baryon number is a conserved quantity, and so net baryons are not created or destroyed – rather, since we start with a given initial net baryon number (from the colliding nuclei, composed of baryons as opposed to anti-baryons), only their density can change. The baryon chemical potential μ_B is therefore

4. In fact, for a time there was excitement that if there was a formal deconfinement phase transition in the early universe, its associated fluctuations would leave an imprint that would be observable today. However, it turns out that in the high- T , low baryon density regime, there is not formally a phase transition, but a crossover.

a measure of the net baryon density, with high- μ_B corresponding to a high concentration of baryons. At higher $\sqrt{s_{\text{NN}}}$, μ_B becomes smaller because the collision produces more $q\bar{q}$ pairs, and so the baryon density is diluted (for fixed baryon number), which results in lower baryon density and thus lower μ_B . Accordingly, the LHC produces high- T and low- μ_B QCD matter, whereas RHIC produces high- T and somewhat higher μ_B QCD matter. In a heavy-ion collision that produces a quark-gluon plasma, the deconfined phase is populated for only a short time $\mathcal{O}(10 \text{ fm}/c)$, before expanding and cooling back to the hadronic phase. In particular, the QCD medium should be thought of (approximately) as having a given T and μ_B at each point in space and time.

Since the transition from hadronic matter to deconfined QCD matter occurs in the non-perturbative regime, there is no known way to analytically calculate the transition temperature T_c or its properties. However, the method of lattice QCD can be used to predict certain properties. Lattice QCD discretizes spacetime, numerically solves the equations of QCD (under certain assumptions, e.g. about the quark masses), and then takes the continuum limit of the lattice spacing $\rightarrow 0$. In particular, it has been used to predict the transition temperature and the nature of the transition in the high- T , low- μ_B regime. Fig 1.4 shows a calculation of $\frac{\varepsilon}{T^4}$ vs. T , which exhibits a sharp rise at a temperature T_c , where $T_c \approx 154 \text{ MeV}$. Note that $\frac{\varepsilon}{T^4}$ is approximately proportional to the number of relativistic degrees of freedom – this rise is precisely the transition from hadronic degrees of freedom to partonic degrees of freedom. Lattice QCD predictions also state that the high- T , low- μ_B transition is not a phase transition, but rather a continuous crossover of phases, at which a transition temperature T_c describes a region where the thermodynamic properties of the medium rapidly change. It is speculated, but not yet known, that this becomes a first-order transition at higher μ_B .

Lattice QCD predictions are currently limited to the low μ_B regime. Other theoretical predictions, such as the phenomenological bag model or various effective theories, are used to predict the nature of the low- T , high- μ_B transition, expected to occur at roughly five times the density of hadronic nuclear matter, $\varepsilon_{\text{hadron}} \approx \frac{\Lambda_{\text{QCD}}}{V_{\text{nucleon}}}$.




Figure 1.4: Lattice QCD calculation of several thermodynamic quantities including $\frac{\epsilon}{T^4}$ vs. T , which exhibits a rapid rise in the degrees of freedom near a crossover temperature T_c [10].

1.3.1 A strongly coupled relativistic fluid

When we heat up QCD matter, deconfinement is achieved as the hadronic bound states are melted. But there are two crucial questions that must be answered if we are to understand the nature of hot QCD matter. First, what is the structure of deconfined QCD matter at a given T ? Does it consist of individual quarks and gluons, or are there quasi-particles, such as screened color centers (e.g. “dressed” quarks or gluons) or perhaps more exotic structures, which are the relevant constituents of the system? If quasi-particles do exist at certain T , at what T do they melt? And what is the structure of QCD matter during the transition? This is a particularly tantalizing question, since it may give us a view of the mechanism of confinement. Second, what is the coupling in the quark-gluon plasma at a given T ? This of course depends on the first question – the coupling between *what*? In the limit of ultra-high temperatures, where we expected to have individual partons, we expect the coupling in the quark-gluon plasma to become weak, since momentum transfers between partons become large and we obtain asymptotic freedom. But what is the case at T that have been measured experimentally, up to $\approx 3-4 T_c$? And if there is a quasi-particle

structure at lower T , such as screened color centers, will the coupling weaken, analogous to the coupling between nucleons?

Experiments have provided a partial answer for those T that have been measured, up to $\approx 3 - 4 T_c$. To do so, the experiments have extracted the quark-gluon plasma's shear-viscosity to entropy-density ratio, η/s , from a category of observables called anisotropic flow. The experimental observable will be described in the next section – here, we discuss the implications of this measurement. The value of η/s in fact directly relates to the coupling in a fluid [11].

Intuitively, viscosity describes the thickness of a fluid. Shear viscosity can be defined classically as follows: Consider two plates of area A separated by a distance y with a fluid in between, with the bottom plate fixed and the top plate free to move. The shear viscosity η is defined from the force F needed to push the top plate at velocity u : $F = \eta A \frac{u}{y}$. From this setup, one can see that shear viscosity arises from friction between neighboring particles in adjacent layers moving at different velocities, caused by molecular diffusion between layers, resulting in momentum transfer. Low viscosity, or for a relativistic fluid η/s , then, is obtained when the coupling in the fluid is strong – a low- η/s fluid is a highly ordered fluid, where the coupling dominates the momentum diffusion.

For the quark-gluon plasma, η/s turns out to be the smallest of any liquid observed.⁵ In fact, the value of η/s in the QGP is very close to a conjectured quantum lower limit that any fluid can achieve [12]. The conjecture uses the AdS/CFT correspondence to compute η/s in a strongly-coupled supersymmetric Yang-Mills $SU(4)$ field theory using a weakly-coupled 5D string theory, and asserts that the lower quantum limit is $\frac{\eta}{s} = \frac{\hbar}{4\pi k_B}$.⁶ Measurements (described in Section 1.4.3) show that $\eta/s \approx 0.1 - 0.2 \frac{\hbar}{k_B}$. It is unclear why QCD appears to coincide with this value, or if there is a deeper reason. Regardless, the small η/s definitively means that the QGP is a strongly-coupled liquid – and a liquid with the lowest η/s ever observed.⁷ This is almost certainly the most important and unexpected discovery to date

5. η is in fact very large – larger than tar. However, the QGP also has a very high entropy density.

6. Subsequent works (e.g. [13]) demonstrate that this bound can be violated.

7. The extraction of η/s is model-dependent, but regardless of details the QGP is the lowest η/s fluid ever observed. The next smallest η/s observed in nature is from strongly-coupled ultra-cold Fermi gases, with a value of η/s several times larger than in the quark-gluon plasma [11].

in the subfield of quark-gluon plasma physics.

We can then summarize what is known about the two crucial questions outlined above. At temperatures above the transition, but not at ultra high- T , the coupling is evidently strong. Here, we may or may not have some quasiparticle structure; we may really have just individual quarks and gluons, exchanging Q^2 such that the coupling remains strong, but with T sufficiently large to melt any quasiparticles. In the ultra high- T limit, we expect a weakly-coupled regime where QCD matter consists of free partons. At temperatures near the transition, we do not know the structure of QCD matter nor its coupling; it may consist of quasiparticles, and the interactions between them will govern the behavior – likely at a weaker coupling.

1.3.2 Physics goals

The quark-gluon plasma is a rich laboratory for the physics of QCD and more [14]. It provides arguably the most compelling system from which to understand how QCD confinement arises. This can be done by studying the structure and coupling of the QCD matter as a function of T , as outlined above, and by a variety of other methods, such as constraining models of parton hadronization. Quark-gluon plasma physics is part of a larger effort to study the structure of hadronic states, including the proton. Perhaps quasi-particle structure in the QGP will deliver insight to how hadrons form and why they have the properties that they do. However, there are a number of complementary questions which also make a strong physics case for studying the quark-gluon plasma.

It is unknown where the spontaneous chiral symmetry restoration transition occurs in QCD, and if anomalous chiral symmetry restoration also occurs. Moreover, it is unknown if there is a deconfinement critical point in the QCD phase diagram. These are major fundamental open questions about the basic behavior of QCD.

We have learned that QCD matter at 10^{12} K is a liquid. In fact, the QGP is the only known relativistic fluid. What are the collective properties of this fluid? The mean free path, the speed of sound, the transport coefficients, and so on? If it has quasi-particle structure, is there interesting physics to be learned from their interactions? Studying the fluid's properties and how they arise is a new angle of physics in its own right.

Why is the η/s value measured in the QGP close to the quantum limit in infinitely coupled supersymmetric $SU(4)$ gauge theory, according to the AdS/CFT correspondence? This observation can be argued to lend credence to the AdS/CFT correspondence, and could motivate deeper insights.

Perhaps above all, studying as rich a system as this is likely to deliver insight to unforeseen questions.

1.4 Heavy-ion physics

There are two main facilities currently performing heavy-ion collisions to create droplets of quark-gluon plasma: the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Lab, and the Large Hadron Collider (LHC) at CERN. RHIC hosts the STAR and PHENIX experiments, while the LHC hosts the ALICE, ATLAS, CMS, and LHCb experiments. The first evidence for such a state was claimed by the CERN heavy-ion program at the SPS [15]. In 2010, RHIC officially claimed to have created deconfined QCD matter at $4 \cdot 10^{12}$ K, and in 2012, the LHC claimed to have created deconfined QCD matter at $5.5 \cdot 10^{12}$ K.

1.4.1 Evolution of the quark-gluon plasma

The sequence of events in a heavy-ion collision is shown in Fig. 1.5. Two incoming nuclei collide, and hard and soft scatterings of partons produce a large energy density. The system then approximately thermalizes – the quark-gluon plasma has formed. The system then cools via hydrodynamic expansion, and when the temperature reaches the confinement transition, partons hadronize and re-scatter in a hadron gas until they free stream. Each element of the evolution will now be discussed below.

Initial state

The incoming nuclei, consisting of Lorentz-boosted collections of nucleons described by a nuclear PDF, impact each other with a given impact parameter. At the earliest times of the collision, hard scatterings of partons take place ($t \sim \frac{\hbar}{E}$). As the collision continues, a large number of inelastic soft scatterings occur, as the large kinetic energy of the accelerated

nuclei is transformed into partonic energy density. The resultant energy density in the collision system allows the partons to become de-confined, in a pre-equilibrium precursor state to the QGP. The nuclei pass through each other with proper collision time $\tau \sim 0.01 \frac{\text{fm}}{c}$, with most of the baryon density propagating close to beam rapidity, and leaving a collision region of high energy density and low- μ_B , which evolves to an approximate equilibrium on a timescale $\tau \approx 0.5 \text{ fm}/c$. This timescale is smaller than the causal time that links one side of the system to the other, which is $\mathcal{O}(10 \text{ fm}/c)$. The mechanism by which the system thermalizes is unknown: whether the system approximately thermalizes in this pre-equilibrium phase, or if the system is “born” into equilibrium. Regardless, at this point, the QGP is said to have formed.

The system then expands and cools, and has been found to be very well described by viscous hydrodynamics. The initial conditions provided to hydrodynamic models are usually described in one of two popular models: a Glauber model, or a color glass condensate (CGC) model. There are a variety of specific implementations of the initial conditions for the evolution, for example MC Glauber or Optical Glauber for Glauber models [17], IP-Glasma [18] for CGC models, or others such as Trento [19].

Glauber models are relatively simple. Using a distribution of nucleons based on measured nuclear density profiles (typically a Woods-Saxon distribution, identical for protons and neutrons), and including event-by-event fluctuations, the Glauber model uses an eikonal approximation (small angle scattering) of the trajectories of independent nucleons and the nucleon-nucleon cross-section to compute the number and location of nucleon interactions. In this picture, nucleons pass through each other with some energy density deposited in




Figure 1.5: The stages of a heavy-ion collision [16].

the collision region, but continue onward (recall that the pp cross-section does not depend strongly on energy at these high energies, so the energy deposition does not have a strong effect on the subsequent scatterings). The Glauber model entirely neglects the fact that nucleons are bound (some quite strongly, such as np pairs), and that there is rich gluonic substructure in nucleons – yet it is still quite accurate. In addition to providing event-by-event initial conditions, the Glauber model is commonly used to determine the impact parameter of a collision, as will be described in the next section. CGC models, on the other hand, model the incoming nucleus as dominated by low- x gluons (due to time dilation) which can be described classically, due to the high gluon occupation number. The pre-equilibrium state is referred to as the Glasma, which describes the evolution of these high-occupancy gluons to thermal equilibrium.

Hydrodynamic expansion

The medium then expands according to relativistic viscous hydrodynamics. This is described by an energy-momentum tensor $T^{\mu\nu}$:

$$T^{\mu\nu} = \varepsilon u^\mu u^\nu - p (g^{\mu\nu} - u^\mu u^\nu) + \pi^{\mu\nu}$$

where ε is the energy density, p is the pressure, u^ν are the velocity fields, $g^{\mu\nu} = (1, -1, -1, -1)$, and $\pi^{\mu\nu}$ is the shear stress tensor, which contains η/s . Given initial conditions supplied by one of the initial state models above, and a lack of external sources, the equations of motion are given by conservation of the energy-momentum tensor:

$$\partial_\mu T^{\mu\nu} = 0.$$

These equations are typically computed numerically in either 2+1 or 3+1 dimensions, by packages such as MUSIC [20], iEBE-VISHNU [21], and CLVisc [22].

One alternative model to hydrodynamic expansion is that of transport models. These are microscopic models which describe the evolution of partons throughout the expansion, and unlike hydrodynamics is able to describe non-equilibrium dynamics. Common transport

models are AMPT [23] or BAMPS [24].

Hadronization

When the energy density becomes sufficiently small, and the temperature is sufficiently cool, the long-distance coupling causes the partons to hadronize. This hadronization process is different in vacuum and in medium, and is accordingly modeled differently in each system. Since hadronization is a non-perturbative process, all hadronization models are phenomenological.

In pp collisions, the two most common models are those of string fragmentation and cluster fragmentation. In string fragmentation models, hadronization is described by a color string connecting $q\bar{q}$ pairs with a linear confinement potential $V \approx kx$, where $k \approx 1$ GeV/fm. As the quarks separate and the potential energy of the string gets large enough, a new $q\bar{q}$ pair forms, and the string breaks into two color singlet hadrons. Similarly, a diquark can be treated as an antiquark, which then allows for the formation of baryons. String fragmentation is implemented in Pythia [25]. In cluster fragmentation models, groups of color-singlet quarks are clustered together, and these clusters decay to groups of hadrons. Cluster fragmentation is implemented in Herwig [26].

In heavy-ion collisions, however, the continuous fluid description of the medium must be translated into discrete hadrons, in a way that satisfies energy and momentum conservation. This is commonly done by Cooper-Frye freeze-out [27] or similar techniques, at a time referred to as “thermal freezeout”. These models describe hadronization by statistical hadronization methods, either by transforming localized thermal states of the medium directly into hadrons on a statistical basis, or by transforming the continuous medium into partons and then allowing nearby partons recombine into hadrons in so-called quark recombination or coalescence models.

The produced hadrons then re-scatter in a hadron gas, and contain excited states and resonances. To describe the evolution of these hadronic interactions, codes such as UrQMD [28] or SMASH [29] are commonly used. When the hadron species become fixed, this is referred to as “chemical freezeout”. And when their four-momenta become fixed, and they free stream, it is called “kinetic freezeout”. This occurs at a time of $t \approx 10$ fm/c. The

final-state particles then travel a distance $\approx 10^{15}$ fm, where they are detected and used to reconstruct various event-by-event or ensemble-averaged observables.

1.4.2 Experimental variables

Before describing specific heavy-ion measurements, it will be useful to define common experimental variables: kinematic variables, including the spatial and momentum coordinates, and centrality, which is the observable used to define the impact parameter.

First, let's define the kinematic variables. In a heavy-ion collision, two beams collide along a beamline, and produce many outgoing particles – we presently describe the coordinates used to discuss those outgoing particles. For cylindrical detectors, as is the case for ALICE, the spatial coordinates of the particles are described using cylindrical coordinates. In particular, we use θ to denote the polar angle from the beam axis, and ϕ to denote the azimuthal angle. The radius r is generally immaterial to describe the particle's kinematics, since the particles propagate in the radial direction.

The three-momentum of a final-state particle in a collision is decomposed into its transverse and longitudinal components relative to the beam-line: $\mathbf{p} = \mathbf{p}_T + \mathbf{p}_L$. The transverse momentum, $p_T \equiv |\mathbf{p}_T|$, serves as a proxy for the momentum transfer $\sqrt{Q^2}$. The longitudinal momentum, on the other hand, is determined not just by the interaction but by the initial longitudinal momentum carried by the parton inside the nucleon. While p_T is approximately conserved in a collision, p_L is not, since two colliding partons will carry a different initial state momentum fraction x of their parent nucleons. Instead of using the longitudinal momentum, particles are typically described by rapidity, $y = \frac{1}{2} \ln \frac{E+cp_L}{E-cp_L}$. Rapidities add for boost β along the longitudinal direction: $y' = y + \text{arctanh}\beta$, which gives the advantage that the shape of a distribution differential in y does not depend on the reference frame. Often, however, the rapidity is approximated in the massless limit by the pseudo-rapidity, $\eta = \frac{1}{2} \ln \frac{|\mathbf{p}|+p_L}{|\mathbf{p}|-p_L}$. Note that pseudo-rapidity diverges along the beam axis, while the rapidity does not, since it is cut off by the particle mass. The pseudo-rapidity has the advantage, however, that it can be written purely in terms of the polar angle: $\eta = -\ln \left[\tan \frac{\theta}{2} \right]$. Intuitively, then, we can think of η as depending on the relative amount of p_T compared to p_L . It is useful also to note that the distance $R = \sqrt{(\phi_1 - \phi_2)^2 + (\eta_1 - \eta_2)^2}$ is invariant

under Lorentz transformations along the beamline. Together, we will typically describe the complete particle three-momentum $\mathbf{p}$ by the coordinates p_T , η , and ϕ .

Aside from the kinematics of the outgoing particles, we need to describe experimentally the overlap of two colliding nuclei. We will consider only the case of spherical nuclei, as is the case for Au and Pb, the main ion species collided at RHIC and the LHC, in which case a single impact parameter b fully describes their overlap. Experimentally, we have no direct access to b . However, “central” collisions with small b produce larger particle multiplicity than “peripheral” collisions with larger b . The nuclear overlap is therefore described by an observable called “centrality”, where 0-10% centrality denotes the 10% smallest b collisions and 90-100% centrality denotes the 10% largest b collisions. Each collision event is classified by a particular centrality value based on the particle multiplicity produced at a certain η , as shown in Fig. 1.6. The particle multiplicity is then correlated with b using the Glauber model, introduced above, and the measured multiplicity is mapped to a given centrality value. For a given centrality, the Glauber model also produces the average number of binary nucleon-nucleon collisions, denoted N_{coll} , and the total number of colliding nucleons, denoted N_{part} . A 0-10% centrality Pb²⁰⁸Pb²⁰⁸ collision at $\sqrt{s_{\text{NN}}} = 5.02$ TeV has $\langle N_{\text{coll}} \rangle = 1583$, $\langle N_{\text{part}} \rangle = 359$, meaning that on average each colliding nucleon undergoes 4-5 scatterings. It is observed that the production of hard particles scales with N_{coll} , whereas the total particle production scales with N_{part} .




Figure 1.6: Number of events as a function of charged particle multiplicity in V0 detector, as measured by ALICE. Events in the highest multiplicity subdivision are assigned centrality values 0-5% [30].

The accuracy of N_{coll} scaling, which will be relevant for the results presented in Chapter 3, is verified by measuring for example the yield of direct photons in Pb–Pb collisions compared to pp collisions. Since photons don’t interact strongly, we expect that their yield in Pb–Pb collisions will be equal to a superposition of pp collisions (neglecting nPDF effects). This is quantified by the observable called the nuclear modification factor, R_{AA} :

$$R_{\text{AA}} = \frac{\frac{1}{\langle T_{\text{AA}} \rangle} \frac{1}{N_{\text{event}}} \frac{d^2 N}{dp_{\text{T}} d\eta} \Big|_{\text{AA}}}{\frac{d^2 \sigma}{dp_{\text{T}} d\eta} \Big|_{\text{pp}}}, \quad (1.1)$$

where $\langle T_{\text{AA}} \rangle = \frac{\langle N_{\text{coll}} \rangle}{\sigma_{\text{inel}}^{NN}}$ is computed in a Glauber model, and has value $\langle T_{\text{AA}} \rangle = 23.4 \pm 0.78$ (sys) mb^{-1} for 0-10% Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV. The nuclear modification factor describes the modification of heavy-ion yields compared to a superposition of pp yields. It has been demonstrated that particles expected to interact negligibly with the medium, such as direct photons, exhibit R_{AA} consistent with 1, providing evidence that N_{coll} scaling is predicted correctly by Glauber.

1.4.3 Signatures of the quark-gluon plasma

Evidence that deconfined QCD matter is produced in heavy-ion collisions arises in several complementary observables. In this section, I highlight three of the most compelling signatures of the quark-gluon plasma: elliptic flow, J/ψ suppression, and jet quenching.

Anisotropic flow

In a mid-central heavy-ion collision, the overlap between the incident nuclei produces an oblong volume of QCD matter, as shown in Fig. 1.7. The medium will then expand in the direction of the steepest pressure gradients. If the medium is liquid-like, it will expand not just radially, but also in the plane transverse to the major axis of the volume, containing the impact parameter and beam axis (bisecting the oblong volume), where the isobar contours are spaced more closely together. The particles near this reaction plane of expansion will acquire an enhancement in p_{T} , due to this rapid expansion. That is, the initial spatial anisotropy of the collision geometry will be translated into a final-state

momentum anisotropy. In particular, this would produce an enhancement of back-to-back particle detection in ϕ ; since the particle p_T spectrum is steeply falling, an enhancement in momentum means that there would be an enhancement in particle yields at any fixed p_T . This azimuthal anisotropy is called elliptic flow. Note that the anisotropy is expected to be largest for more peripheral collisions, and is expected to be present at low- p_T , where particles thermalize and comprise the bulk of the medium.

To measure this effect experimentally, we construct the (Lorentz-invariant) azimuthal distribution of particles relative to the reaction plane. This distribution is expanded in a Fourier expansion in ϕ :

$$E \frac{d^3 N}{dp^3} = \frac{1}{2\pi} \frac{d^2 N}{p_T dp_T dy} \left[1 + 2 \sum_{n=1}^{\infty} v_n \cos(n[\phi - \Psi_n]) \right],$$

where $\phi - \Psi_n$ denotes the azimuthal angle relative to the n -th order reaction plane. The second Fourier component, v_2 , describes the magnitude of the elliptic flow. Higher-order Fourier coefficients denote the contribution of other anisotropies, such as “triangular flow”, which can arise from geometrical fluctuations in the nucleon-nucleon scatterings of the collision. There is a large experimental effort dedicated to measuring these Fourier coefficients, and many techniques have been developed to distinguish these correlations from other correlations (e.g. short-range jet fragmentation correlations), such as using N -particle correlations instead of 2-particle correlations to measure the v_n .

By modeling the initial state and hydrodynamic evolution of the medium, one can




Figure 1.7: Two incident nuclei partially overlap, creating an “almond”-shaped collision region, which then expands and generates elliptic flow [31, 32].

calculate the magnitude of v_n in a given system. In such a fit, η/s is a free parameter in the hydrodynamic evolution. In fact, the v_n are rather sensitive to the value of η/s , since a large η/s will dissipate the initial anisotropy. Accordingly, η/s can be extracted from measurements of v_n , in a model-dependent way (possibly also evolving as a function of T). The result of these experimental measurements and theoretical modeling have demonstrated that η/s is exceedingly small: $\eta/s \approx 0.1 - 0.2 \frac{\hbar}{k_B}$ [34], while the conjectured lower quantum limit is $\eta/s = \frac{1}{4\pi} \frac{\hbar}{k_B}$. Figure 1.8 shows an example of this procedure; to get a good fit at low- p_T , $\eta/s \approx 0.2 \frac{\hbar}{k_B}$ is required at the LHC; slightly smaller values $\eta/s \approx 0.1$ are more appropriate at RHIC [33]. As described in Section 1.3.1, the small extracted value of η/s implies the fluid is strongly coupled. This is inconsistent with a purely hadronic description of matter.

Moreover, since the medium expands locally at a common velocity, in this picture we expect the magnitude of v_2 to be ordered according to particle mass. At the LHC, this is observed experimentally for $p_T < 3$ GeV/c [35]. Further, if hadronization occurs via coalescence of nearby partons, we may expect that baryons have larger v_2 than mesons – “constituent quark scaling”. This has been observed experimentally to approximately hold, and provides further evidence for deconfined matter; the role of coalescence models relative




Figure 1.8: Measurement of the root-mean-square of v_n from ALICE as a function of centrality, along with model predictions for $\eta/s = 0.2 \frac{\hbar}{k_B}$ [33]. Note that the magnitude of v_2 is highest for mid-central collisions, where the collision geometry is expected to generate the largest elliptic flow.

to “traditional” hadronization models is an ongoing research topic.

Quarkonia suppression

In the initial collision, a high- Q^2 scattering can produce $c\bar{c}$ or $b\bar{b}$. These can form “open” heavy flavor hadrons such as D mesons, or they can form quarkonium states such as J/ψ or Υ . In the high- T medium, quarkonia states are predicted to dissociate. This has been observed experimentally, as shown in Fig. 1.9. Moreover, the more tightly bound quarkonia states may be expected to “sequentially melt” as T increases. However, at higher $\sqrt{s_{NN}}$, where higher T is achieved, it has been observed that J/ψ production increases relative to lower T , due to enhanced recombination of independent $c\bar{c}$. In fact, this provides further evidence of deconfined QCD matter – suppression of J/ψ due to melting, and the appearance at higher- T of J/ψ formed not at the hard-scattering, but by statistical recombination near the phase boundary of the thermal medium.

Jet quenching

The hot, dense quark-gluon plasma suppresses the yields of high- p_T particles, due to the interactions with the dense medium. The use of jets to understand the quark-gluon plasma




Figure 1.9: Suppression of J/ψ as measured by ALICE, and several theoretical models [36].

will be discussed in detail in the next chapter.

Other signatures

There are a variety of other confirmed and speculated signatures of the quark-gluon plasma [34, 37]. The measured charged-particle pseudo-rapidity density for central collisions suggests the initial energy density at the LHC reaches $\sim 15 \text{ GeV/fm}^3$, much higher than hadronic matter. The yields of identified hadrons generally follow a thermal distribution, with a temperature compatible with the predicted deconfined phase. Thermal photons also exhibit a thermal distribution corresponding to a large medium T . The content of strange particles is enhanced, given that they can be produced thermally. There is a large ongoing search for signs of chiral symmetry restoration, particularly by measuring and modeling the spectral functions of the ρ meson, and also searches for chiral anomalies, such as the chiral magnetic effect.

One prominent area of research is the study of pp and p-Pb events with high-multiplicity. There have been two particularly unexpected observations in such high-multiplicity systems: a “nearside ridge” in the two-particle correlation similar to that in elliptic flow [38], and an enhancement in the strangeness content [39]. The explanation for these phenomena is not yet clear, and it raises the question of what the minimal system is that can form a quark-gluon plasma, and if we misunderstand something about certain signatures being unique to deconfined QCD matter. It is an important ongoing effort to attempt to understand pp, p-A, and A-A in a single framework.

Nevertheless, the establishment of the existence of a state of deconfined QCD matter created in heavy-ion collisions is by now clear. As the field attempts to move from the discovery phase to the precision phase, interplay between theory and experiment is crucial in order that the observables investigated are measurable, calculable, and useful.

1.5 Jets in heavy-ion collisions

A high- Q^2 scattering can produce a parton with large p_T . As it propagates, this parton will fragment into a shower of daughter partons, mostly via collinear gluon radiation and

the formation of quark-antiquark pairs. When sufficient splittings have occurred such that the shower partons reach low enough energy, the QCD coupling constant α_S becomes large and the partons hadronize. This collimated collection of final state particles is referred to as a jet.⁸ Typically, a jet contains most of its p_T within a “core” of radius $R \approx 0.2$ (where $R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$), and all jet particles are typically within $R < 1$. By momentum conservation, jets are produced approximately back-to-back in ϕ , and are known as “di-jets”.⁹ However, jets are typically not produced in LO $2 \rightarrow 2$ on-shell scatterings, but rather in a highly virtual state, and this virtuality decays as the jet fragments. The initial legs of the shower can be either quarks or gluons; gluons, having a larger color factor by $\frac{9}{4}$, tend to radiate more and therefore have broader and softer showers. In general, large-angle splittings are suppressed since this requires large- Q^2 , and the strong coupling is smaller at large Q^2 . This fact defines the collimated nature of jets.

At the LHC, jet production is dominated by gluon fusion, due to the increasingly gluonic nature of the proton at large $\sqrt{s}$. Moreover, at the higher energies of the LHC, the cross-sections for hard probes are much larger than at RHIC. That is to say, the shape of the p_T -differential jet cross-section as a function of p_T , which can be roughly modeled by an exponential or power law distribution, becomes flatter at higher $\sqrt{s}$. Because of this, the advent of the LHC has marshaled a revolution in the physics of fully reconstructed jets.

In heavy-ion collisions, jets can serve as a probe of the hot QCD medium, since the hard scattering occurs early in the collision,¹⁰ and the jet production spectrum can be computed in perturbative QCD. A jet will then traverse a significant pathlength of the medium, and the effect that the medium has on the jet can be deduced by comparing jet properties in heavy-ion collisions to those in pp collisions. Moreover, the fact that jets carry a large p_T suggests they will be sensitive to a wide range of momentum exchanges with the medium,

8. Jets can also be produced from decays of heavy particles, such as the hadronic decay of the W boson, but these cases will not be relevant here.

9. At LO and with no medium energy loss, di-jets have equal p_T and $\Delta\phi \approx \pi$. Higher-order contributions can involve a third jet, if a high-energy gluon is radiated early in the shower, which breaks $\Delta\phi \approx \pi$. Also, the initial partons can have a nonzero transverse momentum component, denoted k_T , which results in not quite back-to-back in $\Delta\phi$. Note that jets are not opposite in θ , since partons have different x_1, x_2 ; the jet is boosted by $y = \ln \frac{x_1}{x_2}$, known as the “ η -swing”.

10. A 50 GeV jet is produced in $t \sim \frac{\hbar}{E} \approx 0.003$ fm/c, while the QGP thermalizes in $t \approx 0.5$ fm/c.

and thereby can provide insight to the medium at a wide range of resolution scales. Before it was tenable to reliably reconstruct sufficiently many jets in heavy-ion collisions, high- p_T hadrons were used as proxies for jets. However, high- p_T hadrons are only one component of jets, and they arise from a highly biased sample of jets that tend to undergo little medium modification [40]. Fully reconstructed jets have the advantage of being unbiased objects representing high- p_T energy flow through the medium, and moreover jets offer a rich set of unique observables such as in their substructure.

At a qualitative level, it is expected that jets will be modified by the medium, since the QGP has a larger energy density and more degrees of freedom than cold nuclear matter. A wide variety of jet observables are studied to look for significant modifications: inclusive jet observables, which study the overall suppression of jets; jet correlation observables, which associate a jet with one or more other objects; and jet substructure observables, which examine modification to the internal shape or substructure of jets. Significant modification in all of these categories has been observed experimentally in heavy-ion collisions, and serves as one of the strongest pieces of evidence for the existence of deconfined QCD matter. These modifications taken together are generally referred to as “jet quenching.” The current goal is to use this rich set of observables to learn about the structure of this deconfined state by understanding how jets interact with it. A detailed summary of what has been learned from jet measurements is summarized later in this section, but first we describe the basic attributes of reconstructing jets experimentally.

1.5.1 Measuring jets in heavy-ion collisions

Jet reconstruction algorithms

There is not an objective way to define a jet, either in theory or in experiment. Accordingly, then, the precise definition of a jet is whatever we define it to be – the output of a jet reconstruction algorithm. A jet reconstruction algorithm consists of a clustering scheme, used to group together relevant final-state particles, and a recombination scheme, which constructs the jet momentum from these particles. In order to be well-defined theoretically, a jet reconstruction algorithm should be infrared and collinear (IRC) safe, meaning that

the output of the algorithm should not change if a particle with $p_T \rightarrow 0$ is added, or if a particle is split into two collinear particles. A successful algorithm should also be able to apply equally well to experimentally measured particles at the hadron level and in theoretical computations at the parton level.

The simplest jet reconstruction algorithm is a cone algorithm, in which one identifies the jet direction according to its high- p_T particles, and then clusters all particles within a given radius R , e.g. $R_{cone} \approx 0.7$. However, such a clustering algorithm is not IRC safe, and is dependent on the seed. More sophisticated algorithms have been developed to avoid these issues, and the most common jet reconstruction algorithms currently in use are those known as sequential recombination algorithms. Sequential recombination algorithms cluster particles according to two metrics:

$$d_{ij} = \min \left(p_{T_i}^p, p_{T_j}^p \right) \frac{\Delta_{ij}^2}{R^2},$$

$$d_i = p_{T,i}^p,$$

where $p_{T,i}$ is transverse momentum of the i^{th} constituent, $\Delta_{ij} = \Delta\eta^2 + \Delta\phi^2$ is separation between two constituents, R is a parameter of the algorithm, referred to as the jet-resolution parameter, roughly corresponding to the radius of the jet cone in natural angular units, and p is a fixed parameter defining the model. The algorithm is as follows: The d_{ij} are computed for all pairs of particles, and the d_i for all single particles. The global minimum of $\{d_{ij}, d_i\}$ is then computed. If the global minimum measure is d_{ij} , then we cluster the i, j nearest-neighbor particles and combine their momenta. If the lowest measure is d_i , then this particle is deemed a jet and removed from the collection. The d_{ij}, d_i are then re-computed using all available particles, until all jets are found. The jet momentum is then determined by a particular recombination scheme, such as the E -scheme, where the four-momenta of the constituents are added, or the p_T -scheme, where the constituent four-momenta are normalized to be massless. In essence, particles are clustered according to their momentum (depending on the exponent p) and their angular proximity (with collinear particles likely to be clustered together). One can see that the algorithm is infrared safe since a particle

with infinitesimal momentum will immediately be clustered with its nearest particle, only changing the momentum infinitesimally. Similarly, the algorithm is collinear safe since two collinear particles will be immediately combined into one. There are three common selections of p . For $p = 2$, known as the k_T model, clustering begins from low p_T particles, and the jets can have irregular shapes in η, ϕ . For $p = -2$, known as the anti- k_T model, clustering begins from high p_T particles, and jets tend to be somewhat circular in η, ϕ . For $p = 0$, known as the Cambridge/Aachen model, clustering depends only on angular distance. These algorithms have been efficiently implemented in the FastJet package [41, 42].

Heavy-ion background

In a heavy-ion environment, there is a significant complication to reconstruct jets: there is a large background of bulk hadrons produced from the thermally expanding QGP. For example, in a central Pb–Pb collision at $\sqrt{s_{NN}} = 2.76$ TeV, ALICE finds an average jet background of $p_T \approx 24$ GeV/ c per $R = 0.2$ jet. Moreover, for larger R jets – which capture the full jet information – the background dramatically increases, proportional to the jet area; for a $R = 0.4$ jet, the average background is $p_T \approx 100$ GeV/ c . In addition to the large average background present in heavy-ion collisions, there are large event-to-event and intra-event fluctuations in this background. This is particularly challenging, since these background fluctuations distort the measured jet p_T , and at low jet- p_T it becomes impossible to distinguish a “real” jet from a background jet on a jet-by-jet basis. The presence of the heavy-ion background also complicates the study of jet substructure, since on a particle-by-particle basis it is not clear whether the particle arose from the background or from the jet fragmentation.

Several strategies have been developed to cope with the heavy-ion background, depending on the jet observable in question. First, jets are measured with the anti- k_T algorithm so as to start clustering from the high- p_T particles. Jets are reconstructed at relatively small R , typically $R \approx 0.2 - 0.4$, and the average background is subtracted jet-by-jet.¹¹ Sometimes,

11. Different experiments use different background subtraction schemes; some do it before jet-finding, some after; some find an average background over the entire event, while some also subtract flow modulations in the heavy-ion background.

biases are imposed on the jets or their constituents, such as requiring a high- p_T constituent to be present in the jet, or by restricting the minimum p_T of the constituents. These have the drawback of biasing the selection of jets relative to an inclusive sample, which can alter the characteristics of their interaction with the medium. For jet substructure observables, methods have been developed to subtract background on a particle-by-particle basis or a statistical basis. In addition to these jet-by-jet strategies, statistical corrections are used to de-convolute the smearing of the jet p_T by the background fluctuations. This is done using unfolding/deconvolution algorithms, and will be explained in Chapter 3.

The most challenging jet reconstruction in heavy-ion collisions is therefore at low jet p_T . At RHIC, there is a substantially smaller signal-to-background ratio than at the LHC. This has driven innovations to be developed in heavy-ion jet measurements at RHIC, such as mixed-event background subtraction, but it has also limited the ability of RHIC experiments to fully reconstruct jets without imposing strong biases. Full jet reconstruction has therefore been largely dominated by the LHC, with ALICE focusing on the $p_T \lesssim 100$ GeV/c range, and ATLAS and CMS measuring jets at very large p_T up to nearly 1 TeV.

1.5.2 Theory of jet modification in heavy-ion collisions

There are several theoretical approaches to model jet energy loss in heavy-ion collisions. Most commonly, jet energy loss is modeled in perturbative QCD as soft gluon radiation emitted out of the jet cone, induced as the jet travels through the dense color-charged medium. There are several pQCD-based formalisms describing jet energy loss in the QGP [43]. These typically involve the following parameters: the Debye mass, m_D , is the inverse screening length of the medium; the opacity $N \equiv L/\lambda$, where L is the pathlength and λ is the mean free path; the momentum diffusion $\hat{q} \equiv m_D^2/\lambda$, which describes the transverse momentum squared transferred from the jet per unit pathlength.¹²

Jet energy loss is generally believed to be dominated by radiative energy loss as opposed to collisional energy loss. However, models suggest that collisions of the jet partons with

12. The JET Collaboration deduced limits on $\hat{q}$ based on high- p_T single hadron suppression at RHIC and the LHC, which gives the rate of energy loss as $\hat{q} = 5 - 15 \frac{(\text{GeV}/c)^2}{\text{fm}}$. This is inconsistent with cold nuclear matter energy loss, for which $\hat{q}$ is several times smaller [44].

the medium induces gluon radiation to be emitted. In pp, the initial parton in the jet shower is produced off-shell, i.e. $E^2 \neq p^2 + m^2$. The virtuality (commonly denoted $Q^2 \equiv |E^2 - p^2 - m^2|$) that the jet exhibits then decays over the course of the shower. In Pb–Pb, it is predicted that elastic collisions with the medium slow the decay of the parton’s virtuality.

Additionally, quantum interference effects are important to accurately describe the jet energy loss. When a parton splitting in the jet shower occurs, the distance between the daughters must be sufficiently large in order for the medium to independently resolve the two partons. For the time they are too close together to be resolved by the medium, they lose energy coherently as a single parton. This is described by the Landau-Pomeranchuk-Migdal (LPM) effect.

There are four main pQCD formalisms [43]. The BDMPS approach uses a path integral approach to describe $\hat{q}$ by soft gluon radiations. The GLV approach uses the “opacity expansion”, in which the gluon radiation distribution is expanded in the opacity. The Higher Twist approach describes multiple parton scattering by computing power corrections to the “leading-twist” cross-section. The AMY approach uses finite temperature field theory to compute the modified jet fragmentation. Each of these approaches essentially uses different assumptions in pQCD to compute the gluon radiation rate as the jet shower loses virtuality.

There are also approaches to jet energy loss as strongly-coupled interactions with the medium, based on the conjectured AdS/CFT correspondence. This uses dramatically different physics than pQCD-based models, with the jet losing energy by drag force through a continuous liquid, rather than weakly-coupled induced radiations. Unfortunately, the only calculational tool in the strong-coupling approach is using holographic methods for an infinitely-coupled supersymmetric $SU(4)$ Yang-Mills theory, which is different than QCD.

Each of these formalisms has been fairly successful, although it has been challenging to compare them rigorously, since they describe different physical processes, use different approximations (e.g. treating α_s as constant), use different approaches to model the evolution of the medium, and use different input medium parameters. For example, the parton mean free path in the medium remains fairly unconstrained. Precisely determining $\hat{q}$ vs. T is a major outstanding task of heavy-ion jet physics. To properly compare a given jet energy loss model to experimental measurements, one must also model the entirety of the

heavy-ion evolution: initial state, hydrodynamic expansion, and hadronization. Ultimately, however, we aim to understand not jet energy loss, but deconfined QCD matter itself.

1.5.3 Phenomenology of jet modification in heavy-ion collisions

While a precise theoretical description of jet energy loss and the structure of the quark-gluon plasma remains elusive, measurements of jet and high- p_T observables have painted a phenomenological picture of the characteristics of jet modification in heavy-ion collisions. Below, I outline several of the major insights that have been gained.

Jet yields are suppressed

It is predicted that when a jet passes through the quark-gluon plasma, a certain fraction of the jet's energy is radiated outside of the jet cone – the jet loses energy. Since the p_T -differential jet cross-section is a steeply falling function of p_T , this shift in jet energy implies that the yields of jets in a given p_T bin will be suppressed in heavy-ion collisions compared to pp collisions. The observable describing this is the inclusive jet R_{AA} , introduced previously in Eq. 1.1. Inclusive jet R_{AA} is a highly averaged observable: it includes quark and gluon jets, light and heavy flavors, all configurations of medium evolution – typically for a fixed centrality. Figure 1.10 (left) shows the measured R_{AA} values from ALICE, ATLAS, and CMS in Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. The R_{AA} exhibits a strong suppression, corresponding to $\approx 10 - 20\%$ energy loss, with the amount of energy loss decreasing as p_T increases roughly as $\sqrt{p_T}$ (presumably due to smaller α_S of the jet-medium interaction) [45].

It was speculated that some or all of this energy loss may be due to effects of cold nuclear matter, either in the initial state (nPDF) or final state (energy loss). However, Fig. 1.10 (right) shows that R_{pPb} , the ratio of the p-Pb to pp jet spectra (analogous to R_{AA}), exhibits negligible suppression.¹³ This measurement and others confirm that the observed jet R_{AA} is indeed due to energy loss in deconfined QCD matter.

13. Certain measurements of R_{pPb} exhibit a centrality-dependent R_{pPb} suggesting suppression; however, this is generally believed to be due to biases in centrality determination, not due to jet quenching.




Figure 1.10: Left: Jet R_{AA} for $R = 0.2$ in central Pb-Pb collisions measured by ALICE, ATLAS, and CMS at $\sqrt{s_{NN}} = 2.76$ TeV. Right: Ratio of the observable $\Delta_{recoil}(p_{T,jet}^{ch})$ in 0-20% event activity to 50-100% event activity p-Pb collisions measured by ALICE. This ratio measures the “centrality”-dependence of the magnitude of jet energy loss in p-Pb collisions, and is found to be consistent with no energy loss within a bound of $\Delta p_T = 0.4$ GeV/c per jet [46].

Soft energy is distributed to large angles

Given that jets lose energy due to out-of-cone fragmentation, one immediately asks: Where does the energy go? A priori, it is not clear whether the energy can be found outside of the jet cone, or if it is entirely thermalized in the medium. Measurements demonstrate that much of the lost jet energy is re-distributed outside of the jet cone, at large angles. This supports the interpretation that the lost jet energy is only “partially thermalized”, retaining a correlation to the jet.

This is best illustrated by a measurement from CMS in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV [47]. In this measurement, events containing an approximately back-to-back di-jet are examined. The di-jet is characterized by the p_T asymmetry between the leading and subleading di-jet: $A_J = \frac{p_{T,1} - p_{T,2}}{p_{T,1} + p_{T,2}}$, where large- A_J signifies a highly imbalanced di-jet. Then, the p_T of charged particle tracks in the event are projected onto the di-jet axis, with the projection denoted $p_T^{\parallel}$. The sum of $p_T^{\parallel}$ over tracks within a cone $R = 0.8$ is then performed, around both the leading jet and the subleading jet – as well as the sum outside of each $R = 0.8$ cone. The average $\langle p_T^{\parallel} \rangle$ is then plotted in bins of track p_T . This is shown in Fig.




Figure 1.11: CMS momentum balance plot in Pb–Pb, with the left panel corresponding to in-cone particles ($R < 0.8$), and the right panel corresponding to out-of-cone particles ($R > 0.8$). The circular data points display the net projected momentum $\langle p_T^\parallel \rangle$ summed over all p_T bins (at a given A_J) within the specified spatial region (in-cone or out-of-cone). In-cone, there is a large excess of high- p_T particles in the leading-jet direction, while out-of-cone, there is a large excess of low- p_T particles in the subleading jet direction. This is interpreted as the quenched di-jet energy being redirected to soft particles at large angles [47].

1.11. The result is that highly-imbalanced di-jets, i.e. those with at least one jet undergoing significant jet energy loss, have a relative excess of high- p_T particles inside the leading jet cone (compared to inside the subleading jet cone), and an excess of low- p_T particles outside of the subleading jet cone (compared to outside of the leading jet cone). That is to say, the subleading jet, which has lost energy, contains an excess of soft particles outside of its jet cone. And the magnitude of this effect was shown to be larger in Pb–Pb compared to pp (not shown here). Jet energy loss therefore results in soft energy being re-distributed to large angles.

This picture is corroborated by measurements of the jet “fragmentation function” $D(z) \equiv \frac{1}{N_{\text{jet}}} \frac{dn_{\text{ch}}}{dz}$, where z is the fraction of jet p_T carried by a jet constituent.¹⁴ Figure 1.12 shows a measurement of the modification to $D(z)$ in Pb–Pb compared to pp: $R_{Dz} \equiv \frac{D(z)_{\text{PbPb}}}{D(z)_{\text{pp}}}$ [48].

14. Note that this is not the true jet fragmentation function, but rather a measured quantity – it may contain not only jet fragments but also fragments from the heavy-ion background, particularly those that have recoiled from jet-medium scatterings.




Figure 1.12: Ratio of jet fragmentation observables $R_{Dz} \equiv \frac{D(z)_{PbPb}}{D(z)_{pp}}$ measured by ATLAS [48].

This demonstrates an excess of soft particles in Pb–Pb jets relative to pp jets. The excess at high- z will be discussed in the following subsection.

While the observation of the soft wide-angle excess is clear, the exact cause is not known. It is debated whether the large-angle soft excess is due to wide-angle radiation emitted by the jet, or whether it is due to medium particles recoiling from jet-medium scatterings.¹⁵

The fragmentation pattern of a jet impacts modification

It has also become clear that the initial fragmentation pattern of a jet is highly correlated with its amount of quenching. In particular, the “vacuum-like” properties of the initial jet are influential in determining how the jet will interact with the medium. There are two particular examples of this: (i) jets with wide-angle hard splittings have been observed to lose more energy than jets with collinear hard splittings, and (ii) gluon-like jets lose more energy than quark-like jets.

¹⁵. In addition to broadening of the jet energy, one may wonder if the angular coplanarity of di-jets is modified in heavy-ion collisions. In particular, it could be that a jet propagating through the medium undergoes a hard scattering with a medium scattering center (a quasi-particle) and the jet gets deflected in ϕ , resulting in the di-jet having a $\Delta\phi$ imbalance. Measurements at RHIC and the LHC investigated this effect, and with the current experimental precision, show no evidence for large-angle jet scatterings in the medium [49].

The hard splitting of a jet can be found using the Soft Drop technique [52, 53]. To do this, we find a jet using the anti- k_T algorithm, and then re-cluster the jet constituents according to the Cambridge-Aachen algorithm in order to produce an angularly ordered tree, similar to a parton shower. We then unwind the last clustering step, and check the Soft Drop condition $z > z_{cut} \left(\frac{\Delta R}{R_0} \right)^\beta$, where z_{cut} and β are parameters of the Soft Drop algorithm, typically taken as $z_{cut} = 0.1, \beta = 0$. If the condition is not passed, we discard the soft sub-jet and repeat another unwinding step. If the condition is passed, we deem the two subjets as a hard splitting of the jet, and characterize them by the shared momentum fraction $z_g = \frac{\min(p_{T,1}, p_{T,2})}{p_{T,1} + p_{T,2}}$, where $z_g \rightarrow 0.5$ denotes a symmetric splitting, and $z_g \rightarrow 0$ denotes an asymmetric splitting. Measurements by ALICE, shown in Fig. 1.13, plot the z_g distribution in Pb–Pb collisions in two categories of ΔR separation between the hard splitting sub-jets: $\Delta R < 0.1$, which are collinear splittings, and $\Delta R > 0.2$, which are large-angle splittings [50, 51]. The measurements show that collinear splittings exhibit only a small difference to pp, whereas the wide-angle splittings exhibit a strong suppression for symmetric splittings. This suggests that jets with a wide-angle hard splitting tend to lose energy in Pb–Pb collisions. This is interpreted as due to the fact that jets with a wide-angle hard splitting can be more readily resolved by the medium as two independent




Figure 1.13: Left: Distribution of z_g for collinear splittings $\Delta R < 0.1$ in Pb–Pb collisions measured by ALICE, and the ratio to Pythia [50]. Right: Distribution of z_g for wide-angle splittings $\Delta R > 0.2$ in Pb–Pb collisions measured by ALICE, and the ratio to Pythia [51].

partons, whereas jets with collinear hard splittings lose energy coherently as a single parton. It should be noted that while we know jets traverse the QGP, it is not obvious to what extent the fragmentation shower of the jet takes place inside the medium.

In addition to “groomed” jet substructure measurements, constituent-based jet shape measurements have also been performed. For example, ALICE has measured the radial moment, $g \equiv \sum_{i \in \text{jet}} \frac{p_{T,i}}{p_{T,\text{jet}}} \Delta R_{\text{jet},i}$, which is a measure of the jet’s radial momentum profile, and the momentum dispersion, $p_T D \equiv \frac{\sqrt{\sum_{i \in \text{jet}} p_{T,i}^2}}{\sum_{i \in \text{jet}} p_{T,i}}$, which is a measure of the dispersion of constituent momentum inside a jet [54]. A small radial moment corresponds to a more collimated jet, i.e. a more quark-like jet. A small momentum dispersion corresponds to a jet having softer fragmentation. Figure 1.14 shows the measured distributions of the radial moment and momentum dispersion. In Pb–Pb collisions, the radial moment exhibits a shift towards smaller g , and the momentum dispersion exhibits a shift towards larger $p_T D$. That




Figure 1.14: Left: Distribution of jet radial moment in Pb–Pb collisions measured by ALICE, and the ratio to Pythia. Right: Distribution of jet momentum dispersion in Pb–Pb collisions measured by ALICE, and the ratio to Pythia [54].

is, jets in Pb–Pb collisions are more collimated and harder fragmenting than jets in pp. This suggests that gluon-like jets, which have broader and softer fragmentation, lose more energy than quark-like jets – and so for a given p_T , the fraction of narrow, quark-like jets is enhanced.

Medium recoil is important to understand

Lastly, in order to accurately describe jet modification in heavy-ion collisions, it is important to understand the role of the jet’s effect on the medium. When a jet propagates through the medium, collisions of the jet shower partons with medium partons transfers momentum to medium particles in the direction of the jet, creating a correlated background. This is known as medium recoil. Medium recoil is believed to affect certain observables more than others, for example it is expected to have a larger impact on $R = 0.4$ jets than $R = 0.2$ jets. As an illustration, Fig. 1.15 shows the jet mass measured by ALICE in Pb–Pb collisions, along with two predictions by the jet energy loss Monte Carlo JEWEL [55]. The JEWEL prediction is seen to dramatically differ depending on whether medium recoil is included in the computation. A precise description of medium recoil is dependent on the full description of jet energy loss in the medium, and is therefore not yet understood.




Figure 1.15: Distribution of jet mass in Pb–Pb collisions measured by ALICE, and the comparison to several theoretical models [55]. The JEWEL prediction “recoil on”, which includes medium recoil particles in the jet mass computation, substantially differs from the JEWEL prediction “recoil off”, which neglects medium recoil particles.

Chapter 2

The ALICE Detector

ALICE is a large multi-purpose detector designed to study ultra-relativistic heavy-ion collisions at the LHC [56, 57]. There is a wide range of observables of interest in heavy-ion collisions, from the elliptic flow of identified hadrons to jet quenching to the suppression of quarkonia states, and more. In order to achieve its physics goals, the detector system must have the ability to:

- Precisely determine particle p_T over a wide range, $150 \text{ MeV}/c < p_T < 100 \text{ GeV}/c$
- Perform accurate particle identification, particularly at low- p_T
- Maintain performance in a high-multiplicity environment of several thousand particles per unit of rapidity

Accordingly, ALICE features a central barrel tracking detector at mid-rapidity, as well as a forward muon spectrometer and a variety of additional detectors, housed inside a $B = 0.5 \text{ T}$ solenoid magnet parallel to the beamline. The main tracking system consists of an inner tracking system (ITS) of 6 silicon layers, which allow for precise primary and secondary vertex determination, followed (radially outward) by a gas time projection chamber (TPC), which is the hallmark of ALICE. The charged particle tracking system provides high-precision p_T resolution and particle identification (PID) at low- p_T , and the TPC is able to maintain this performance in a high-multiplicity environment. The low- p_T precision and PID capabilities of ALICE in a high-multiplicity environment are unrivaled at

the LHC. Beyond the TPC sits the transition radiation detector (TRD) and time-of-flight (TOF) detector, which are used for PID and to aid in tracking. For part of the azimuth, electromagnetic calorimeters (EMCal, DCal, PHOS) measure direct photons, electrons, and π^0 . At forward rapidity, a muon spectrometer consisting of a set of absorbers and tracking chambers allows for the measurement of low- p_T J/ψ . Several scintillators are used for multiplicity determination and triggering. Nearly all of the detectors are in the 0.5 T solenoidal magnetic field to determine particle momentum and charge. The most relevant detectors for this thesis are the tracking system, which will be described in detail in Section 2.2, and the electromagnetic calorimeter, which will be described in detail in Section 2.3.

In total, ALICE consists of the following sub-detectors, which can be seen in Fig. 2.1:

1. Inner tracking system, a 6-layer silicon tracking system described in Section 2.2.1.
 - (a) Silicon pixel detector (SPD)
 - (b) Silicon drift detector (SDD)
 - (c) Silicon strip detector (SSD)




Figure 2.1: The ALICE detector, with all components labeled. The right-hand side is the “C”-side, and the left-hand side if the “A”-side.

2. Time projection chamber (TPC), the world's largest gas TPC, described in Section 2.2.2.
3. Time-of-flight (TOF), a multi-gap resistive plate chamber (MRPC) consisting of a stack of parallel plates and gas gaps, which measures particle time-of-flight to ≈ 100 ps, and is used for particle identification at intermediate p_T .
4. Transition radiation detector (TRD), which uses transition radiation (EM radiation from a particle passing through an inhomogeneous medium) to identify electrons in conjunction with the TPC.
5. Muon chamber and muon trigger (MCH, MTR), located on the “C-side” of the detector, which measure muons with $p_T^\mu > 4$ GeV/c at forward rapidity with a combination of several layers of absorbers (to stop hadrons and low- p_T muons) and tracking chambers, with goal to measure the suppression of low- p_T $J/\psi \rightarrow \mu^+\mu^-$ and other quarkonia.¹
6. Electromagnetic calorimeter (EMCAL, DCAL), a Pb-scintillator sampling calorimeter described in Section 2.3.
7. Photon spectrometer (PHOS) and charged particle veto (CPV), a highly-segmented PbWO_4 crystal calorimeter for precision photon measurements.
8. High-momentum particle identification detector (HMPID), a ring-imaging Cerenkov detector (the ring size depends on the particle velocity) that complements the p_T reach of PID in the TPC and TOF.
9. T0, a pair of Cherenkov counters surrounding the beam pipe at $-3.28 < \eta < -2.97$ and $4.61 < \eta < 4.92$ to measure the initial time of the collision event to ≈ 50 ps, and to provide vertex and centrality information.
10. V0, a scintillator array on either side of the interaction point, which is used to construct the minimum bias trigger, and provides centrality information.

1. Due to the constraint of having thick absorbers, one can only measure muons with $p_T^\mu > 4$ GeV/c [58]. That is, one can only measure J/ψ with relatively large p . Therefore, in order to measure low- p_T J/ψ , one must measure those J/ψ with a large p_z , i.e. at large η .

11. Zero degree calorimeter (ZDC), which consists of a tungsten calorimeter neutron detector (ZN) and brass calorimeter proton detector (ZP) located ≈ 116 m from interaction point, in order to measure the spectator nucleons of a collision for centrality determination.
12. Forward multiplicity detector (FMD), a silicon strip detector to measure charged-particle multiplicity at $-3.4 < \eta < -1.7$ and $1.7 < \eta < 5.0$.
13. Photon multiplicity detector (PMD), a preshower converter and gas counter to measure the multiplicity and position of photons at $2.3 < \eta < 3.7$.
14. Cosmic ray detector (ACORDE), a plastic scintillator located outside of the L3 magnet, to study high-energy cosmic rays.
15. Diffractive detector (AD), two sets of plastic scintillators to improve the measurement of pp diffractive scatterings.

The ALICE triggering system contains three hardware-level trigger levels (L0: $1.2\mu\text{s}$, L1: $6.5\mu\text{s}$, L2: $88\mu\text{s}$) and one software-level High-Level Trigger (HLT) responsible for online processing and reconstruction, and which compresses event information by a factor ≈ 10 . Heavy-ion physics observables require a large sample of minimum bias events, which are triggered by the V0 scintillators, but also rare probes such as high- p_T photons or forward muons. Accordingly, the electromagnetic calorimeters and muon spectrometer have the capability to trigger these rare events. The readout time per event is ≈ 1 ms, so ALICE accordingly records data at rates up to ≈ 1 kHz, depending on data-taking conditions.

2.1 The Large Hadron Collider

The Large Hadron Collider is the most powerful particle accelerator in history, delivering pp beams up to center-of-mass energies of $\sqrt{s} = 13$ TeV and Pb–Pb beams up to center-of-mass energy per-nucleon of $\sqrt{s_{\text{NN}}} = 5.02$ TeV [59]. The facility is home to four major experiments: ATLAS, CMS, ALICE, and LHCb. The physics programs of ATLAS and CMS are focused on the high-energy frontier of particle physics, although they have contributed

notably to heavy-ion physics as well, particularly in the high- p_T regime. LHCb, on the other hand, focuses on b -quark physics and CP-violation – although they have recently begun participating in the heavy-ion program as well. For the bulk of its runtime, the LHC delivers pp beams at its maximal energy for the high-energy frontier physics program. For approximately one month per year, however, the facility delivers beams dedicated to the heavy-ion program. This has consisted of Pb–Pb and pp beams at $\sqrt{s_{NN}} = 2.76$ TeV, as well as Pb–Pb, p–Pb, Pb–p, and pp beams at $\sqrt{s_{NN}} = 5.02$ TeV. A brief run of Xe–Xe nuclei at $\sqrt{s_{NN}} = 5.44$ TeV was also performed in 2017.

The LHC consists of two rings of counter-circulating hadron beams, spanning a 26.7 km tunnel located 45–170 m underground.² The beams are accelerated by an RF cavity and steered by superconducting dipole magnets. Quadrupole focusing magnets control the transverse spread of the beam.

The injection chain is shown in Fig. 2.2. Proton beams are first extracted from hydrogen gas into a linear accelerator Linac2, and then injected to the Proton Synchrotron Booster, where they are accelerated to 1.4 GeV. Lead ions, on the other hand, are vaporized and accelerated in LINAC3 and then the Low Energy Ion Ring. Then, the beams (both pp and Pb–Pb) are injected into the Proton Synchrotron (PS), where they are accelerated to 25 GeV. The PS injects the beams into the Super Proton Synchrotron (SPS), which accelerates the beams to 450 GeV. Finally, the SPS injects the beams into the LHC, which accelerates them to their nominal energy with an RF cavity. The maximum energy obtained by the LHC is limited by the strength of the ≈ 8 T magnetic field delivered by the superconducting dipole magnets that steer the beam around the ring.

Proton beams circulate in bunches of $\approx 10^5$ protons per bunch, with a beam width of ≈ 64 μm . The bunches are typically separated by 25 ns, allowing ≈ 2800 bunches to circulate simultaneously. Given that the beams circulate at $v \approx c$, each bunch transits the 26.7 km ring at ≈ 11 kHz. This yields a bunch crossing rate of up to ≈ 31 MHz. Since there can be multiple collisions per bunch crossing (up to ≈ 50 in ATLAS and CMS), this yields

2. The tunnel consists of eight straight sections and eight arcs. This is not the optimal design for a hadron collider, but rather is an artifact of the LEP ep collider for which the tunnel was originally built, in which synchrotron radiation losses are more prominent.

a delivered collision rate of up to hundreds of MHz, or a design luminosity of $\mathcal{L} \approx 10^{34} \text{ s}^{-1}\text{cm}^{-2}$.

ALICE, however, is not optimized for such high-rate collisions. The TPC and muon spectrometer cannot safely handle pp rates of more than $\approx 700 \text{ kHz}$ [58]. The luminosity delivered to ALICE is therefore “leveled” or diluted by the LHC before delivering collisions in ALICE. For Pb–Pb data taking, the instantaneous luminosity delivered to ALICE in 2015 was $\mathcal{L} \approx 10^{27} \text{ s}^{-1}\text{cm}^{-2}$ ($\approx 8 \text{ kHz}$). During pp data-taking, the instantaneous luminosity delivered to ALICE varies depending on whether we are collecting minimum bias data ($\mathcal{L} \approx 10^{29} \text{ s}^{-1}\text{cm}^{-2}$ or $\approx 10 \text{ kHz}$) or rare triggers ($\mathcal{L} \approx 10^{31} \text{ s}^{-1}\text{cm}^{-2}$ or $\approx 200 \text{ kHz}$ or higher). Note that the detector readout rate is $\approx 1 \text{ kHz}$ in pp collisions, and $\approx 400 \text{ Hz}$ in Pb–Pb collisions. The delivered luminosity is measured by ALICE with a Van der Meer (vdM) Scan, which measures the beam luminosity by scanning the overlap of the two beams in the transverse plane, and observing the collision rate.




Figure 2.2: The LHC accelerator chain [60].

2.2 ALICE tracking system

There are three main tasks of the ALICE tracking system: (1) To determine the spatial location of charged particles, (2) To determine the momentum of charged particles, and (3) To determine the type of each charged particle. By definition, a tracking system must determine the trajectories of particles. In order to determine the momentum of particles, or typically rather p_T , the tracking system measures the curvature of a track in the B-field. To do this effectively, a long lever arm is necessary. In order to uniquely identify a particle, one must determine its mass m and charge q . The charge q can easily be deduced from the sign of curvature in the B-field. The mass can then be determined by charged particle ionization $\frac{dE}{dx}$ in the tracking system.³ To do this effectively, a large number of space points along the trajectory is necessary.

To achieve these goals, the ALICE tracking system consists of a 6-layer silicon Inner Tracking System and a large gas Time Projection Chamber, fully spanning the range $|\eta| < 0.9$, $0 < \phi < 2\pi$ and containing a small material budget $\approx 13\%X_0$. The tracking efficiency and track p_T resolution for the track selection relevant to this analysis is shown in Chapter 3.

2.2.1 ITS

The ITS has several purposes: To measure the primary vertex of the collision, to extend the lever arm of the tracking system, to provide $\frac{dE}{dx}$ measurements at low p_T , and to measure secondary vertices from long-lived unstable particles such as heavy flavor hadrons.⁴ For the purpose of this thesis, the relevant elements are the identification of the primary vertex and the extension of the tracking system lever arm.

The ITS consists of three separate silicon detectors, each comprised of two layers: the Silicon Pixel Detector (SPD), Silicon Drift Detector (SDD), and Silicon Strip Detector (SSD). The innermost and most crucial layers, the SPD, extends to $-2 < \eta < 2$ and has its

3. As described above, a variety of other PID techniques (time of flight, transition radiation, Cerenkov, etc.) are also employed by ALICE.

4. A variety of unstable particles or resonances are produced and decay within a cm scale, such as K_S^0 and Λ , or the mm scale, such as D mesons and b -hadrons.

inner layer 3.9 cm from beam. In order to cope with high track density (tens of tracks per cm^2), the sensor has fine segmentation $50\mu\text{m} \times 425\mu\text{m}$ and $\approx 10\text{M}$ readout channels. This results in a primary vertex resolution of $\approx 12\mu\text{m}$. The SDD and SSD layers extend to an outer radius of 43 cm, and enable sufficient ITS-TPC matching performance.

2.2.2 TPC

The ALICE TPC is the main tracking detector of ALICE, with the ability to identify tracks of charged particles down to $p_T \approx 150 \text{ MeV}/c$ based on their curvature and energy deposition. The TPC spans a radial distance 85-250 cm from the beamline, with an active volume of $\sim 92 \text{ m}^3$, as shown in Fig. 2.3 (left). When a charged particle traverses the TPC, it ionizes the TPC gas, and the resulting electrons are drifted by an electric field to readout endcaps. The readout consists of multi-wire proportional chambers to pre-amplify the detected charged (known as gas amplification), and is divided into 72 sectors (18 inner/outer on each endcap), with over half a million readout pads.

The ALICE TPC allows up to 159 space points along a track's trajectory. The $x - y$ coordinates of the track (perpendicular to the beamline) are determined as the drift electrons induce signal on a 2D readout pad structure. The z -coordinate of the track (along the beamline) is then determined by the time of the drifted electron signal. Tracks are reconstructed using an iterative “inward-outward-inward” Kalman filter based approach.




Figure 2.3: Left: A schematic of the ALICE TPC. Right: dE/dx separation is used to identify charged particles at a given momentum [61].

Typically, the tracking efficiency is 80 – 90%; below $p_T \approx 0.5$ GeV/c, the tracking efficiency drops due to energy loss from the detector material.

The TPC performs PID based on the $\frac{dE}{dx}$ of a track, using a truncated mean of the distribution of cluster energies of the track. This allows for the separation of pions from electrons up to several GeV/c, as well as the identification of pions from kaons and protons up to several GeV/c, as shown in Fig. 2.3 (right).

In high-rate gas TPCs, ion back-flow (IBF) from the gas amplification region to the drift volume (from the same E -field that drifts ionized electrons to the readout) distorts the drift field, and deteriorates tracking and PID performance. To minimize IBF, a structure known as a gating grid is employed. When the gating grid is in the open configuration, electrons can pass through to the gas amplification region with high efficiency. When the gating grid is closed, the ions are then collected before they drift back into the active volume of the TPC. However, the gate must be closed for a relatively long time (~ 200 μ s gate closure per ~ 100 μ s event drift time) due to the slow drift of the ions, limiting the readout rate to approximately 3 kHz. An upgrade to the TPC readout will be discussed in Section 2.4.

2.3 ALICE EMCal

The ALICE EMCal is designed to measure photons (direct photons and decay photons, predominantly from π^0 and η mesons) and electrons over a large range in energy, $E \approx 0.3 - 100$ GeV. The basic principle of the design is that e^+, e^-, γ deposit all of their energy via electromagnetic showers in the calorimeter, and that one can distinguish photons from e^+, e^- by extrapolating charged particle tracks from tracking system to the EMCal in order to determine if the EMCal deposition came from a neutral or charged particle. Moreover, one can distinguish direct photons from merged decay photons by the transverse shape of the electromagnetic shower.

The EMCal is a Pb-scintillator sampling calorimeter, consisting of 77 alternating layers of Pb absorber and polystyrene scintillator. In a sampling calorimeter only part of the energy is detected in the scintillators, but can be calibrated to translate to actual energy. The benefit of such a design is that the calorimeter contains enough material to stop and

contain all high- p_T photons and electrons. The acceptance of the EMCal spans $-0.7 < \eta < 0.7$, $80^\circ < \phi < 188.137^\circ$ and consists of 12,288 towers (“cells”) divided into 12 Super Modules (SMs), each comprised of $24(\phi) \times 48(\eta)$ towers, except two 1/3 SMs comprised of $8(\phi) \times 48(\eta)$ towers. Additionally, a back-to-back arm called the DCal (“di-jet calorimeter”) was commissioned in 2015, covering the regions $0.22 < |\eta| < 0.7$, $260^\circ < \phi < 320^\circ$ and $-0.7 < \eta < 0.7$, $320^\circ < \phi < \sim 327^\circ$ and consisting of 5,376 towers divided into 8 SMs. The calorimeters are located 4.36 m from the beamline.

Each tower has a depth of 24.6 cm, and a transverse size of $\approx 6.0 \times 6.0$ cm, organized in an approximately projective geometry relative to the interaction point. The radiation length of the EMCal is $X_0 \approx 12.3$ mm, meaning that the EMCal contains a total thickness of $20X_0$. The hadronic interaction length of the EMCal is $\lambda \approx 24.6$ cm, meaning that the EMCal contains a total thickness of 1λ . The Moliere radius of the EMCal is $r_M = 3.20$ cm. The Moliere radius describes the transverse size of the electromagnetic shower, specifically it corresponds to the radius of a cylinder containing 90% of the shower. In angular units, the cell size is $(\Delta\eta, \Delta\phi) \approx (0.014, 0.014)$.

The energy resolution of electromagnetic calorimeters is typically parameterized as

$$\frac{\sigma}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c,$$

where a denotes a stochastic term due to statistical fluctuations in the shower deposition, b denotes an electronic noise term, and c denotes an irreducible imperfection term. The direct sum notation means to add the contributions in quadrature. For the ALICE EMCal, a test beam campaign measured $a = 11.3\%$, $b = 4.8\%$, $c = 1.68\%$ [62]. The stochastic term is relatively large due to the fact that this is a sampling calorimeter. Note that the relative energy resolution improves with higher energy, which is opposite to the behavior of tracking systems (where it is determined by curvature).

In addition to photon and electron depositions, charged particles also deposit ionizing radiation in the calorimeter. The average MIP deposition in the EMCal is ≈ 280 MeV, so in fact at low energies, where the tail of the MIP deposition is prominent, there is a comparable contribution of charged and neutral particles in the calorimeter. Additionally,

all hadrons (charged and neutral) also can deposit energy by hadronically scattering from the calorimeter material. Typically, photon depositions are distinguished from hadronic interactions by a combination of the transverse shower shape and matching to charged particle tracks.

In data reconstruction, groups of several cells are clustered together into a single entity intended to represent a particle. Typically, cell energies considered are those above ≈ 100 MeV, and cluster energies are considered above ≈ 300 MeV. Note also that above $p_T > 6$ GeV/c, the two photons from π^0 decay start merging appreciably into a single cluster.⁵ Note moreover that at close particle spacings (e.g. high multiplicities), multiple particles can fall into a single cluster, and the 1-1 correspondence of truth-level particles to detector-level objects is broken, unlike in the tracking system. This is discussed further in Chapter 3.

The EMCal is also used as a photon and jet trigger for ALICE. As described earlier in Chapter 2, the ALICE TPC in LHC Run 2 has a maximal readout rate of ≈ 400 Hz in Pb–Pb collisions. The LHC, however, is capable of delivering Pb–Pb beams at an interaction rate of ≈ 8 kHz. In a minimum-bias triggered scheme, the delivered interactions are randomly sampled and read out according to the TPC-limited rate. By using a fast detector such as a calorimeter, however, one can read out an enhanced fraction of those delivered events satisfying a trigger condition. In particular, high- p_T jets can be read out by triggering on a large energy deposition in the calorimeter. In this way, triggering allows us in principle to read out all the delivered high- p_T events, but only a fraction of the total events.

A jet trigger was implemented in the EMCal based on energy deposited into a given rectangular patch of EMCal cells. A sliding window of 16×16 cells was used, with the patches sliding by one subregion, defined as 8×8 cells. The EMCal has $8(\phi) \times 6(\eta)$ such jet patches, and the DCal has $5(\phi) \times 4(\eta)$, in both cases excluding the 1/3-SMs. Background subtraction in Run 2 is performed online event-by-event using the median background patch energy in the opposite calorimeter, with the leading background patch excluded. The background patches are fixed (non-sliding) patches of the same size as jet

5. These can be identified by their transverse shower shape, which can, for example, be split into two subclusters to calculate invariant mass.

patches. If a background-subtracted jet patch energy exceeded 20 GeV, the jet trigger fires. The jet trigger patch covers an area corresponding to a jet cone of $R \approx 0.13$. The idea of the median subtraction scheme is to subtract a background using a data-driven method (with the background signal constructed from the same detector in which the signal is measured), and being roughly opposite in ϕ , so that the flow background will be similar. The median technique also avoids aging problems with the V0, which was used previously to estimate the background. Note that in the case of a di-jet being present, it will not dramatically bias the background subtraction because by taking the median, we remove the highest-energy patches. A similar gamma trigger was simultaneously implemented using 4×4 cell patches, and a 10 GeV threshold.

The EMCal jet trigger is located at the trigger level L1. The EMCal trigger data stream is as follows: Data from cells is read out in 2×2 cell units called FastORs. At this point, individual cell information is not propagated, but rather only the summed FastOR signal. These FastORs send their signals to several localized units called Trigger Region Units (TRUs). The TRUs send their data to a single Summary Trigger Unit (STU).

During the 2015 Pb–Pb data-taking period, the EMCal jet and gamma triggers were allocated ~ 8 Hz of bandwidth out of the total rate of ≈ 400 Hz. Given the delivered rate of 8 kHz, the trigger must have a rejection factor of ≈ 1000 in order to read out all delivered events satisfying its trigger condition, which is ≈ 50 times larger rejection factor than minimum bias. The trigger thresholds were selected accordingly.

In order to check whether the triggered events capture all of the delivered high- p_T jets, one must compare the triggered jet spectrum with the minimum bias jet spectrum, for samples of equal luminosity. That is, the efficiency is equal to the probability of the trigger firing, given a minimum bias event with a jet of a given p_T : $\varepsilon(p_T) = P(\text{Trig} | \text{MB}, p_T^{jet})$. To obtain a sample of events of equal luminosity of minimum bias and triggered events (i.e. a sample of events in which the minimum bias trigger and jet trigger both examined every event), we can exploit the fact that in ALICE the jet trigger requires as a prerequisite the minimum bias trigger. Therefore, if we know the downscaling factor for minimum bias triggers, we know precisely how many minimum bias events the jet trigger examined. So if we scale the minimum bias sample up by its downscaling factor, this is the baseline for our

trigger efficiency. Taking the ratio of the triggered jet p_T spectrum to this scaled minimum bias spectrum yields the efficiency $\varepsilon(p_T)$. When the trigger efficiency is less than 100%, we record a sample of jets that may systematically differ from MB jets. In particular, one must examine the centrality bias, neutral energy bias, fragmentation bias, multiplicity bias, jet shape bias, and other trigger biases in order to understand the triggered sample. Studies performed in the course of this thesis demonstrated that the Pb–Pb jet trigger bias extends up to $p_T \approx 100$ GeV/c.

2.4 ALICE Upgrades

ALICE has organized an ambitious upgrade plan to be installed during the LHC Long Shutdown 2 in 2018-2020. The detector will be upgraded in order to achieve continuous readout at 50 kHz in Pb–Pb collisions, thereby running in an untriggered mode, and dramatically enhancing (by a factor ≈ 50) the capability to gather minimum-bias data useful for soft probes. There are several major aspects to this upgrade program:

1. Upgrade of the TPC [61, 63]: This is detailed below in 2.4.1.
2. Upgrade of the ITS [64]: The entire ITS system will be replaced with a new 7-layer system with a faster readout, more precise resolution, and less material. The pixel size will be reduced from $50\mu\text{m} \times 425\mu\text{m}$ to $27\mu\text{m} \times 29\mu\text{m}$, and the inner layer will be located closer to the beamline, moving from 3.9 cm to 2.2 cm.
3. And several more, which include the upgrade of the online-offline reconstruction infrastructure (“O²”) [65], upgrade of the Muon Forward Tracker [66], and Trigger and Readout upgrades [67].

2.4.1 TPC Upgrade

For LHC Run 3, an upgrade of the ALICE TPC to operate at 50 kHz with a continuous-readout is planned [61, 63, 68]. In order to operate continuously, the TPC gating grid in its current form must be eliminated in order to remove its resultant dead time. However, this demands an alternate solution to minimize IBF from the gas amplification region to a level

acceptable from the perspective of distortion corrections, such that track reconstruction and analysis have comparable performance to a gating grid solution. One possible solution is to use micro-pattern gas detectors (MPGDs), such as gas electron multipliers (GEMs) [69] and micro-mesh gaseous structures (MMGs) [70], which have intrinsically low IBF. Like all MPGDs, these have a small anode-cathode distance containing a very large electric field, allowing large gas amplification over a small distance and thus short time. Crucially, these two technologies have the benefit of inherently suppressing IBF while operating continuously. To allow these devices to operate under stable conditions while achieving sufficient IBF suppression, a stack of several layers is necessary. Multi-layer MPGD designs allow multiple IBF-suppressing layers as well as flexibility in operational voltages and alignment, with only a small loss in electron transparency. Simulations for the ALICE TPC [10] have shown that at the foreseen gain of 2000 (Ne-CO₂-N₂ 90-10-5), with IBF as high as 2% and energy resolution of 14% (σ/E) or better (for ⁵⁵Fe X-rays), TPC space charge distortions can be corrected to an acceptable level in regards to TPC track finding, PID capability, and momentum resolution.

In ALICE, a substantial R&D effort investigated several possible designs, and the upgrade choice was selected to be a 4-GEM stack configuration. In the course of this R&D effort, several alternate design choices were considered. The main alternate design considered was a 2-GEM/MMG hybrid design [71]. These R&D efforts are described in Appendix A.1. A second alternate design choice was to use a multi-layer extended gating grid [72, 73], which enables a quasi-continuous readout. Simulations were performed for such a design, as detailed in Appendix A.2, and laboratory tests remain ongoing. Each of these designs remain promising, but R&D was incomplete at the time the design choice was required by ALICE.

Chapter 3

Inclusive jet measurements in Pb–Pb collisions

The ALICE Collaboration performs two types of jet reconstruction: “charged jets”, using charged particle tracks from the ALICE tracking system, and “full jets”, which also include neutral particle information from the ALICE EMCal. Charged jets represent only a subset of the jet particles, whereas full jets include all jet constituents, and coincide with the traditional definition of jets. Most of the published ALICE jet results use charged jets, since tracking is the strength of ALICE, and they are technically simpler to measure.¹ However, full jets offer a clear benefit when comparing to theoretical predictions, since full jets are not dependent on a description of the charged particle fraction of jets, and can be directly compared to theoretical calculations.

ALICE has previously measured the inclusive full jet R_{AA} at $\sqrt{s_{NN}} = 2.76$ TeV for jets with radius $R = 0.2$ [74]. The purpose of this thesis is to perform a similar measurement at $\sqrt{s_{NN}} = 5.02$ TeV, and to extend the jet radius to $R = 0.2 - 0.4$. This allows us to measure the R -dependence of inclusive jet R_{AA} , as well as the jet cross-section ratio $\sigma_{R=0.2}/\sigma_{R=0.4}$, which is an inclusive jet shape observable. The analysis also contains several technical improvements to the analysis strategy intended to improve the accuracy of the measurement.

1. Charged jets also have a particular strength for constituent-based jet shape observables, since each jet constituent is unambiguously associated with a single track.

Measurements of the jet R_{AA} at $\sqrt{s_{NN}} = 2.76$ TeV were performed by ALICE, ATLAS, and CMS [74–76]. CMS measured the jet R_{AA} over several jet R from $R = 0.2 - 0.4$, and showed no significant R -dependence. ATLAS measured the central-to-peripheral ratio R_{CP} from $R = 0.2 - 0.5$, and showed a modification of jet R_{CP} with R . Moreover, a measurement of hadron-jet correlations by ALICE showed no significant modification in the R -dependence of energy loss of hadron-triggered jets from $R = 0.2 - 0.5$. Taken together, the picture of the R -dependence of jet energy loss in the range $R = 0.2 - 0.5$ is unclear. At $\sqrt{s_{NN}} = 5.02$ TeV, there is one existing measurement of the jet R_{AA} by ATLAS, performed only for $R = 0.4$ jets with $p_T^{\text{jet}} > 100$ GeV/c [77]. The measurement presented in this thesis is the first measurement of jet R_{AA} at $\sqrt{s_{NN}} = 5.02$ TeV at low p_T^{jet} , and the first at $\sqrt{s_{NN}} = 5.02$ TeV to measure the R -dependence of the inclusive jet spectra.

The methods of this analysis are based heavily on the ALICE inclusive full jet measurement in Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV [74], and the ALICE charged jet suppression measurement in Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV [78]. Additionally, this analysis relies on the charged jet suppression measurement in Pb–Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV [79] for tracking studies in the analyzed dataset.

This chapter is organized as follows. Section 3.1 describes the dataset and analysis objects that are analyzed in this measurement, and their quality assurance. Section 3.2 describes the jet reconstruction procedure. Section 3.3 describes two performance studies carried out in the course of this measurement. Section 3.4 describes the de-convolution procedure to correct the measured jet p_T spectrum for detector and background effects. Section 3.5 presents the systematic uncertainties, and Section 3.6 presents the results. Section 3.7 presents and discusses comparisons of the results to several theoretical predictions.

3.1 Experimental data

3.1.1 Analysis selections

Datasets

In November-December 2015, ALICE measured Pb–Pb collisions delivered by the LHC at $\sqrt{s_{\text{NN}}} = 5.02$ TeV, known internally as the LHC15o data-taking period. This analysis utilizes a data sample of approximately 4.6×10^6 collected minimum bias events in the centrality range 0-10% (after event selection), acquired over 50 runs with globally good tracking detectors and EMCal performance.²

The analysis also uses a Pythia $p_{\text{T,hard}}$ MC production (Pythia8, Monash 2013 tune) with a full GEANT3 ALICE detector simulation, known as the LHC16j5 dataset. The production consists of 20 $p_{\text{T,hard}}$ bins, each populated with approximately 700,000 events, with bin edges: [5, 7, 9, 12, 16, 21, 28, 36, 45, 57, 70, 85, 99, 115, 132, 150, 169, 190, 212, 235, 235+]. The appendix contains a detailed description of how these $p_{\text{T,hard}}$ bins are combined together. The MC is anchored run-by-run to data-taking runs, and so the good runlist in this analysis is defined as the largest existing subset of the good measured runs above, which consists of 48 runs.

Event selection

Figure 3.1 shows the event selection criteria implemented according to the centralized class AliEventCuts. This class implements period-specific standard event cuts including:

- Primary vertex reconstruction: The number of vertex contributors is required to be > 0 , i.e. there is a successfully reconstructed vertex.
- Primary vertex position relative to the interaction point: $-10 \text{ cm} < V_z < 10 \text{ cm}$
- Primary vertex quality:
 - The distance between the SPD vertex and the track vertex is required to be within

². Minimum bias refers to selection of events with the minimal experimentally achievable event trigger, typically including all non-diffractive collisions, and excluding some fraction of diffractive collisions.

certain proximity, by absolute distance as well as relative to the resolution of the SPD and track vertices: $\Delta v_z < 0.2$ cm, $10\sigma_{SPD}$, $20\sigma_{track}$.

- If the SPD vertex only has its z coordinate reconstructed, then the SPD vertex is required to have sufficient resolution: $\sigma_{SPD} < 0.25$ cm.

Basic pileup checks on V0 timing and ZDC timing are also performed. Additionally, we implement a standardized set of out-of-bunch pileup cuts tailored to the LHC15o period:

- Cut on the correlation between the total number of tracks and the number of TPC only tracks.
- Cut on the correlation between the number of TPC+ITS tracks and the number of tracks matched to TOF within the bunch crossing (since the precise timing resolution of TOF can accurately distinguish different bunch crossings).

These cuts take advantage of the fact that different detectors are affected differently by out-of-bunch pileup, and one can therefore correlate multiplicity in different detectors in order to detect events with pileup.




Figure 3.1: Number of minimum bias events passing each event selection criteria in AliEvent-Cuts, over the centrality range 0-90%.

EMCal cluster selection

Three EMCal cell-level corrections are applied in the analysis before the cells are clustered together to resemble detected particles:

- Energy calibration, based on an iterative π^0 mass calibration.
- Bad channel removal, based on the average energy and occupancy of cells.
- Time calibration, in order to account for differences in readout time such as cabling length, as shown in Fig. 3.2.

The cells are then clustered together using a cluster algorithm which finds a local maximum starting from a minimum seed $E_{\text{seed}} = 300$ MeV, and clusters all adjacent cells with minimum $E_{\text{cell}} = 100$ MeV until a local minimum is reached. The cluster energy is taken as the sum of its cell energies. Exotic clusters are then removed by requiring $F_{\text{cross}} < 0.97$ for clusters with $E_{\text{cluster}} > 4$ GeV, where

$$F_{\text{cross}} \equiv 1 - \frac{E_{\text{cross}}}{E_{\text{cell}}},$$

and E_{cross} is the sum of the cell energy of the four adjacent cells sharing a full edge with the leading cell.




Figure 3.2: Distribution of EMCal cell time before calibration (left) and after calibration (right).




Figure 3.3: Distributions of $\Delta\phi$ (left) and $\Delta\eta$ (right) of matched cluster-track pairs after propagation of tracks to the EMCal. Note that the $\Delta\phi$ distribution is broader, due to the ϕ coordinate being impacted by the B-field.

A nonlinearity correction is then applied to the cluster energy, since the calorimeter response becomes nonlinear at both low and high energies. The correction is based on test beam data, determined using clusterization thresholds $E_{\text{cell}} = 50$ MeV, $E_{\text{seed}} = 100$ MeV. The nonlinearity-corrected energy is referred to as $E_{\text{cluster}}^{\text{nonlincorr}}$.

The cluster energy is then modified in an attempt to approximately remove the contribution of charged particles to the cluster energy. To do so, all accepted tracks are propagated to the EMCal by assuming the distance of closest approach (DCA) to the vertex as the starting point of propagation, and a PID-based mass hypothesis. The tracks are propagated to the average shower depth of the EMCal, $R = 440$ cm. Each track is allowed to match geometrically to at most one cluster, while clusters are allowed to have multiple matching tracks. EMCal clusters with $E_{\text{cluster}}^{\text{nonlincorr}} > 150$ MeV are used to assign matches. Figure 3.3 shows the $\Delta\phi$ and $\Delta\eta$ distributions of matched cluster-track pairs.

If a track is matched within $\Delta\phi < 0.3$, $\Delta\eta < 0.15$, then a hadronic correction is applied to the cluster:

$$E_{\text{cluster}}^{\text{hadcorr}} = E_{\text{cluster}}^{\text{nonlincorr}} - \Delta E,$$

with

$$\Delta E = c \sum_i p_i^{\text{track}},$$

where the i spans over all tracks matched to the cluster, p_i^{track} is the track 3-momentum, and c is the speed of light.

After the above cuts and corrections have been performed, two requirements define accepted clusters:

- The cluster has a hadronically-corrected energy $E_{\text{cluster}}^{\text{hadcorr}} > 300 \text{ MeV}$.³
- The cluster time (on the leading cell) satisfies $t_{\text{clus}} \in [-50 \text{ ns}, 100 \text{ ns}]$.

The cluster time cut is motivated by examining the distribution of cluster time, and selecting the main signal region, as seen in Fig. 3.4.




Figure 3.4: Cluster time distribution as a function of cluster energy. Note: Full centrality range 0-90% is plotted.

3. Note that the threshold uses the total cluster energy, rather than E_T , which is used for jet reconstruction.

Track selection

This analysis uses “hybrid” tracks in which two classes of tracks are included:

- Global tracks, including at least one SPD hit.
- Complementary tracks, which do not require an SPD hit, and for which the track is re-fit to be constrained to the primary vertex.

These two classes of tracks are each defined by a specific set of tracking cuts: χ^2 of the track fit, the number of crossed rows, and more. The minimum p_T^{track} is 150 MeV/c, and track acceptance is allowed over the TPC: $-0.9 < \eta < 0.9$, $0 < \phi < 2\pi$.

In the LHC15o period, significant TPC distortions were discovered, with larger magnitude in higher interaction rate runs. Distortion maps were produced by the ALICE tracking experts using the ITS and TRD to constrain TPC tracks, and distortion corrections were implemented in reconstruction. The ALICE charged jet analysis has demonstrated that the interaction-rate dependence of the tracking performance for hybrid tracks is well described by simulation [79].




Figure 3.5: Measured jet spectrum (background subtracted but not unfolded) for $\eta < 0$ and $\eta > 0$, over the centrality range 0-10%. No significant difference is observed between the two, suggesting fake high- p_T tracks are not present in the track selection we use.

A concern was also raised by the ALICE Data Preparation Group regarding fake high- p_T tracks for certain loose track selections in this dataset, caused by matching of uncorrelated ITS and TPC segments due to TPC distortions. The effect is known to increase with interaction-rate, and to exhibit a difference between $\eta < 0$ and $\eta > 0$. It was demonstrated that global tracks have no such problem. For complementary tracks, which have relatively loose selection criteria, no obvious effect was observed (see also Fig 3.8, which shows no obvious problem), but the study was not conclusive. However, the IR-dependent studies in [79] demonstrate that there is no dependence of the high- p_T track rate on the interaction rate. Moreover, we compare the measured jet spectrum at $\eta < 0$ and $\eta > 0$ in Fig. 3.5, and find no significant difference. We can therefore safely conclude that fake high- p_T tracks do not pose an issue for this analysis.

Jet selection

Jets are constructed with FastJet 3.2.1, using the anti- k_T algorithm and the p_T recombination scheme [41, 42]. The jet R considered are $R = 0.2, 0.4$. The examined p_T^{jet} range is $p_T^{\text{jet}} \in [40 \text{ GeV}/c, 140 \text{ GeV}/c]$ for $R = 0.2$ jets, and $p_T^{\text{jet}} \in [60 \text{ GeV}/c, 140 \text{ GeV}/c]$ for $R = 0.4$ jets. Jets are only considered in the EMCal, since the DCal is too small to contain even an $R = 0.3$ jet. See Section 4.1 for a detailed description of jet reconstruction.

The following requirements describe the criteria for a jet to be accepted:

- The center of the jet must be within the fiducial volume of the EMCal, i.e. a distance R from any edge of the EMCal.
- The jet must contain a track with $p_T^{\text{track}} > p_{T,\text{lead,ch}}$, where we take $p_{T,\text{lead,ch}} = 5 \text{ GeV}/c$ or $p_{T,\text{lead,ch}} = 7 \text{ GeV}/c$, depending on the result considered.⁴
- The jet must not contain any tracks with $p_T^{\text{track}} > 100 \text{ GeV}/c$, since tracking is not reliable for such tracks.
- The area of the jet must be $A > 0.6\pi R^2$, since signal anti- k_T jets tend to be circular, whereas background jets can have irregular shapes.

4. We do not extend this requirement to allow a neutral cluster to satisfy the 5 GeV/c requirement, since there is not a one-to-one correspondence between truth-level particles and clusters.

3.1.2 Data quality assurance

Run-by-run quality assurance (QA) was performed on both the LHC15o dataset and the LHC16j5 MC dataset in order to verify consistent quality of the data. In what follows, the basic quality assurance plots summed over the full measured Pb–Pb dataset are summarized. Note that in many cases preliminary calibrations were used in the plots below, and the cuts employed are at times slightly different from those in the final analysis.

Tracks

Figure 3.6 shows the ϕ distribution and η distribution of hybrid tracks, which demonstrate that the hybrid track selection provides an approximately uniform ϕ -acceptance of tracks. Figure 3.7 shows the $\eta - \phi$ distribution of hybrid tracks, which corroborates the uniformity observed in Fig. 3.6. Figure 3.8 shows the p_T^{track} distribution and track p_T resolution (based on varying the track fitting parameters) of hybrid tracks, which demonstrate that the complementary tracks behave similarly to the global tracks.




Figure 3.6: Left: ϕ distribution of global and complementary tracks, as well as their sum. Right: η distribution of global and complementary tracks, as well as their sum. Note: Full centrality range 0-90% is plotted.




Figure 3.7: $\eta - \phi$ occupancy distribution of hybrid tracks, over the centrality range 0-90%. Left: All accepted tracks. Right: Accepted tracks with $p_T^{\text{track}} > 10$ GeV/c. Note that the empty strips of constant ϕ correspond to sector boundaries in the TPC; high- p_T tracks, which have nearly straight trajectories, cannot be reliably tracked if they coincide with a sector boundary.




Figure 3.8: Left: p_T^{track} distribution of global and complementary tracks, as well as their sum, over the centrality range 0-90%. Right: Track p_T resolution of global and complementary tracks, over the centrality range 0-90%.

EMCal clusters

Figure 3.9 shows the ϕ, η distribution of EMCal clusters, in which the Supermodule boundaries can be seen. Figure 3.10 shows the $\eta - \phi$ occupancy distribution of EMCal clusters, and the fraction of clusters with a given number of matched tracks within $\Delta\eta < 0.015$, $\Delta\phi < 0.03$. Figure 3.11 shows the cell energy spectrum after bad channel removal and energy calibration, and exhibits a smooth spectrum as expected. Figure 3.12 shows the cluster energy spectra of EMCal clusters, and their ratio to DCal and PHOS clusters. Note that bump in the ratio of EMCal to PHOS cluster spectra near 3 GeV prompted the investigation described in Section 3.3.1.




Figure 3.9: Left: ϕ distribution of EMCal clusters. Right: η distribution of EMCal clusters. Note: Full centrality range 0-90% is plotted.




Figure 3.10: Left: $\eta - \phi$ distribution of EMCal, DCal, and PHOS clusters. Note: Full centrality range 0-90% is plotted. Right: Fraction of clusters with a given number of matched tracks within $\Delta\eta < 0.015$, $\Delta\phi < 0.03$, for 0-10% centrality, as a function of non-linearity-corrected cluster energy.




Figure 3.11: Right: Cell energy spectrum after bad channel removal and energy calibration, 0-90% centrality.




Figure 3.12: Left: Cluster spectra of EMCAL and DCal, and their ratio. Right: Cluster spectra of EMCAL, DCal, and PHOS, and the ratio to PHOS. Note: Full centrality range 0-90% is plotted. Note that the spectra are not corrected for geometrical acceptance, which is why the ratio does not plateau at 1.

Jets

Figure 3.13 shows the $\eta - \phi$ distribution of $R = 0.2$ full jets in the EMCal fiducial volume. Figure 3.14 shows the calorimeter energy fraction of $R = 0.2$ full jets, and the jet area distribution. Note that at low- p_T^{jet} , where we expect to have many combinatorial jets, the jet area distribution extends to small area, whereas at high- p_T^{jet} , where we expect to have no combinatorial jets, the area distribution narrows near $A = \pi R^2$. Figure 3.15 shows the distribution of p_T^{track} and z of the leading charged hadron in accepted jets. A detailed description of jet reconstruction is presented in Section 3.3.




Figure 3.13: $\eta - \phi$ distribution of $R = 0.2$ full jets in the EMCal fiducial volumes. Note: 0-90% centrality.




Figure 3.14: Left: Calorimeter energy fraction of accepted $R = 0.2$ full jets. Right: Jet area of $R = 0.2$ full jets, before area cut. Note: 0-10% centrality.




Figure 3.15: Left: Leading track distribution of $R = 0.2$ full jets, as a function of jet p_T (before leading hadron requirement). Right: Leading track z over the range $p_{T,\text{det}}^{\text{jet}} \in [40, 140]$ GeV/c. Note: 0-10% centrality.

3.2 Jet reconstruction

3.2.1 Jet clustering

In the ALICE scheme of full jet reconstruction, we must combine tracking information and EMCal information to measure the “full” jet energy. In this way, we can measure charged particles in the tracking system ($\pi^\pm, K^\pm, p^\pm, e^\pm, \mu^\pm$), and also measure the majority of neutral particles (γ , including direct photons and decay photons mostly from π^0) in the EMCal. We fail, however, to reliably measure K_L^0 or $n/\bar{n}$, since these can only interact hadronically, and the EMCal is too thin to consistently measure them. Our strategy is to model these missing neutral particles via Pythia in order to report a result containing the full jet energy. In order to account for the fact that charged particles also deposit energy in the EMCal, we follow the approach taken in [74] to employ a hadronic correction method in which tracks and clusters are geometrically matched, and if a cluster has one or more matched tracks, p_T^{track} is subtracted from p_T^{cluster} as described in Section 3.1.

Jets are constructed using the anti- k_T sequential recombination algorithm. The jet constituents are then combined using the p_T recombination scheme,⁵

$$p_T^{\text{jet}} = \sum_i p_{T,i}^{\text{track}} + \sum_j p_{T,j}^{\text{cluster}}.$$

The tracking system directly measures p_T^{track} . The calorimeter, however, measures the energy of the cluster. The cluster energy E_{cluster} can be related to the transverse energy $E_T^{\text{cluster}} \equiv \sqrt{p_T^2 + m^2}$ by

$$\frac{E_{\text{cluster}}}{E_T^{\text{cluster}}} = \cosh \eta.$$

We assume that clusters are massless, leading to:

$$p_T^{\text{cluster}} = \frac{E_{\text{cluster}}}{\cosh \eta},$$

5. Unlike the standard E-scheme, in which the four-vectors are combined to form a jet four-vector, the p_T -scheme imposes a re-scaling on the four-vectors to make the energy equal to the 3-momentum. For ALICE, this is a natural scheme, since we directly measure p_T^{track} , and we assume that clusters are massless (the majority of them arise from photons).

i.e. p_T^{cluster} can be obtained from the measured E_{cluster} and η . We always use the hadronically-corrected cluster energy in jet reconstruction.

Note that the η, ϕ values used in the jet finder are the η, ϕ at the EMCal for EMCal clusters (i.e. we assume the clusters are neutral), and the initial η, ϕ at production for tracks. The fact that we do not have exact knowledge of which particles should be clustered together in the jet can cause the improper inclusion of charged particles that make a deposit in the EMCal. In such cases, if the charged particle was successfully tracked, then its contribution will anyway be removed by the hadronic correction – but if it is untracked, its p_T will be mistakenly clustered into the jet (e.g. an untracked electron which originated at an η, ϕ outside of the jet cone). The response matrix in Section 3.4 takes this effect into account, however.

3.2.2 Background subtraction

The background density ρ is determined each event, and used to subtract the average background from each jet in that event:

$$p_{T,\text{corr}}^{\text{jet}} = p_T^{\text{jet}} - \rho A.$$

To determine ρ for full jets, we are unable to directly compute the full-jet p_T -density in the calorimeters, since the partial acceptance of the calorimeters admits only a small number of jets per event [74]. Instead, to compute ρ in each event, we first find $R = 0.4 k_T$ charged jets, and exclude the two leading jets from the collection. We then compute the median p_T -density of the remaining sample:

$$\rho_{\text{charged}} = \text{med} \left(\frac{p_T^i}{A_i} \right).$$

The full-jet p_T -density is then determined by measuring the ratio of possible jet constituents in full jets compared to charged jets. Specifically, we apply a scaling factor s :

$$s(C) = \frac{\left(\sum p_{T,\text{track}}^{\text{calo}} + p_{T,\text{cluster}}^{\text{calo}} \right) / A_{\text{calo}}}{\sum p_{T,\text{track}}^{\text{TPC}} / A_{\text{TPC}}}$$

where the sum in the numerator spans the acceptance of the calorimeter, the sum in the denominator spans the acceptance of the TPC, A is the $\eta - \phi$ acceptance of the labeled detector, and C is the event centrality. This data-driven method is expected to be an accurate measure of the scaling since it precisely uses the constituent objects that enter the jet finder, and thereby accounts for detector inefficiencies (tracking inefficiency, dead channels, sector boundaries, etc.). Note that the scale factor s is computed as a function of centrality, since medium effects alter the composition of neutral to charged energy (e.g. low- p_T thermal photons). The mean of the event-by-event scale factors at each centrality is parameterized with a second-order polynomial over the range $C \in [0, 50]$, as shown in Fig. 3.16.

The full-jet average background density is then

$$\rho(C) = s(C) \times \rho_{charged},$$

which can be seen in Fig 3.17.

In data, we compute $s(C)$ with precisely the same analysis cuts that will be used in the final analysis. Note that the effect of the clusterization cell thresholds and the hadronic correction will be large on the background scale factor, since these are large corrections in the soft sector.

In principle one doesn't need to subtract the background of a jet, but could put this in the unfolding procedure (as long as an embedding-based unfolding is used). However, if we fail to take advantage of the event-by-event information that we know, the overall uncertainty will be larger. Figure 3.17 demonstrates that the variance of ρ is quite significant.

3.2.3 Background fluctuations

To study jet-by-jet fluctuations in the background, in each event we generate a random (η, ϕ) within the fiducial calorimeter acceptance, and compare the sum of constituents in an $R = 0.2$ cone to the expected average background in that cone:

$$\delta_{p_T} = \left(\sum_i p_T^{\text{track}} + p_T^{\text{cluster}} \right) - \rho \pi R^2,$$




Figure 3.16: Left: Background scale factor as a function of centrality, histogrammed over all events. Right: Mean background scale factor, parameterized by a second-order polynomial.




Figure 3.17: Full jet average background density, $\rho(C)$.

where i spans the random-cone. The width of the δ_{p_T} distribution is a measure of the size of the background fluctuations. Figures 3.18, 3.19, 3.20 show the measured δ_{p_T} for $R = 0.2, 0.3, 0.4$, respectively.

Given a measurement of δ_{p_T} , we can in principle select a p_T range (e.g. 5σ) above which we can be assured that nearly all of the jets are real jets. Table 3.1 summarizes the widths of the δ_{p_T} distributions for each R .

The mean of the δ_{p_T} distribution should be close to zero if ρ was determined accurately. However, the distribution is expected to have a slight positive skew, as previously demonstrated by modeling the multiple contributions to the background fluctuations: random fluctuations in particle number and p_T , jet contributions, and flow contributions [80]. One can also study the background fluctuations by embedding a jet or a single high- p_T track




Figure 3.18: δp_T as a function of centrality for $R = 0.2$ jets, using the random cone method.

into the background, and measuring the difference between the reconstructed jet p_T and the probe's true p_T [74]. However, in the present analysis the background fluctuations will not be explicitly used except to consider at what p_T^{jet} to report measurements, and we therefore do not pursue them further.

0-10% Centrality	$\sigma_{\delta p_T}$ (GeV/c)	$5\sigma_{\delta p_T}$ (GeV/c)
$R = 0.2$	6.5	32.5
$R = 0.3$	10.9	54.5
$R = 0.4$	16.1	80.5

Table 3.1: Standard deviation of δp_T for each R in 0-10% centrality, as well as $5\sigma_{\delta p_T}$.




Figure 3.19: δp_T as a function of centrality for $R = 0.3$ jets, using the random cone method.




Figure 3.20: δp_T as a function of centrality for $R = 0.4$ jets, using the random cone method.

3.2.4 Impact of jet cuts on the measured spectra

In this section, we present the impact that the area cut, the leading hadron requirement, and the 100 GeV/c track prohibition have on the measured jet spectra. The fraction of measured jets that pass each cut, which we refer to as the detector-level cut efficiency, is plotted in Fig. 3.21 as a function of measured jet p_T for $R = 0.2$ for both simulated pp data and measured Pb–Pb data. Note that the jet p_T has not been unfolded, and the detector-level efficiency curves should be interpreted with this in mind – they show the effect of the cuts on the measured spectra, rather than the likelihood of truth-level jets to pass the cut, which will be shown in the next section as the jet reconstruction efficiency.

We see, as expected, that the most impactful cut is the 5 GeV/c leading track requirement. In the pp det-level case, we see that this cut is nevertheless a small bias above $p_{T,\text{det}}^{\text{jet}} \approx 40$ GeV/c. Note that the location of the turn-on of the 5 GeV/c leading track efficiency in the Pb–Pb case is not the “true” turn-on, since, for example, background fluctuations tend to smear the spectrum to higher- p_T compared to the unfolded spectrum. Nevertheless, we observe that at high- p_T , the efficiency of this cut reaches $\approx 90 - 95\%$, with the rejected jets presumably containing a neutral leading constituent. The area cut in Pb–Pb is negligible except at very low p_T . The maximum track p_T cut is seen to cut only a small number of jets at large p_T . In the Pb–Pb case, this must be interpreted with the caveat that some of these measured jets have been smeared up in p_T by background fluctuations, and so trivially can’t contain a 100 GeV/c track. However, the pp case shows a similarly small effect of the cut up to ≈ 150 GeV/c. Figure 3.22 shows the z_{leading} distribution of charged tracks for Pb–Pb jets with $120 < p_{T,\text{det}}^{\text{jet}} < 140$ GeV/c, and we see that less than a few percent of jets are removed by this cut. This can also be seen in Fig. 3.15. We therefore expect that while these removed jets may have different quenching properties (since e.g. harder fragmenting jets are more likely to be quark jets), there are too few of them to make a significant impact. For the present analysis, we utilize $p_{T,\text{det}}^{\text{jet}}$ up to 120 GeV/c.

The detector-level jet cut efficiencies for $R = 0.3$ and $R = 0.4$ are shown in Fig. 3.23 and Fig. 3.24. Note that in the Pb–Pb case the leading track requirement removes more

jets at higher p_T compared to $R = 0.2$ since background fluctuations are larger for larger R , resulting in higher- p_T fake jets (and larger upward smearing of jet p_T). Note also that the measured jet energy scale depends on R , with larger R jets containing more of the true jet energy. The area and maximum track p_T cuts have similarly small effect as for $R = 0.2$.




Figure 3.21: Left: Fraction of simulated pp det-level jets that pass each jet selection cut, for $R = 0.2$. Right: Fraction of measured Pb–Pb jets that pass each jet selection cut, for $R = 0.2$ in 0-10% centrality.




Figure 3.22: Charged track $z_{leading}$ distribution in Pb–Pb data for jets with $120 < p_{T,det}^{jet} < 140$ GeV/c, for $R = 0.2$ in 0-10% centrality. Jets with $z_{leading}$ larger than ≈ 0.8 on this plot would be removed by the 100 GeV/c track cut (and approximately corrected for by the jet reconstruction efficiency).




Figure 3.23: Left: Fraction of simulated pp det-level jets that pass each jet selection cut, for $R = 0.3$. Right: Fraction of measured Pb–Pb jets that pass each jet selection cut, for $R = 0.3$ in 0-10% centrality.




Figure 3.24: Left: Fraction of simulated pp det-level jets that pass each jet selection cut, for $R = 0.4$. Right: Fraction of measured Pb–Pb jets that pass each jet selection cut, for $R = 0.4$ in 0-10% centrality.

3.2.5 Jet performance

Jet reconstruction performance is typically estimated by three quantities:

- Jet energy scale (JES) shift: $\Delta_{\text{JES}} = \left\langle \frac{p_{\text{T,det}}^{\text{jet}} - p_{\text{T,gen}}^{\text{jet}}}{p_{\text{T,gen}}^{\text{jet}}} \right\rangle$. That is, for a fixed $p_{\text{T,gen}}^{\text{jet}}$, the mean measured $p_{\text{T,det}}^{\text{jet}}$.
- Jet energy resolution (JER): $\text{JER} = \frac{\sigma(p_{\text{T,det}}^{\text{jet}})}{p_{\text{T,gen}}^{\text{jet}}}$. That is, for a fixed $p_{\text{T,gen}}^{\text{jet}}$, the width of the distribution of $p_{\text{T,det}}^{\text{jet}}$ that is measured.
- Jet reconstruction efficiency: Given a truth-level jet with $p_{\text{T,gen}}^{\text{jet}}$, the probability that we will reconstruct it as an accepted jet at any $p_{\text{T,det}}^{\text{jet}}$.

The pp jet reconstruction performance is shown below for $R = 0.2$ full jets: Figure 3.25 shows the jet energy scale shift and the distribution of JES shifts for various $p_{\text{T,gen}}^{\text{jet}}$ intervals, and Figure 3.26 shows the jet energy resolution and the jet reconstruction efficiency, where we require both the det-level and truth-level jets to contain a 5 GeV/c charged particle. Unlike single particle observables, the measured jet energy has large fluctuations due to the fact that jet energy is an observable involving many particles, and the single-particle tracking efficiency leads to a stochastic fluctuation in the number of jet particles reconstructed. Note that the jet energy scale shift decreases at lower p_{T} , since the accepted jets contain a 5 GeV/c leading charged particle, and therefore are ensured to reconstruct a significant fraction of their energy. The jet reconstruction efficiency is determined by the various cuts specified in Section 3.1 defining accepted jets, and is dominated by the 5 GeV/c leading track requirement, since in certain cases we fail to track the leading charged hadron [74]. Note that the pp response approximately, but not exactly, describes the detector effects in jet reconstruction relevant for this analysis.

The Pb–Pb jet reconstruction performance from embedding pp MC events into Pb–Pb data (as described in detail in Section 3.4) is shown below: Figure 3.27 shows the jet energy scale shift for $R = 0.2, 0.3, 0.4$ jets and the distribution of JES shifts for various $p_{\text{T,gen}}^{\text{jet}}$ intervals. Note that the mean JES shift depends on R primarily due to background fluctuations, which can be seen by examining the mean values of the $\delta_{p_{\text{T}}}$ distributions in Section 3.2.3. Similarly, we see that the background fluctuations smear the JES distributions compared to




Figure 3.25: Left: Mean jet energy scale shift for $R = 0.2$ jets in pp. Right: Jet energy scale shift distribution in pp for $R = 0.2$ jets for various $p_{T,\text{gen}}^{\text{jet}}$ intervals. Note that the 5 GeV/c leading charged track requirement is imposed at both truth-level and det-level.




Figure 3.26: Left: Jet energy resolution for $R = 0.2$ jets in pp. Right: Jet reconstruction efficiency for $R = 0.2$ jets in pp. Note that both the det-level and truth-level jets have a 5 GeV/c charged particle requirement.

the pp case. Figure 3.28 shows the jet energy resolution and the jet reconstruction efficiency for $R = 0.2, 0.3, 0.4$ jets, where we require the truth-level jet to contain a 5 GeV/c charged particle. At low p_T , the JER is dominated by background fluctuations (which leads to broad JER), while at high p_T the JER is dominated by detector effects. Note that the Pb–Pb jet reconstruction efficiency is slightly smaller than the pp jet reconstruction efficiency; Section 3.4 contains a detailed description of how the jet reconstruction efficiency is computed.




Figure 3.27: Left: Mean jet energy scale shift for $R = 0.2, 0.3, 0.4$ jets in Pb–Pb. Right: Jet energy scale shift distribution in Pb–Pb for $R = 0.2$ jets for various $p_{T,\text{gen}}^{\text{jet}}$ intervals.




Figure 3.28: Left: Jet energy resolution for $R = 0.2, 0.3, 0.4$ jets in Pb–Pb. Right: Jet reconstruction efficiency for $R = 0.2, 0.3, 0.4$ jets in Pb–Pb. Note that both the det-level and truth-level jets have a 5 GeV/c charged particle requirement.

3.3 Performance studies

Several performance studies were done in order to investigate potential modifications to the ALICE full jet analysis strategy at 2.76 TeV [74]. This section outlines the two investigations that result in the two main changes from the previous full jet spectrum measurement: Using an embedding-based response matrix, and raising the EMCal cell thresholds relative to those used in [74].

3.3.1 Centrality-dependence of the EMCal cluster spectrum

A large centrality-dependent “bump” in the ratio of the cluster energy spectrum of the EMCal relative to PHOS was observed near $E_{\text{cluster}} \approx 3.5$ GeV, as shown in Fig. 3.29, with the largest discrepancy in central events. This is not to be seen necessarily as a problem, since the two calorimeters have many differences including their granularity and detector material (the EMCal is a Pb-scintillator sampling calorimeter, while PHOS is a PbWO_4 crystal calorimeter), and therefore different hadronic responses. Nevertheless, being unanticipated, the observation warranted further investigation.

The magnitude of the bump was observed to be noticeably reduced (though not eradicated) by raising the EMCal clusterization thresholds, applying the hadronic correction to EMCal clusters, or rejecting EMCal clusters with matched tracks. Moreover, it was found that the bump is essentially completely removed if a cluster shape cut is placed on EMCal clusters in order to select the single-photon dominated region, as shown in Fig. 3.30. The relevant cluster shape parameter for the EMCal is $M_{02} \equiv \lambda_0^2$, where λ_0 is the eigenvalue of the major axis of the spatial distribution of cells logarithmically weighted by their energy. This parameter has been studied in detail in pp collisions, where it has been demonstrated to be an effective selector of single photons, which appear in a sharply peaked distribution around $\lambda_0^2 \approx 0.25$. It has also been studied to a limited extent in Pb–Pb collisions. This indicates that the cause of the bump is due to clusters that are not single photons.

In fact, when we compare different selections of λ_0^2 in the EMCal itself, we also observe the bump, as shown in Fig 3.31. This allows us to pursue the issue in a way that is independent of PHOS.

Note that in general, there is a complicated mixture of particles comprising the cluster spectrum, depending on the cluster energy: single photons, merged photons, hadronic minimum ionizing depositions, partially contained hadronic showers, electrons, and more. In central Pb–Pb, there are additionally a large number of uncorrelated overlaps of multiple such particles, in addition to a change in particle composition. In fact, the λ_0^2 distribution is observed to be dramatically different in central Pb–Pb compared to pp, as seen in Fig. 3.32 (bear in mind the distribution depends dramatically on the cell thresholds and




Figure 3.29: Cluster energy spectra of EMCal, DCal, and PHOS, as well as the ratio of EMCal and DCal to PHOS. Each of the four plots shows a different centrality range. The spectra are not normalized for acceptance.




Figure 3.30: Cluster energy spectra of EMCal, DCal, and PHOS, as well as the ratio of EMCal and DCal to PHOS, 0-10% centrality, with cluster shape cuts applied to select the single-photon dominated region. The spectra are not normalized for acceptance.




Figure 3.31: EMCal cluster energy spectra for two selections of λ_0^2 . Left: 0-10% centrality. Right: 50-90% centrality, with cluster shape cuts applied to select the single-photon dominated region. The spectra are normalized such that $\int_{10 \text{ GeV}}^{20 \text{ GeV}} \frac{dN}{dE} dE = 1$.

track-matching selections).

Certain centrality-dependent effects, such as the enhancement of uncorrelated particle overlaps, are expected to be present in HIJING. We find that HIJING qualitatively produces such bump structures in central but not peripheral collisions, but for λ_0^2 values different than




Figure 3.32: EMCAL λ_0^2 distributions for a variety of cluster energies. Left: 0-10% centrality. Right: 50-90% centrality. The spectra are normalized such that $\int_{0.1}^{0.4} \frac{dN}{d\lambda_0^2} d\lambda_0^2 = 1$. Clusterization thresholds $E_{\text{cell}} = 100$ MeV, $E_{\text{seed}} = 300$ MeV.

that observed in data, and with a smaller bump magnitude, and slightly different shape, shown in Fig. 3.33. In addition, upon examining the HIJING simulation by truth-level cluster contributors, we find that clusters consisting of overlaps of multiple particles, as well as certain other cluster types, exhibit a bump. One should bear in mind that the particle composition in HIJING is expected to disagree significantly with measured data, and the referenced plots here have not been re-weighted for such effects. Moreover, one should bear in mind that the MC description of λ_0^2 in pp is imperfect, and that it is a very sensitive parameter due to its logarithmic weighting. Nevertheless, the fact that HIJING qualitatively generates the features we observe in data is convincing evidence that the bump is a legitimate physical effect. We do not speculate on the quantitative cause, since it can be due to a complicated mixture of particle composition effects and uncorrelated particle overlap effects.

The presence of the bump has the potential to alter the detector response in the calorimeter as a function of centrality, namely by measuring additional energy in central Pb-Pb events (e.g. an under-threshold particle overlapping with another particle). Such effects would not be taken into account in a jet measurement by a pp-based response matrix. However, since the effects are quite complicated, attempting to model them explicitly would likely result in large systematic uncertainties. Instead, we employ a data-driven correction. In particular, we can use a data-driven jet response matrix by embedding a pp MC jet into a Pb-Pb data background. This procedure will be described in detail in Section 3.4. This




Figure 3.33: EMCal cluster energy spectra for two selections of λ_0^2 , using LHC16g1 (HIJING+GEANT3). Left: 0-10% centrality. Right: 50-90% centrality, with cluster shape cuts applied to select the single-photon dominated region. The spectra are normalized such that $\int_{10 \text{ GeV}}^{20 \text{ GeV}} \frac{dN}{dE} dE = 1$.

approach will automatically contain the correct magnitude and energy dependence of the bump effects since we embed into measured data. Note that the bump still may have a significant impact on the raw reconstructed jet, since there are more clusters near 3.5 GeV in central Pb–Pb relative to pp, and some of these may be charged particles that should not be overcounted. The data-driven embedding approach is intended to correct for any such effects. Note that since we embed a pp probe jet, implicit to this approach is the assumption that the hadronic response of the calorimeter is well-described in pp, which has been shown to be the case (see [81], Section 6.3.3).

An alternate idea to avoid these centrality-dependent bump effects is to employ a λ_0^2 cut in jet reconstruction to select the single-photon dominated region. That is, instead of selecting all clusters and correcting for the charged particle depositions, one would predominantly select photon clusters, and reject hadronic clusters and other large- λ_0^2 contributors, thereby bypassing unknown effects of the ill-understood bump region. There are several complications that would need to be addressed, such as the increased dependence on the MC description of λ_0^2 and biases to the λ_0^2 distribution inside a jet, as well as the need for a scheme to accept merged π^0 . That is, the single photon efficiency of the cut, as well as




Figure 3.34: The fraction of all clusters remaining after a cut $0.1 < \lambda_0^2 < 0.4$ is plotted as a function of cluster energy, in both central and peripheral events. Note that there is a large dip in the cluster efficiency in central relative to peripheral events.

the merged π^0 efficiency and hadron contamination would need to be addressed. Figure 3.34 suggests that the single photon efficiency is not very high, i.e. the large number of combinatorial overlaps in central Pb–Pb collisions would mean that a significant fraction of photons would be removed from our jet reconstruction by applying such a cut. The method was therefore not further pursued.

Alternately, one could in principle improve the precision of the jet measurement by leveraging the event-by-event information we measure, in particular cluster-track matching and the cluster shape. By taking these cluster properties into account, we could in principle narrow the JER because on a jet-by-jet basis we gain information about the overestimation of the jet energy (i.e. some jets may consist of clusters with only a small amount of overcounting, while others may have a large overcounting). This method was however not pursued.

3.3.2 EMCal cell thresholds

The previous full jet spectrum measurement in ALICE used cell clusterization thresholds of $E_{\text{cell}} = 50$ MeV, $E_{\text{seed}} = 100$ MeV. However, subsequent studies suggest that these cell thresholds are too low and may contain noise that is unaccounted for by the MC. Accordingly, this analysis investigated three sets of cell thresholds to examine if they result in differences in the performance of an inclusive jet measurement: $E_{\text{cell}} = 50$ MeV, $E_{\text{seed}} = 100$ MeV; $E_{\text{cell}} = 100$ MeV, $E_{\text{seed}} = 300$ MeV; $E_{\text{cell}} = 150$ MeV, $E_{\text{seed}} = 300$ MeV.

Raising the cell thresholds is not expected to impact jet reconstruction dramatically, since a relatively small fraction of the jet p_T consists of soft contributions. Nevertheless, the contribution may be significant. In order to study the impact of changing the cell thresholds, Pythia jets are embedded into MB Pb–Pb data, following the procedures described in Section 3.4, with the idea to evaluate the jet energy scale shift and jet energy resolution, in order to gain and approximate understanding of the effect of the cell thresholds on jet reconstruction. Ultimately, however, the most important metric is the final uncertainty on the unfolded jet spectra. Each of these metrics will be presented below. Unless otherwise noted, no timing cuts are applied in these studies.

Note that removing soft contributions is expected to have a large effect on the combinatorial background. Accordingly, the cell thresholds may also impact track-matching significantly, and the hadronic correction. The present studies are performed using the standard hadronic correction technique from the previous analysis. Below are listed the background scale factors for the various EMCal cell thresholds (the background scale factor plots are included in Appendix B):

- $E_{\text{cell}} = 50$ MeV, $E_{\text{seed}} = 100$ MeV: $s(C) = 1.785 - 0.00815C + 0.000064C^2$.
- $E_{\text{cell}} = 100$ MeV, $E_{\text{seed}} = 300$ MeV: $s(C) = 1.443 - 0.00490C + 0.000054C^2$.
- $E_{\text{cell}} = 150$ MeV, $E_{\text{seed}} = 300$ MeV: $s(C) = 1.397 - 0.00355C + 0.000038C^2$.

The δ_{p_T} width is slightly reduced with higher cell thresholds, as shown in Table 3.2; the δ_{p_T} distributions are included in Appendix B. Note that $5\sigma_{\delta_{p_T}}$ is listed for each case to roughly indicate where combinatorial jets become negligible.

0-10% Centrality	E_{cell} (MeV)	$\sigma_{\delta_{p_T}}$ (GeV/c)	$5\sigma_{\delta_{p_T}}$ (GeV/c)
$R = 0.2$	50	6.9	34.5
	100	6.5	32.5
	150	6.3	31.5
$R = 0.3$	50	11.8	59.0
	100	10.9	54.5
	150	10.6	53.0
$R = 0.4$	50	17.5	87.5
	100	16.0	80.0
	150	15.6	78.0

Table 3.2: Standard deviation of δ_{p_T} for each R in 0-10% centrality for various EMCAL cell thresholds, as well as $5\sigma_{\delta_{p_T}}$. Note that cell time cuts $t_{\text{cell}} \in [-50\text{ns}, 50\text{ns}]$ were applied for this δ_{p_T} computation.

The jet energy scale shift is plotted in Fig. 3.35 and as expected shows that a small amount of jet energy is removed when the cell thresholds are increased. The jet energy resolution is plotted in Fig. 3.36, and shows that at low p_T^{jet} , the jet energy resolution is improved by $\approx 1 - 2\%$. The jet spectra were then unfolded using SVD unfolding [82]. The main result is shown in Fig. 3.37, with statistical uncertainties propagated through the unfolding algorithm. Details of the unfolding procedure for these studies can be found in Appendix B.

From these, we can see that the effect of the cell thresholds is small. The benefit of raising the cell thresholds is to reduce potential noise, and as a secondary benefit, to reduce background fluctuations. The drawback is a slightly increased reliance on simulation to




Figure 3.35: Jet energy scale shift for 0-10% centrality $R = 0.2$ jets, for three sets of EMCAL cell thresholds.




Figure 3.36: Jet energy resolution for 0-10% centrality $R = 0.2$ jets, for three sets of EMCal cell thresholds.




Figure 3.37: Left: Unfolded spectra for 0-10% centrality $R = 0.2$ jets, for three sets of EMCal cell thresholds. Right: The corresponding statistical unfolding uncertainties.

describe the low-energy contributions that we neglect. We see, in fact, that there is no significant difference in unfolding uncertainties between the various cell thresholds. Note, however, that these studies were done with no timing cuts. Applying timing cuts is expected to induce a small disagreement between data and MC at low cluster energy, since the timing resolution of low-energy clusters becomes increasingly broad. This effect is expected to be strongest for the lowest cell thresholds. In accordance with the above considerations, we therefore select $E_{\text{cell}} = 100$ MeV, $E_{\text{seed}} = 300$ MeV for the analysis.

3.4 Unfolding the jet p_T spectrum

The reconstructed p_T^{jet} fails to account for a number of effects, such as:

- Background fluctuations
- Detector effects
 - Tracking inefficiency
 - Missing long-lived neutral particles: n, K_L^0
 - Track p_T resolution
 - Gaps in acceptance
 - Material loss in front of EMCal
 - Hadronic correction over- and under-subtraction
 - EMCal energy spectrum “bump”
 - Untracked charged particles outside of the jet cone depositing in the EMCal
 - Approximating the p_T of π^0, η by the E_T of their decay photons

We seek to correct the measured jet spectrum to the “truth”-level jet spectrum at the hadron-level.⁶ In order to correct the p_T^{jet} spectrum simultaneously for detector effects and background fluctuations, we generate a response matrix that simultaneously describes both, and then use a statistical unfolding procedure to correct the measured p_T^{jet} to the true jet p_T . To do this, we embed a Pythia jet into Pb–Pb data. This approach produces a single response matrix, and doesn’t rely on the assumption of [74] that the response factorizes into separate components for detector effects and background fluctuations. Moreover, this approach ensures that the detector response more accurately reflects the Pb–Pb response rather than assuming the pp response. The embedding approach therefore has several potential advantages compared to a factorized approach:

6. Note that we correct the jet p_T to include the “missing” long-lived neutral particles, since while in general we wish to minimize the model-dependence of our measurement, these particles in fact give significant depositions in the calorimeter (the EMCal has $\lambda \approx 1$), and so we must model them in the detector response.

- It ensures we capture any centrality-dependent effects of the detector response or analysis strategy:
 - The centrality-dependent enhancement of the EMCal cluster spectrum, including particle overlaps in the calorimeter as well as the Pb–Pb particle composition, as discussed in Section 3.3.
 - It ensures the effect of the hadronic correction is equivalent in data and in the response.⁷
 - It accounts for untracked charged particles which originate outside of the jet cone and deposit an accepted cluster in the EMCal.
- It ensures any correlations between the background and detector acceptances/inefficiencies is properly accounted for (e.g. if a jet is measured in a region that contains a large number of bad EMCal channels, it will be statistically accounted for).
- It ensures that any residual background that remains after the event-by-event background subtraction will be accounted for.

The magnitude of these effects is expected not to be large,⁸ but may still be significant.

3.4.1 Constructing the response matrix

Embedding details

We seek to build a response matrix describing how, given a truth-level jet with a given p_T , we will reconstruct that jet’s p_T at detector-level. To construct the response matrix, we embed a Pythia event (from LHC16j5), which contains both truth-level and detector-level information, into Pb–Pb data after the detector-level reconstruction has been run

7. Note that the algorithm being the same in data and in the response is not sufficient, but rather we want the effect of the correction to be equivalent in data and in the response. For example, if the algorithm over-subtracts energy in Pb–Pb (i.e. background tracks matching to jet clusters), but a pp detector response is used, then the JES will be too low.

8. The full jet analysis at 2.76 TeV used a factorized approach for the main result, but performed an embedding study which demonstrated that within unfolding uncertainties, the two approaches are consistent. Details on the performed procedure are somewhat sparse, however. It remains possible that the cluster “bump” shifts the JES upward (e.g. overlaps of under-threshold MIP depositions with photon clusters), while hadronic over-subtraction (due to combinatorial overlaps) shifts the JES downward.




Figure 3.38: Left: Tracking efficiency of all charged physical primary particles in HIJING, for different centralities. Note the dip in tracking efficiency at $p_T \approx 2$ GeV/c from tracks crossing a TPC boundary. Right: Ratio of central to peripheral tracking efficiencies in HIJING. The ratio can be seen to be approximately independent of p_T .

individually on both. For tracks, this means the set of tracks in the “hybrid” event is the sum of all tracks in both events individually, without any additional track finding. For EMCal depositions, on the other hand, we re-cluster cells from both events into a single set of clusters.⁹ Typically, each Pythia event is sampled on order 30 times.

The truth-level jet is constructed from the primary particles of the pp event, defined as all particles with a proper decay length longer than 1 cm, but no daughters of these particles [83]. This includes the strange hadrons that decay weakly in the detector – but not their decay products.

The tracking efficiency is known to be slightly worse in Pb–Pb compared to pp. In order to examine the centrality-dependence of the tracking efficiency, we follow the approach in [78], and examine the tracking efficiency in HIJING as a function of centrality. This can be seen in Fig. 3.38. By comparing the tracking efficiency in central to peripheral HIJING detector simulations, we accordingly approximate this effect by randomly rejecting 2% of the pp tracks, independent of p_T .

Additionally, the tracking efficiency in Pythia is known to be slightly incorrect due to the fact that Pythia uses a particle composition slightly different than nature. In particular, in the few GeV/c range, Pythia under-predicts the strangeness content, and since strange particles are typically excluded by tracking, the efficiency of Pythia is too high by up to

9. Note: We do not combine overlapping cell hits together.

6%, depending on p_T . The ALICE inclusive hadron R_{AA} analysis computed a correction to the Pythia tracking efficiency by re-weighting according to measured particle compositions (Section 2.3.2 of [84]). However, it has been shown that the strangeness content within jets is significantly smaller than the inclusive strangeness content, by a factor $\approx 5 - 10$. We therefore ignore this effect.

We assume that the combinatorial background has negligible impact from the pp event, and therefore we compute the event-by-event ρ_{ch} using only Pb–Pb tracks, and we apply the background scale factor obtained in Pb–Pb MB data.

We then perform jet finding on the hybrid event at detector-level, as well as the pp truth-level. We employ a geometrical matching procedure between the hybrid jets and particle-level jets: If an $R = 0.2$ accepted hybrid jet and an accepted probe jet are within $R < 0.25$, and they are both the closest jets to each other, then the jets are matched, and the response matrix is incremented at $(p_{T,det}^{jet}, p_{T,gen}^{jet})$. For $R = 0.4$ jets, we use $R < 0.45$. Note that the matching candidates consist of “accepted” jets, i.e. those satisfying the leading track requirement. This leading track requirement nullifies the need for further criteria such as a shared momentum fraction requirement in order to generate unique and accurate matches.

Jet reconstruction efficiency and jet matching efficiency

We must also compute the efficiency of successfully reconstructing accepted jets, known as the jet reconstruction efficiency. In particular, we would like to compute the efficiency to reconstruct true jets with a 5 GeV/c leading charged hadron bias. This quantity will be used to correct the unfolded spectrum for the fact that we fail to measure a certain fraction of jets. The jet reconstruction efficiency can be computed as

$$\varepsilon(p_{T,gen}^{jet}) = \frac{N_{\text{matched}}(p_{T,gen}^{jet})}{N_{\text{truth}}(p_{T,gen}^{jet})},$$

where N_{matched} is the number of accepted detector-level jets matched to accepted probe jets (where the probe jets are also required to contain a 5 GeV/c leading charged hadron) out of N_{truth} probe jets (also with the 5 GeV/c leading charged hadron requirement). Note that

this quantity does not explicitly include the bias of the 5 GeV/c leading charged hadron requirement, but only the probability to reconstruct an accepted jet given a truth-level jet with a 5 GeV/c leading charged hadron.

In fact, we desire not just the jet reconstruction efficiency, but also the false positive rate at which we measure accepted jets that have no matching probe. That is, we would like ε to account not just for cases where we fail to measure an accepted jet when a truth-level jet originated inside our acceptance, but also we would like to account for the fact that we may measure an accepted jet which did not originate from a truth-level jet in our acceptance, because the truth-level jet was generated slightly outside of our geometrical acceptance. For the case of p_T smearing, we neglect this at present, since the p_T resolution is $< 1\%$ at 5 GeV/c (and, if we removed the leading track requirement from the jet acceptance while matching, we would be plagued by uncorrelated Pb–Pb jets matching to the truth jet). For the case of geometrical acceptance contamination, we can reasonably assume that the jet response will describe these jets essentially as accurately as the in-acceptance jets, and they therefore pose no problem as long as we properly correct for them in the efficiency. We wish, then, to include the fake rate in our jet reconstruction efficiency correction, since these cases may occur in the measured data. Accordingly, then, we must carefully define the acceptances in the ratio above in order to account for possible geometrical acceptance contamination. For the denominator, the geometrical acceptance should be defined as the same as the acceptance used for the measurement, i.e. the EMCal fiducial acceptance. For the matched case, however, we should require the hybrid detector-level jet to be inside the EMCal fiducial acceptance, but we should allow the probe jet to have no restriction on acceptance. More precisely, we want that the fake-corrected number of matches should be equal to the efficiency-corrected number of true jets: $N_{matched} \times \varepsilon_{fake} = N_{truth} \times \varepsilon_{eff}$. If we follow the prescription outlined above, we therefore see that $\varepsilon \equiv \frac{\varepsilon_{eff}}{\varepsilon_{fake}}$, as intended. After the spectrum is unfolded, it must be corrected for this efficiency, $\varepsilon \left(p_{T,gen}^{jet} \right)$.

Note that in order for this quantity to be the jet reconstruction efficiency, we need that the jet matching efficiency is 100%. However, in the Pb–Pb embedding environment, it is difficult to achieve a matching efficiency of 100%, since some criteria need to be imposed to suppress combinatorial jets (in our case, the leading track requirement). Therefore, to

avoid artificially suppressing the jet reconstruction efficiency by the non-100% matching efficiency, and to avoid attempting a sub-percent-level understanding of the jet matching in the Pb–Pb embedding, we instead use the jet reconstruction efficiency as determined in the pp simulation alone (with 2% reduced tracking efficiency). The response still accounts for all of the Pb–Pb effects such as background smearing – but we independently evaluate the jet reconstruction efficiency in pp. In principle, there are several ways that the Pb–Pb jet reconstruction efficiency can be lower than the pp efficiency:

1. Geometrical acceptance migration: The Pb–Pb background can pull the jets in and out of acceptance. However, migration into the acceptance and out of the acceptance should occur at equal rates, and we therefore don’t expect a net difference from the pp geometrical acceptance migration.
2. The jet area cut. The effect of the area cut was examined by determining the jet reconstruction efficiency in the embedding case with and without the area cut, and measuring the difference. The area cut was found to have negligible impact in the reported p_T range, much less than 1%, for all R .
3. The leading track requirement: If we fail to reconstruct the embedded pp jet, but there is a background jet that passes the leading track requirement, we can get a fake efficiency. However, the potential presence of combinatorial jets in the embedding case is irrelevant for the jet reconstruction efficiency, since the efficiency is only applied at truth-level, after unfolding, so we are operating under the expectation that there are no combinatorial jets contributing in the unfolded spectrum in the ranges we will report. So the question of whether or not there are combinatorial jets that can give fake efficiency in the matching is immaterial, since we in fact wouldn’t want to count those in the efficiency anyway. Rather the efficiency is only a correction of the yield of the true unfolded jets.
4. The high- p_T track veto: If we fail to reconstruct the embedded pp jet, but there is a Pb–Pb 100 GeV/c track there, we can get a fake efficiency. This however is virtually impossible, and we can safely assume it is pp-like.




Figure 3.39: Left: Fine-binned response matrix for $R = 0.2$, over the full range of $p_{T,\text{det}}^{\text{jet}}$ and $p_{T,\text{gen}}^{\text{jet}}$. Right: Re-binned response matrix for $R = 0.2$, used in the unfolding, with selections imposed on $p_{T,\text{det}}^{\text{jet}}$ and $p_{T,\text{gen}}^{\text{jet}}$ as described in the text.

We therefore assume that the Pb–Pb jet reconstruction efficiency can be described by the pp jet reconstruction efficiency (with 2% reduced track efficiency), without further corrections. Note also that in the 2.76 TeV analysis, the pp jet reconstruction efficiency was also used, and implicitly relied on the same assumptions.

Response matrix details and kinematic efficiency

The response matrix is generated in a fine binning, with 1 GeV/c bin widths on both axes. This is then re-binned into a more coarse binning to be used in the actual unfolding: 5 GeV/c bin width for $p_{T,\text{det}}^{\text{jet}}$, 10 GeV/c width for $p_{T,\text{gen}}^{\text{jet}}$. Figure 3.39 shows the fine-binned and re-binned response matrices for $R = 0.2$. Corresponding plots for $R = 0.4$ are provided in Appendix B. It is important that the number of truth-level bins is sufficiently larger than the number of detector-level bins, in order that the input data can meaningfully constrain the unfolded result.

In general, we seek to unfold the measured spectrum over a fixed window of $p_T^{\text{det}} \in [p_{T,\text{min}}^{\text{det}}, p_{T,\text{max}}^{\text{det}}]$ that we expect to be above the combinatorial background. In principle, the response matrix should unfold combinatorial jets to low p_T , but we want to minimize reliance on this. Accordingly, the response matrix used in the unfolding procedure should extend only over that same range in $p_{T,\text{det}}^{\text{jet}}$, and over a large range of $p_{T,\text{gen}}^{\text{jet}}$ so as to include all truth-level bins that yield any contribution to the $p_{T,\text{det}}^{\text{jet}}$ range (although the high- p_T

edge doesn't really matter, since there are very few such jets populating the measured $p_{T,\text{det}}^{\text{jet}}$ range). For our purposes, we use $p_T^{\text{truth}} \in [5, 250 \text{ GeV}/c]$. Note that one should not unfold using the full $p_{T,\text{det}}^{\text{jet}}$ range (of both the response matrix and the measured spectrum), since we are not interested in the region of combinatorial jets, and this can de-stabilize the unfolding process. On the other hand, one does not want to set too high a $p_{T,\text{det}}^{\text{jet}}$ threshold, in order to keep the kinematic efficiency high. For $R = 0.2$, we use $p_{T,\text{det}}^{\text{jet}} > 20 \text{ GeV}/c$. For $R = 0.4$, we use $p_{T,\text{det}}^{\text{jet}} > 35 \text{ GeV}/c$. Each of these corresponds to $\approx 2 - 3 \times \sigma_{\delta_{p_T}}$, which, in combination with the leading charged hadron requirement, is expected to result in a largely background-free region. Any residual combinatorial jets will still be unfolded to low- p_T by the response matrix. For the upper $p_{T,\text{det}}^{\text{jet}}$ limit, we take $p_{T,\text{det}}^{\text{jet}} < 120 \text{ GeV}/c$, where the $p_T^{\text{track}} < 100 \text{ GeV}/c$ requirement is seen to have negligible bias.

Truncating the response matrix in $p_{T,\text{det}}^{\text{jet}}$ loses information, in particular the fraction of truth-level jets that migrate outside of the measured det-level window. That is, given a truth-level jet with $p_{T,\text{gen}}^{\text{jet}}$ that is successfully reconstructed at detector-level, and a fixed measured window $p_{T,\text{det}}^{\text{jet}} \in [p_{T,\text{min}}, p_{T,\text{max}}]$, there is only a certain probability that the truth-level jet will be reconstructed within the measured window. This probability is described by the kinematic efficiency, $\varepsilon_{\text{kin}}(p_{T,\text{gen}}^{\text{jet}})$. The result must be corrected for this effect, since kinematically we only have the chance to measure a certain fraction of each $p_{T,\text{gen}}^{\text{jet}}$ yield. Implicit to this correction is the assumption that those jets that fall outside our kinematic range are not quenched differently than those that do fall in our kinematic range. Figure 3.40 shows the kinematic efficiency for $R = 0.2, 0.4$, which is obtained as the ratio of the $p_{T,\text{gen}}^{\text{jet}}$ projection of the response matrix after truncating in $p_{T,\text{det}}^{\text{jet}}$ to the $p_{T,\text{gen}}^{\text{jet}}$ projection before truncating.

The response matrix is normalized so as to preserve the number of jets upon unfolding. That is, each truth-level jet should map with probability 1 to a detector-level jet. To ensure this, given the un-normalized response matrix we project the det-level yield for each $p_{T,\text{gen}}^{\text{jet}}$, and normalize that yield to 1. That is, we compute $\text{Proj}(p_{T,\text{gen}}^{\text{jet}})$ for each $p_{T,\text{gen}}^{\text{jet}}$, and normalize each bin $(p_{T,\text{det}}^{\text{jet}}, p_{T,\text{gen}}^{\text{jet}})$ by $\frac{1}{\text{Proj}(p_T^{\text{truth}})}$. If one then plots the truth axis projection, its amplitude is uniformly at 1. Note that our unfolding approach only works correctly when a jet is smeared (i.e. the probe jet overlapping with a “background” jet),




Figure 3.40: Left: Kinematic efficiency for $R = 0.2$ jets for the range $p_{T,\text{det}}^{\text{jet}} \in [20, 120]$ GeV/c. Right: Kinematic efficiency for $R = 0.4$ jets for the range $p_{T,\text{det}}^{\text{jet}} \in [35, 120]$ GeV/c.

but it does not work correctly when two “real” jets overlap (which sometimes occurs in the measured data), since the unfolding procedure conserves the number of jets. Since this is the case, we want to exclude the possibility of unfolding a real-real jet overlap to the low- p_T jet (e.g. to unfold a $p_{T,\text{det}}^{\text{jet}} = 100$ GeV/c to $p_{T,\text{gen}}^{\text{jet}} = 5$ GeV/c). However, this possibility only contributes significantly for very small $p_{T,\text{gen}}^{\text{jet}}$, and we already impose $p_{T,\text{gen}}^{\text{jet}} > 5$ GeV/c by the leading hadron requirement, so we can safely neglect this possibility.

3.4.2 Performing the unfolding

With the response matrix generated in the embedding procedure, we wish to unfold the measured Pb-Pb jet spectrum in order to produce a truth-level quenched spectrum.

The unfolded result is produced over the full unfolded range $p_T^{\text{truth}} \in [p_{T,\text{min}}^{\text{gen}}, p_{T,\text{max}}^{\text{gen}}]$. However, we intend to report the result only over a limited range of p_T over which we believe the input data meaningfully constrains the unfolded result. That is, we want the kinematic efficiency to be reasonably large, and we want to be confident we are in a region unaffected by combinatorial jets. This is mainly important at the low- p_T edge, where the kinematic efficiency becomes small, and the MC correction is correspondingly large (and sensitive to feed-in due to the steeply falling spectrum). We expect this to be approximately at the average truth-level p_T corresponding to the minimum $p_{T,\text{det}}^{\text{jet}}$. The sensitivity of the various selected ranges will be addressed in Section 3.5.

We employ the SVD unfolding algorithm [82], in which the inverse of the response matrix

is approximated by an expansion with weight d_k for the k -th term. The parameter k is used as the regularization parameter to forbid high-frequency variations in the unfolded result. One should select a value of k such that d_k is approximately 1, but not so large that the correlation coefficients of nearby bins are anti-correlated with each other, which indicates the presence of unphysical fluctuations. The unfolded result should not actually converge with larger k – rather, increasing k too much will make the distribution spuriously oscillate. Note that the SVD algorithm requires that a prior input spectrum be provided.

We perform the unfolding using the RooUnfold package [85]. In RooUnfold, the response matrix should be provided un-normalized (it will be normalized internally), and the RooUnfoldResponse object should be provided with the full $p_{T,\text{gen}}^{\text{jet}}$ spectrum before truncating the response matrix in order to correct for the kinematic efficiency. The measured input binning for $R = 0.2$ is 5 GeV/c intervals from 20 GeV/c to 120 GeV/c, and for $R = 0.4$ is 5 GeV/c intervals from 35 GeV/c to 120 GeV/c. The unfolded output binning for both radii is [5, 10, 20, 30, 40, 50, 60, 70, 80, 100, 120, 140, 190, 250]. These choices are selected in order to increase the number of input constraints, while balancing with the need to keep statistical fluctuations not too large.

The SVD unfolding algorithm also performs statistical error propagation. Note that the statistical uncertainties may therefore be partially correlated. We use the default error propagation in RooUnfold, which computes the statistical errors according to the square root of the diagonal elements of the covariance matrix. Systematic uncertainties on the unfolding procedure will be discussed in Section 3.5.

As an illustration of the unfolding procedure we plot here the unfolding plots for $R = 0.2$. The corresponding plots for $R = 0.4$ are provided in Appendix B. Figure 3.41 shows the unfolded result as a function of k and the d-vector, which suggests that $k = 4$ is a reasonable solution. Figure 3.42 shows the correlation coefficients, which measure the degree of correlation between the bins. We want that when k approaches our selected value, far away bins are not correlated, for reasons of stability and smoothness.

Once we obtain a solution, we must verify in a data-driven way that the solution is robust. We do this by a refolding test and a “self-closure” test. The refolding test consists of the following: Generate a response matrix (from MC data sample 1) and unfold the




Figure 3.41: Left: Unfolded result for $R = 0.2$ jets as a function of k . Right: D-vector for $R = 0.2$ jets.




Figure 3.42: Pearson correlation coefficients for $R = 0.2$ jets for $k = 2, 3, 4, 5$ from left to right.

measured distribution (using all data for the measured distribution), then apply a response matrix (from MC data sample 2) to the unfolded result, and compare the re-folded solution to the measured distribution. The “self-closure” test consists of the following: From the full embedded sample, take the matched detector-level jet spectrum, and smear each data point with a Gaussian according to the statistical uncertainties on the measured Pb–Pb data. Then, unfold this spectrum using the response matrix, and compare the result to the truth-level Pythia jet spectrum.

Figure 3.43 shows the results of the refolding test and self-closure test. The analogous plots for $R = 0.4$ are shown in Appendix B. Note that this test does not test the validity of the response matrix (this is determined rather by the accuracy of the simulation and the embedding procedure), or the effect of combinatorial jets. Section 3.5 addresses a variety of systematic uncertainties associated with this unfolding correction.




Figure 3.43: Left: Re-folding test for $R = 0.2$ jets. Right: “Self-closure” test comparing Pb-Pb embedded det-level data (statistically smeared according to uncertainties in real data) with Pythia truth. Both are performed for the main result $k = 4$.

3.4.3 Thermal closure test

The re-folding and self-closure tests check the mathematical consistency of the unfolding framework, but it does not test whether the unfolded solution is physically correct. In particular, the rejection of combinatorial jets in the final result must be verified. To do so, we perform a test using Pythia events in a simulated thermal background. We perform the entire analysis on the “hybrid” jets clustered from the combination of Pythia detector-level particles and the thermal background particles: Construct the “hybrid” detector-level jet spectrum, build the response matrix, and unfold the “hybrid” jets – and compare to the truth-level Pythia spectrum. Since the background does not have any jet component, this test is able to verify whether the analysis procedure indeed recovers the correct solution. That is, the test ensures that if I have a jet signal and a background, my analysis procedure will correctly produce the jet spectrum, and not be contaminated by the background. In principle, one could also use a simulated background with a known jet component – but it cannot be Pb–Pb data itself, since we do not know a priori which jets are “true” jets, and which are combinatorial jets.¹⁰

10. This point has a tendency to be confusing, so I elaborate a bit more: Embedding Pythia into Pb–Pb background and performing a folding/unfolding test doesn’t help achieve our goal, since this is just a technical check of the unfolding procedure: We give a truth spectrum, fold it, and unfold it. That spectrum could be

We follow a similar analysis procedure to that in [74], embedding Pythia events into a simple thermal background. Note that we only consider a single hard-scattering here, which is a more stringent test than if we had a realistic number of hard-scatterings, since the contribution of the hard spectrum compared to the background is smaller in our case, which should enhance susceptibility to combinatorial jets. For a more rigorous approach, one should use a model that combines a known jet spectrum with an accurate heavy-ion background (i.e. having N_{coll} Pythia hard scatterings, but in a way such that the hard spectrum does not overpopulate the low- p_T background distribution).

We implement the following procedure: Generate N particles with p_T taken from a Gamma distribution:

$$f_{\Gamma}(p_T; \alpha, \beta) \sim p_T^{\alpha-1} e^{-p_T/\beta}$$

We model N as a Gaussian, and select $\alpha = 2$ which is typical for fitting p_T spectra. We then choose the free parameters $\bar{N}, \sigma_N, \beta$ in order to roughly fit the δ_{p_T} distribution in 0-10% Pb-Pb data (and to be very roughly compatible with $N, \langle p_T \rangle$ from data). We then perform jet-finding on the hybrid event to get a “detector-level” jet spectrum, and to fill the response matrix: $\left(p_{T,\text{truth}}^{\text{pythia}}, p_{T,\text{det}}^{\text{pythia+thermal}} \right)$.

For the case $R = 0.2$, with $p_{T,\text{lead,ch}} = 5 \text{ GeV}/c$, we choose $\bar{N} = 3500, \sigma_N = 500, \beta = 0.4$. Figure 3.44 shows the δ_{p_T} distribution in the thermal model, as well as the JER of the hybrid jets, both of which are seen to agree reasonably well with 0-10% Pb-Pb data. The ρ distribution is not expected to agree very well with Pb-Pb data, but does not pose a problem, since the average background is subtracted as usual. Figure 3.46 shows the result of the closure test: the ratio of the unfolded hybrid jet spectrum to the Pythia truth spectrum. We see that down to approximately $p_T = 30 \text{ GeV}/c$, the ratio is within a few percent of unity, supporting the claim that the analysis procedure is correct.

anything. This has nothing to do with the question of whether there are combinatorial jets in our measured spectrum. To test for combinatorial jets in this approach, we would have to embed Pythia into Pb-Pb and then form a det-level spectrum, and unfold that with our previously obtained response matrix. However, this is not possible since there would be many Pb-Pb jets in this sample, so the det-level spectrum would have many jets that we don’t have a handle on. Rather, we must know exactly the background and hard distributions.

We repeat the procedure for the case $R = 0.4$, with $p_{T,\text{lead,ch}} = 5 \text{ GeV}/c$, we choose $\overline{N} = 4000, \sigma_N = 500, \beta = 0.5$. Note that the δ_{p_T} distribution cannot be fit as a function of R with this simple model, so we must re-tune the parameters [80]. Figure 3.45 shows the δ_{p_T} distribution in the thermal model, as well as the JER of the hybrid jets, both of which are seen to agree reasonably well with 0-10% Pb-Pb data. Figure 3.47 (left) shows the result of the closure test, which fails to successfully unfold. We instead try to repeat the procedure for $R = 0.4$ with an increased leading track threshold, $p_{T,\text{lead,ch}} = 7 \text{ GeV}/c$. Figure 3.47 (right) shows the result of the closure test, which is seen to agree with unity down to at least 60 GeV/c. We therefore require a 7 GeV/c leading track for $R = 0.4$ jets.




Figure 3.44: Left: Distribution of δ_{p_T} in the thermal closure test for $R = 0.2$. Right: JER of the hybrid jets in the thermal closure test for $R = 0.2$.




Figure 3.45: Left: Distribution of δ_{p_T} in the thermal closure test for $R = 0.4$. Right: JER of the hybrid jets in the thermal closure test for $R = 0.4$.




Figure 3.46: Closure test for $R = 0.2$, with $p_{T,\text{lead,ch}} = 5$ GeV/c: Ratio of the unfolded hybrid jet spectrum to the pythia truth spectrum.




Figure 3.47: Left: Closure test for $R = 0.4$, with $p_{T,\text{lead,ch}} = 5$ GeV/c: Ratio of the unfolded hybrid jet spectrum to the Pythia truth spectrum. Right: Closure test for $R = 0.4$, with $p_{T,\text{lead,ch}} = 7$ GeV/c: Ratio of the unfolded hybrid jet spectrum to the Pythia truth spectrum.

3.5 Systematics

Following [74], we categorize two classes of systematic uncertainties: correlated uncertainties and shape uncertainties. Correlated uncertainties encompass detector effects such as uncertainty on the tracking efficiency and uncertainty on the EMCal response, which are positively correlated among all p_T^{jet} bins. Shape uncertainties refer to systematic unfolding uncertainties, which alter the shape of the final p_T^{jet} spectrum. The dominant systematic uncertainties in this analysis are the uncertainty in the tracking efficiency and the systematic uncertainty in the unfolding procedure. Note that in general we are interested in computing the uncertainties on the jet spectrum, not the uncertainty on the jet p_T scale.

3.5.1 Correlated uncertainties

The dominant correlated uncertainty is the uncertainty on the tracking efficiency, since correcting for unmeasured tracks is a major effect of the unfolding procedure. It is estimated that for hybrid tracks, the uncertainty on the tracking efficiency is approximately 4% [79]. This number is attributed to two contributions: variation in the track selection parameters, and variation in the ITS-TPC matching requirements. In order to assign a systematic to the final result, we construct a response matrix using the same techniques as for the final result except that we randomly reject an additional 4% of tracks in jet finding (that is, in addition to the 2% rejection used in the main result). We also compute the jet reconstruction efficiency with this extra 4% suppression applied. This response matrix is then used to unfold the same measured result as used in the final result. This result is corrected for the jet reconstruction efficiency, and compared to the main result, with the differences in each bin taken as the uncertainty, shown in Fig. 3.48 (left).

We include also a systematic uncertainty associated with the choice of jet matching procedure, varying from pure geometrical matching to an MC-fraction based approach. The shared momentum fraction requirement ensures that the matched jet contains “enough” of the MC jet. To enforce a shared momentum fraction requirement, however, we must compare det-level momentum quantities. Therefore, we use a scheme of jet matching where we match the “combined” jet geometrically to the nearest pp-embedded detector-level, and

require that the jets be within $\Delta R < R$ and that the combined jet must contain at least 50% of the tracks of the pp detector-level charged jet, as measured by p_T (a charged jet is used just for simplicity). Additionally, the pp detector-level charged jet is matched to its corresponding pp truth-level jet. This then defines a matching between the combined jet and a truth-level jet, and is used to fill the response matrix. This is shown in Fig. 3.48 (right), and gives an uncertainty of approximately 2%, except increasing to 6% in the case of $R = 0.4$.

Note that there is no need to compute an additional uncertainty due to the background fluctuations, since we do not explicitly use δ_{p_T} in the unfolding procedure. Moreover, we need not perform a correction in our response matrix for flow effects biasing the background fluctuations due to the leading hadron requirement, since our embedding procedure naturally includes these effects. We therefore do not included a systematic uncertainty associated with the effect.

There are several additional correlated uncertainties, which slightly increase the total correlated uncertainty. We do not perform studies to evaluate these individually for the present analysis, but rather rely on previous studies whose results should hold here as well (see [74] for example). The uncertainty on the tracking p_T resolution is assumed to be




Figure 3.48: Correlated uncertainty variations for $R = 0.2$ jets. Left: Variation of the tracking efficiency by 4% compared to the main result. Right: Variation of the jet matching technique from geometrical matching to MC-fraction based matching.

	Relative uncertainty (%) for $p_T \in [A, B]$ GeV/c						
	40-50	50-60	60-70	70-80	80-100	100-120	120-140
$R = 0.2, 5$ GeV/c							
Tracking efficiency	5.2	5.8	6.5	6.7	7.3	7.5	7.7
Jet matching	2	2	2	2	2	2	2
Track p_T resolution	1	1	1	1	1	1	1
EMCal energy response	4.4	4.4	4.4	4.4	4.4	4.4	4.4
EMCal hadronic response	0.3	0.7	1.1	1.3	1.5	1.8	2.1
Total corr. uncertainty	7.2	7.6	8.2	8.4	8.9	9.2	9.4
$R = 0.2, 7$ GeV/c							
Tracking efficiency	4.5	5.0	5.6	6.1	6.8	7.5	8.0
Jet matching	2	2	2	2	2	2	2
Track p_T resolution	1	1	1	1	1	1	1
EMCal energy response	4.4	4.4	4.4	4.4	4.4	4.4	4.4
EMCal hadronic response	0.2	0.2	0.2	1.0	2.0	3.0	3.9
Total corr. uncertainty	6.7	7.0	7.5	7.9	8.6	9.5	10.2
$R = 0.4, 7$ GeV/c							
Tracking efficiency			13.7	12.6	10.7	9.1	7.7
Jet matching			6	2	2	2	2
Track p_T resolution			1	1	1	1	1
EMCal energy response			4.4	4.4	4.4	4.4	4.4
EMCal hadronic response			1.9	3.2	4.5	5.9	7.4
Total corr. uncertainty			15.7	13.9	12.6	11.9	11.8

Table 3.3: Correlated uncertainties on the jet spectrum.

approximately 1%. The MC description of the EMCal energy response has several sources of systematic uncertainty: the energy calibration scale, the cluster energy resolution, and the non-linearity correction. We assume an uncertainty of 4.4% as used in [74]. We also rely on the MC to describe well the pp-like EMCal hadronic response, and use the systematics found in the 5.02 TeV pp full jet analysis.

Table 3.3 shows the contributions of the various correlated uncertainties. These uncertainties are expected to be largely independent, so we sum them in quadrature.

3.5.2 Shape uncertainties

We perform several systematic variations on the unfolding procedure to assign a shape uncertainty arising from the unfolding regularization procedure:

- Variation of the unfolding algorithm

- Variation of the regularization parameter
- Variation of the prior
- Variation of the input range

We unfold with a Bayes-inspired iterative unfolding algorithm [86] in order to assign a systematic associated with the choice of unfolding algorithm, as shown in Fig. 3.49. The Bayes approach uses a non-linear iterative procedure which converges to the solution – it generates a truth, tries folding it with the response matrix, and evaluates how close it is to the measured distribution – then updates and repeats. The number of iterations determines the regularization, with fewer iterations imposing a smoother solution. One should choose the number of iterations (nIter) to be a small value for which solution converges.

In the SVD unfolding, we vary the regularization parameter k one unit above and below the nominal solution, shown in Fig. 3.49, and take the maximum difference between either variation and the main result as the systematic attributed to the regularization parameter.

The SVD algorithm requires a prior distribution as input, which for the main result we use as the projection of the response matrix onto the truth axis (before normalization). As a systematic, we vary this input prior by scaling the main prior by $p_T^{\pm 0.5}$, and take the maximum difference between either variation and the main result as the systematic for this effect.

Additionally, we assign a systematic corresponding to uncertainty about the influence of combinatorial jets on the unfolding process. We select an input $p_{T,\text{det}}^{\text{jet}}$ range which excludes most combinatorial jets, in combination with the leading charged hadron requirement. However, there may still be some residual combinatorial jets present in the input measured spectrum, particularly for large jet R . This is not necessarily a problem, since the combinatorial jets can be unfolded to low p_T by the response matrix, but we nevertheless want to minimize such contributions. To assign a systematic uncertainty associated with this effect, we vary the measured input range ± 5 GeV/c around the nominal value for each R , and take the maximum variation per bin as the systematic, shown in Fig. 3.49.

The total shape uncertainty is chosen as the standard deviation of the systematics due to each of the variations, since they each comprise independent measurements of the same

underlying systematic uncertainty in the regularization. That is, the total shape uncertainty for each bin is $\sqrt{\frac{\sum_{i=1}^3 \sigma_i^2}{4}}$, where σ_i is the systematic due to a single one of the four variations described above.¹¹ Table 3.4 shows the systematic uncertainties in each bin for each R .




Figure 3.49: Shape uncertainty variations for $R = 0.2$ jets. Top left: Variation of the regularization algorithm. Top right: Variation of the SVD regularization parameter. Bottom left: Variation of the prior. Bottom right: Variation of input range.

11. Note that we divide by 4, not 5, since the main measurement is assumed to be the mean.

	Relative uncertainty (%) for $p_T \in [A, B]$ GeV/c						
	40-50	50-60	60-70	70-80	80-100	100-120	120-140
$R = 0.2, 5$ GeV/c							
Unfolding method	8.6	2.1	3.2	5.1	4.7	4.5	5.6
Reg. parameter	4.4	5.2	4.9	3.8	1.9	2.2	2.4
Prior	1.7	1.5	1.0	1.8	0.5	2.9	5.5
Input p_T range	0.5	0.6	0.5	0.3	0.8	1.3	1.7
Total shape uncertainty	4.9	2.9	3.0	3.3	2.6	3.0	4.2
$R = 0.2, 7$ GeV/c							
Unfolding method	7.1	1.3	3.2	4.7	4.3	4.6	7.0
Reg. parameter	3.5	4.2	3.9	2.9	1.0	1.0	2.2
Prior	0.6	0.4	0.4	2.0	0.8	2.4	4.9
Input p_T range	0.5	0.5	0.7	0.2	0.4	0.9	1.2
Total shape uncertainty	4.0	2.2	2.6	2.9	2.3	2.7	4.5
$R = 0.4, 7$ GeV/c							
Unfolding method			35.5	30.3	17.0	5.6	25.7
Reg. parameter			6.3	8.0	2.3	7.4	14.1
Prior			0.7	1.1	1.2	0.5	1.7
Input p_T			1.5	2.6	3.5	2.8	1.1
Total shape uncertainty			18.0	15.7	8.8	4.9	14.7

Table 3.4: Shape uncertainties on the jet spectrum.

3.5.3 Uncertainties on the jet cross-section ratio

We compute systematics on the jet cross-section ratio by making the same variations as above on both spectra simultaneously, and compare the varied jet cross-section ratio to the main result. Table 3.5 shows the systematics considered. For the correlated uncertainties, we only explicitly vary the tracking efficiency uncertainty, since other variations sometimes de-stabilize the unfolding and result in artificially inflated uncertainties. We then assume that the other correlated uncertainties on the spectra scale approximately as the tracking efficiency. We denote these as “other” in Table 3.5.

Note that in simultaneously varying the numerator and denominator by the systematic variations described in the previous section, we in fact over-estimate the error. That is, given a quantity x where we choose a systematic variation σ_x , we perform the variation on $\frac{1}{x}$ as $\frac{1}{x-\sigma_x}$, whereas in principle we should take $\sigma_{1/x}$. This overestimates the error by a factor

$$\frac{\frac{1}{x - \sigma_x}}{\frac{1}{x} - \sigma_{1/x}} = \frac{1}{(1 - \sigma_x/x)^2},$$

which can result in the “cancelled” uncertainties being larger than those treated independently when the uncertainties become large. We therefore take the minimum of the conservative variation with cancellation, and treating the uncertainties independently.

It is important to note that the statistical uncertainties are partially correlated, due to error propagation through the unfolding procedure. This likely results in a conservative statistical uncertainty estimation, since there may be significant cancellation between the two radii that is not taken into account. Additionally, we do not use statistically independent samples to form the ratio, since we do not have the precision to do so – and so the numerator and denominator are statistically correlated with each other, which we also do not take into account.

	Relative uncertainty (%) for $p_T \in [A, B]$ GeV/c				
	60-70	70-80	80-100	100-120	120-140
$R = 0.2/R = 0.4$					
Tracking efficiency	7.1	5.8	3.5	1.4	0.3
Other	4.2	3.7	2.4	1.0	0.2
Total corr. uncertainty	8.2	6.9	4.2	1.7	0.4
Unfolding method	35.6	30.7	17.5	7.2	26.6
Regularization parameter	5.7	7.3	4.2	1.8	4.3
Prior	0.6	0.8	0.9	2.8	5.4
Input p_T range	1.2	3.0	3.4	2.1	0.2
Total shape uncertainty	18.0	15.9	9.2	4.1	13.7

Table 3.5: Systematic uncertainties on the jet cross-section ratio.

3.6 Results

In a separate analysis, the fully corrected pp full jet cross-section was measured in order to serve as a reference for R_{AA} .¹² We report the results of this analysis below for $R = 0.2, 0.4$, as well as a comparison to the NLO event generator POWHEG with Pythia8 showering as a pp reference. In pp, the jets are reported for the range $p_T^{\text{jet}} \in [20, 140]$ GeV/c.

We report fully corrected Pb–Pb jet spectra for $R = 0.2, 0.4$ in 0-10% centrality, as well as R_{AA} using the measured pp reference. In addition, we report the Pb–Pb jet cross-section ratio $R = 0.2/R = 0.4$. The $R = 0.2$ jets are reported for the range $p_T^{\text{jet}} \in [40, 140]$ GeV/c. The $R = 0.4$ jets are reported for the range $p_T^{\text{jet}} \in [60, 140]$ GeV/c. The reported ranges are selected based on being sufficiently far above the combinatorial background, as well as having high kinematic efficiency. Note that the reported range extends higher than the measured range because the kinematic efficiency remains high at larger p_T due to the JES shift.

All spectra are reported for jets that satisfy a leading charged hadron requirement, either 5 GeV/c or 7 GeV/c, which is specified on each plot.¹³ All results are unfolded for detector and background effects, and are reported at the hadron-level.

3.6.1 Jet spectra

pp

The pp jet cross-sections are reported differentially in p_T and η as:

$$\frac{d^2\sigma_{jet}}{dp_T d\eta}.$$

Since we measure experimentally the yield $\frac{d^2N}{dp_T d\eta}$, we must divide by the integrated luminosity in order to get the cross-section:

12. The pp analysis was led by Dr. Eliane Epple (Yale University) in collaboration with the author, using the analysis strategies and infrastructure developed for the Pb–Pb analysis.

13. We do not attempt to correct to a fully inclusive spectrum, in order to avoid model-dependence.

$$\frac{d^2\sigma_{jet}}{dp_T d\eta} = \frac{1}{\mathcal{L}} \frac{d^2N}{dp_T d\eta}.$$

In particular, we measure the yield for events with $|z_{\text{vertex}}| < 10$ cm. We therefore want to divide by $\mathcal{L}_{10}$, the luminosity delivered within this $|z_{\text{vertex}}|$ range. So we form:

$$\frac{d^2\sigma_{jet}}{dp_T d\eta} = \frac{\sigma_{\text{MB}}^{\text{visible}}}{N_{\text{MB}}^{10}} \frac{d^2N^{10}}{dp_T d\eta}.$$

Note that the cross-section is independent of the $|z_{\text{vertex}}|$ requirement. Since the differential yield is taken within 10 cm, and the number of events is taken within 10 cm, this correctly gives the total jet cross-section (implicitly assuming that the jet cross-section does not depend at which $|z_{\text{vertex}}|$ the collision occurred).

Further, to obtain the value N_{MB}^{10} , we must correct for the vertex efficiency, since the number of events we accept, N_{acc} , does not account for events that fired the MB trigger, but failed to reconstruct a vertex. That is,

$$N_{\text{MB}}^{10} = N_{\text{acc}} \times \frac{N_{\text{MB}}^{10, \text{tot}}}{N_{\text{MB}}^{10, \text{w/vertex}}}.$$

We further assume that the vertex efficiency is independent of $|z_{\text{vertex}}|$, and therefore we use:

$$N_{\text{MB}}^{10} = N_{\text{acc}} \times \frac{N_{\text{MB}}^{\text{tot}}}{N_{\text{MB}}^{\text{w/vertex}}}.$$

For the present analysis, we obtain the value

$$\frac{N_{\text{MB}}^{\text{tot}}}{N_{\text{MB}}^{\text{w/vertex}}} = 1.054.$$

For ALICE Run 2 at $\sqrt{s_{\text{NN}}} = 5.02$ TeV, the MB trigger is defined as V0AND, requiring a coincidence hit in V0A and V0C. The cross-section for this trigger has been determined to be [87]:

$$\sigma_{\text{MB}}^{\text{visible}} = 51.2 \pm 1.2 \text{ mb (sys)}$$

The uncertainty on the luminosity is determined by this uncertainty (to be taken as 5% for 2017 pp data until the internal ALICE measurement is finalized).

The jet spectra are unfolded and corrected for the kinematic efficiency and jet reconstruction efficiency. The results are corrected for the partial ϕ acceptance of the EMCal by multiplying by the factor:

$$\frac{2\pi}{\Delta\phi_{EMCal} - 2R}.$$

And since the results are reported differentially in η , the spectra are divided by the fiducial η window:

$$\Delta\eta_{EMCal} - 2R.$$

Figure 3.50 shows the unfolded pp full jet spectrum for $R = 0.2$ jets. Figure 3.51 shows the unfolded pp full jet spectrum for $R = 0.4$ jets. Note that a 5 GeV/c leading track




Figure 3.50: Unfolded pp full jet spectrum for $R = 0.2$ jets, along with POWHEG+Pythia reference.




Figure 3.51: Unfolded pp full jet spectrum for $R = 0.4$ jets, along with POWHEG+Pythia reference.

requirement is used here for both $R = 0.2$ and $R = 0.4$.

The POWHEG+Pythia jet cross-section $\frac{d^2 \sigma_{jet}^{pp}}{dp_T d\eta}$ is also plotted for comparison. The POWHEG reference is produced by POWHEG-BOX-V2 at $\sqrt{s_{NN}} = 5.020$ TeV via the jet pair production process [88–90].¹⁴ PDF set CT14nlo is used, along with the settings `bornktmin=1` and `bornsupfact=70`. Pythia 8.223 tune ATLAS-A14 is used for fragmentation; merging with Pythia is done as in [91]. The same set of primary particles is used as described in Section 3.4 [83]. Two theoretical uncertainties are computed for this reference spectrum, both in regard to the POWHEG event generation: PDF uncertainty, computed as in [91], and scale uncertainty, which is computed by varying the renormalization and factorization scales. The total theoretical uncertainty on the cross-section is obtained by adding these two contributions in quadrature. Note that POWHEG+Pythia was shown to describe reasonably well preliminary results of the ALICE pp charged jet cross-section at $\sqrt{s_{NN}} = 5.02$ TeV, and the ATLAS pp cross-section at $\sqrt{s_{NN}} = 7$ TeV [92].

¹⁴. Thanks to Ritsuya Hosokawa (University of Tsukuba) for producing the POWHEG+Pythia spectra.

Note that the POWHEG+Pythia jets also contain a 5 GeV/c leading charged track requirement. The ratio of the reference spectrum with and without the 5 GeV/c leading charged hadron requirement is shown in Fig. 3.52.




Figure 3.52: Ratio of the POWHEG+Pythia pp jet cross-section with and without the 5 GeV/c leading charged particle requirement. Only statistical errors are plotted, but are smaller than the markers.

Pb–Pb

The Pb–Pb jet spectra are reported differentially in p_T and η as:

$$\frac{1}{\langle T_{AA} \rangle} \frac{1}{N_{event}} \frac{d^2 N_{jet}^{AA}}{dp_T d\eta},$$

where $\langle T_{AA} \rangle$ is the average nuclear thickness:

$$\langle T_{AA} \rangle = \frac{\langle N_{coll} \rangle}{\sigma_{inel}^{NN}},$$

computed in a Glauber model. For 0-10% Pb–Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV, ALICE computes preliminary values of

$$\langle T_{AA} \rangle = 23.4 \pm 0.78 \text{ (sys) mb}^{-1},$$

$$\langle N_{coll} \rangle = 1583 \pm 38 \text{ (sys)}.$$

The jet spectra are unfolded and corrected for the kinematic efficiency and jet reconstruction efficiency. The results are corrected for the partial ϕ acceptance of the EMCal by multiplying by the factor:

$$\frac{2\pi}{\Delta\phi_{EMCal} - 2R}.$$

And since the results are reported differentially in η , the spectra are divided by the fiducial η window:

$$\Delta\eta_{EMCal} - 2R.$$

Note that we assume the MB trigger efficiency for 0-10% events is 100%.

Figure 3.53 shows the unfolded full jet spectra for pp and Pb–Pb for $R = 0.2$ jets. Figure 3.54 shows the unfolded full jet spectrum for pp and Pb–Pb for $R = 0.4$ jets. Note that a leading track bias of 5 GeV/c is required for the $R = 0.2$ spectra (both pp and Pb–Pb), while a 7 GeV/c bias is required for the $R = 0.4$ spectra (both pp and Pb–Pb).




Figure 3.53: Unfolded full jet spectrum for $R = 0.2$ jets in pp and Pb-Pb.




Figure 3.54: Unfolded full jet spectrum for $R = 0.4$ jets in pp and Pb-Pb.

3.6.2 Jet R_{AA}

The jet R_{AA} is reported as:

$$R_{AA} = \frac{\frac{1}{\langle T_{AA} \rangle} \frac{1}{N_{event}} \frac{d^2 N}{dp_T d\eta} \Big|_{AA}}{\frac{d^2 \sigma}{dp_T d\eta} \Big|_{pp}},$$

namely the ratio of the Pb–Pb and pp spectra plotted in Figs. 3.53 and 3.54. Since the measured Pb–Pb spectra only report jets satisfying the leading charged hadron requirement, it is most natural to apply the same requirement for the pp reference.

Figure 3.55 shows the unfolded full jet R_{AA} for $R = 0.2$ jets. Figure 3.56 shows the unfolded full jet R_{AA} for $R = 0.4$ jets. The reference spectrum is the measured pp jet spectrum described above in Section 3.6.1. The uncertainties in the Pb–Pb and pp spectra are combined in quadrature, including the pp luminosity uncertainty and Pb–Pb $\langle T_{AA} \rangle$ uncertainty.




Figure 3.55: Jet R_{AA} for $R = 0.2$ full jets. The combined $\langle T_{AA} \rangle$ uncertainty and pp luminosity uncertainty of 6% is shown as a band on the dashed line at $R_{AA} = 1$.




Figure 3.56: Jet R_{AA} for $R = 0.4$ full jets. The combined $\langle T_{AA} \rangle$ uncertainty and pp luminosity uncertainty of 6% is shown as a band on the dashed line at $R_{AA} = 1$.

3.6.3 Jet cross-section ratio

Figure 3.57 shows the pp jet cross section ratio $\sigma_{R=0.2}/\sigma_{R=0.4}$. Figure 3.58 shows the Pb–Pb jet cross section ratio $\sigma_{R=0.2}/\sigma_{R=0.4}$. Figure 3.59 shows the pp and Pb–Pb jet cross section ratios $\sigma_{R=0.2}/\sigma_{R=0.4}$ plotted together. Note that for the pp jet cross-section ratios, we use a 5 GeV/c leading track bias for both radii, whereas for the Pb–Pb jet cross-section ratio we use a 7 GeV/c leading track bias for both radii.

The ratios are built using the spectra shown above, with the full statistics.¹⁵ The computation of the uncertainties is discussed in Section 3.5. Note that the correlated uncertainties largely cancel.




Figure 3.57: Unfolded pp jet cross section ratio $\sigma_{R=0.2}/\sigma_{R=0.4}$.

¹⁵ Ideally one would want to split the sample in two statistically independent samples, in order to avoid correlating the statistical errors – but we do not have sufficient statistics for that, and furthermore the statistical error bars are likely correlated between the two samples due to the unfolding uncertainties.




Figure 3.58: Unfolded Pb-Pb jet cross section ratio $\sigma_{R=0.2}/\sigma_{R=0.4}$.




Figure 3.59: Unfolded pp and Pb-Pb jet cross section ratios $\sigma_{R=0.2}/\sigma_{R=0.4}$.

3.7 Comparison to theory

The jet R_{AA} results in Section 3.6 exhibit strong suppression of jet yields in Pb–Pb compared to pp, which has also been observed in previous measurements. The measurements presented, however, are the first jet R_{AA} measurements at $\sqrt{s_{NN}} = 5.02$ TeV at low jet p_T (i.e. $p_T < 100$ GeV/c), and the first inclusive jet measurements by ALICE extending to $R = 0.4$ at any collision energy. Since the mechanism of jet energy loss in the quark-gluon plasma is not precisely known, the main new physics to be extracted from these measurements is therefore through the quantitative comparison of the experimental data to theoretical models. We compare four theoretical predictions to the jet R_{AA} and jet cross-section ratio: JEWEL, the Linear Boltzmann Transport (LBT) model, the Soft Collinear Effective Theory with Glauber gluons (SCET_G) model, and the Hybrid model. Each of these models contains different physics of jet quenching, and their comparison to the data will be discussed in detail in this section.

The R_{AA} predictions of these models are compared to the measured data in Figures 3.60 and 3.61 for $R = 0.2$ and $R = 0.4$, respectively. The jet cross-section ratio predictions are compared to the measured data in Figure 3.62. The predictions are all computed using the anti- k_T jet algorithm with $|\eta| < 0.7 - R$.

3.7.1 Summary of theoretical models

JEWEL

JEWEL is a Monte Carlo implementation of BDMPS¹⁶ jet energy loss to leading log (LL) accuracy [93]. Starting with a parton shower, each parton interacts with thermal medium particles via radiative and collisional energy loss. The partons are eventually hadronized using Pythia. JEWEL allows the option to include the recoiling thermal medium particles in the jet energy (“recoil on”), or to ignore the recoiling medium particles (“recoil off”) [94]. In the case of including the recoils, the medium particles do not interact again with the medium – that is, they free-stream and do not thermalize. It is often believed that the

16. As discussed in Section 1.5.2, BDMPS is a particular pQCD formalism of jet energy loss, which occurs via momentum diffusion (described by a parameter $\hat{q}$) from many soft gluon radiations, and destructive interference of emissions that occur closely together (the LPM effect).




Figure 3.60: Jet R_{AA} for $R = 0.2$ full jets compared to JEWEL, LBT, SCET_G, and Hybrid model predictions.




Figure 3.61: Jet R_{AA} for $R = 0.4$ full jets compared to JEWEL, LBT, SCET_G, and Hybrid model predictions.




Figure 3.62: Unfolded Pb-Pb jet cross section ratio $\sigma_{R=0.2}/\sigma_{R=0.4}$ compared to theoretical predictions from JEWEL, LBT, SCET_G, and the Hybrid model.

“true” JEWEL prediction should be in between these two options, since in fact the recoiling medium particles are believed to partially thermalize. If recoils are not included, there is no need to perform background subtraction in JEWEL (since there are no medium particles). If recoils are included, however, there is a need to perform background subtraction – this can be done by several different schemes, of which we select the option “4MomSub” recommended as the default. JEWEL contains several free parameters that are fixed by independent measurements: The medium temperature and initial time, which are taken from a hydrodynamic calculation fitting to measured soft particle production, and the Debye mass, which is fixed by comparing to inclusive hadron suppression data from RHIC. It is therefore argued that JEWEL contains “zero” free parameters, since they are all fixed by independent measurements. The JEWEL predictions were generated internally (courtesy of Ritsuya Hosokawa (University of Tsukuba) [95]), and use $T = 590$ MeV and $t_0 = 0.4$ fm/c [96]. Note that these predictions do not include systematic errors, but rather only statistical uncertainties. Leading track requirements are applied as in the measured data.

LBT model

The Linear Boltzmann Transport (LBT) model implements pQCD energy loss based on gluon radiation induced by elastic scattering, and describes the evolution of recoiling medium particles with the thermal medium [97]. To perform this, the model implements linear Boltzmann equations to describe the transport of jet and recoil partons through the QGP. It is assumed that interactions between the jet shower and recoiling medium partons is negligible (this is the meaning of “linear” in LBT). The rate of induced gluon radiation is determined from a Higher Twist approach. An effective strong coupling constant α_s is taken as a free parameter fit to experimental data.

The model calculations are provided in [98]. The initial jet shower is produced by Pythia 8, neglecting nuclear modification of the PDF. The jet shower partons then interact with the evolving QGP through the LBT model described above. The QGP evolution itself is described by the CLVisc 3+1D hydrodynamic model, including event-by-event initial conditions. Hadronization of jet and medium partons is described by a parton recombination model as outlined in [98]. The calculation uses a background subtraction scheme described in [98], with only the need to address background from the pp event and jet-induced medium response. No systematic uncertainties were provided for this calculation. No leading track requirement is applied.

SCET_G model

The approach of Soft Collinear Effective Theory with Glauber gluons (SCET_G) to calculate jet energy loss in heavy-ion collisions describes interactions of jet partons with the hot QCD medium in an effective field theory via the exchange of “Glauber” gluons, which gives rise to in-medium collinear splitting functions [99]. The SCET_G approach builds on the approach of Soft Collinear Effective Theory (SCET) to compute a resummation of the pp jet cross-section in jet R ,¹⁷ but additionally includes in-medium modification. The theory contains a free parameter, g , the coupling between the Glauber gluons to collinear jet partons in the

17. In pp collisions, the cross-section is an expansion in $\alpha_s^n \ln^n R$. For small-radius jets, a fixed-order calculation in α_s is therefore not sufficient, since the jet R terms may give non-negligible contributions at larger n . There is therefore a need to “re-sum” the jet R contributions, which is to say, to evaluate the effect of the jet R contribution summed over all orders in n .

EFT, which is determined by fitting to experimental data. The model has been shown to describe certain jet substructure observables [100].

The predictions are performed by Dr. Haitao Li (LANL), according to [99] but with further improvements [101]. The pp jet cross-section is computed to NLO in α_s , and with a LL resummation in jet R . Medium effects are computed at NLO, but without (yet) a resummation in jet R (resulting in large systematic uncertainties for $R = 0.2$). The in-medium splitting functions described above include radiative processes, but these predictions do not yet include collisional energy loss. Note that this could have significant impact particularly on the larger radius jets, where it may increase suppression. The EFT coupling constant between the medium and jets is $g = 2.0$. The medium is evolved using 2+1D viscous hydrodynamics. For pp the CT14nlo PDF is used, and for Pb–Pb, the nCTEQ15FullNuc PDF is used. Energy loss in cold nuclear matter is also taken into account. The plotted error band represents the systematic uncertainty obtained by scale variations. No leading track requirement is applied.

Hybrid model

In the Hybrid model [102–105], partons are produced by vacuum pQCD, and shower according to vacuum pQCD (unmodified by the medium). In between these hard splittings, parton energy loss is modeled according to a gauge-gravity duality computation in $N = 4$ Supersymmetric Yang-Mills at infinitely strong coupling and large N_c . The energy loss is given by:

$$\frac{dE}{dx} = -\frac{4}{\pi} E_{in} \frac{x^2}{x_{\text{therm}}^2} \frac{1}{\sqrt{x_{\text{therm}}^2 - x^2}},$$

where E_{in} is the initial parton energy, and $x_{\text{therm}} = (E_{in}^{1/3}/T^{4/3})/2\kappa_{sc}$ is the maximal distance that a parton with E_{in} can travel in the plasma. The parameter κ_{sc} describing the amount of energy loss is the main free parameter in the model, and is fit to ATLAS and CMS hadron and jet data [40]. The model also includes a parameter K describing the transverse momentum broadening of the jet shower, and a parameter L_{res} describing the scale at which the medium can resolve two split partons. Note that in strongly-coupled energy loss models, energy loss is not determined by “out-of-cone” radiation, as in pQCD

models, but rather the lost energy is assumed to be largely thermalized into the medium. However, momentum is still conserved, and the model accounts for this and includes a wake in the direction of the jet [104].

The predictions are provided by Dr. Daniel Pablos (McGill University) [106]. A leading track cut of 5 GeV/c was applied for all predictions. The medium evolution is modeled by a hydrodynamic expansion. Two values of L_{res} are provided, $L_{res} = 0$ and $L_{res} = \frac{2}{\pi T}$. The plotted error bands represent the combination of statistical and systematic uncertainties.

3.8 Discussion and outlook

All models exhibit strong suppression, and produce the same qualitative trend of R_{AA} as a function of p_T , with R_{AA} increasing as p_T increases, and with a slowing increase as p_T increases.

In the case $R = 0.2$, we see that JEWEL under-predicts the jet R_{AA} , and appears to be inconsistent with the data regardless of whether medium recoils are included. We see that there is no significant difference between the recoil on or recoil off option in JEWEL for $R = 0.2$; we expect in general a smaller impact from medium recoil in smaller radius jets. The LBT model describes the data more consistently, although it has slight tension with the data. Note that the dominant systematic uncertainties in the data are positively correlated between p_T bins. Moreover, note that neither the JEWEL nor LBT predictions include systematic uncertainties. The SCET_G prediction is fully consistent with the data, although it suffers from large systematic uncertainties due to a lack of in-medium $\ln R$ re-summation in this calculation. The Hybrid model describes the trend of the data reasonably well, although like the LBT model, exhibits slight tension particularly in the $p_T < 100$ GeV/c range. The shapes of the p_T -dependence differ between the model predictions, most notably between SCET_G and the others. It should be noted that JEWEL has no free parameters in the fit, and so it faces the strictest test of all the models presented.

For the case $R = 0.4$, we see that the LBT and SCET_G models are consistent with the data, whereas JEWEL without recoils appears to be inconsistent with the data. The Hybrid model and JEWEL with recoils exhibit a slight hint of tension with the data. Note that the

SCET_G calculation did not include collisional energy loss, which the authors anticipate to increase the suppression for $R = 0.4$. The uncertainties of the measured data are larger than the $R = 0.2$ case, and the p_T range is more restricted as well, which makes distinguishing the models more challenging. Moreover, the R_{AA} spectrum exhibits a weaker p_T -dependence at higher p_T , which makes it further challenging to distinguish the models. On the other hand, the model predictions span a wider range of R_{AA} than in the case of $R = 0.2$, which highlights the importance of the measurement of the R -dependence of the jet R_{AA} . It is worth noting that our results are consistent with the $R = 0.4$ jet R_{AA} measured by ATLAS at $\sqrt{s_{NN}} = 5.02$ TeV in the p_T range where they overlap; these are the only other existing jet R_{AA} measurements at $\sqrt{s_{NN}} = 5.02$ TeV [77].

The jet cross-section ratio $R = 0.2/R = 0.4$ is consistent with all theoretical predictions, within even larger uncertainties. The ratio is smaller than 1 due to the fact that $R = 0.2$ jets contain a smaller fraction of the “true” jet energy than $R = 0.4$ jets. Note that the Pb–Pb jet cross-section ratio is built with a leading track requirement of 7 GeV/c for both radii. This bias is expected to be fairly small, biasing the jet sample by excluding on order $\approx 10\%$ of jets from the R_{AA} computation. We do not use statistically independent samples, since we do not have the precision to do so. Note that the statistical uncertainties are partially correlated, due to error propagation through the unfolding procedure, and that the statistical errors are therefore also partially correlated to the shape uncertainties. This likely results in a conservative statistical uncertainty estimation, since there may be significant cancellation between the two radii that is not taken into account.

From these comparisons, two general conclusions are apparent. First, most (but not all) of the models can describe the R_{AA} reasonably well, but with a hint of tension. This necessitates investigation of complementary jet observables (jet substructure, heavy-flavor jets, and others) and the need for global fits with fixed free parameters as e.g. in JEWEL. There is also a need to standardize the ingredients of jet energy loss calculations other than the jet energy loss part: the description of the initial state, the input jet spectrum, the hydrodynamic evolution of the medium, and the model of hadronization. The predictions considered above typically use different strategies for each of these pieces, which raises the question of how significant their differences are for the final results. The JETSCAPE

project [107] is an important ongoing development that seeks to systematize these elements in a flexible heavy-ion event generator that includes a variety of theoretical predictions. From the experimental side, the measurement of full jets as opposed to charged jets is an important step to meaningfully compare experiment to theory.

Second, the $R = 0.4 R_{AA}$ and jet cross-section ratio have large uncertainties and therefore limited power to distinguish theories. The results demonstrate that the modification to the jet shape between $R = 0.2 - 0.4$ cannot be too large, yet this observable evidently lacks the precision to make a definitive statement about the models. This necessitates a need for increased statistics to stabilize the unfolding procedure, as well as alternate jet shape studies to reach the necessary level of precision. ALICE is particularly well-suited for jet substructure measurements, and including soft particles could be very important for capturing the medium response effects.

One striking observation is that both pQCD-based energy loss models and the Hybrid model agree with the data. These are very different physics explanations – the pQCD-based models rely on a weakly-coupled description of the jet-medium interaction, whereas the Hybrid model uses a strongly-coupled jet-medium interaction. Does the weak coupling vs. strong coupling description have implications for the quasiparticle nature of deconfined QCD matter? If so, what can we learn about them? Can this ultimately give us information on what temperatures and couplings they exist at? The answers to these questions may hold the key to understanding confinement in QCD.

It is an exciting and open question whether heavy-ion jet physics will be able to answer these challenges. To do so, new precision and new experimental strategies will be necessary. One possibility to search for quasiparticles is to look for modification in the angular deflection of di-jets with high-statistics measurements [108]. New experimental observables, particularly in jet substructure, may also lead to new insights [109]. Perhaps the application of advanced machine learning techniques will guide new understanding. Measurements of heavy-flavor jets may also yield a definitive description of the flavor-dependence of jet energy loss, which could be a crucial input. There is no shortage of interesting measurements still to make – the question will be whether we are clever enough to look in the right places, and if nature is kind enough to leave sufficient hints for us to discover.

Appendix A

TPC upgrade R&D

A.1 2-GEM+MMG

A configuration of two GEMs and one MMG as a gain configuration for TPC gas amplification was investigated in laboratory tests and a test beam environment. The goal of this design, similar to the 4-GEM design, is to minimize the buildup of space charge in the drift volume of such detectors in order to eliminate the standard gating grid and its resultant dead time, while preserving good tracking and particle identification performance. A characterization of the performance of this design in terms of IBF, energy resolution, and stability is reported in [71].

A test beam campaign was conducted at CERN in November-December 2014, using two 2-GEM/MMG 21x26 cm chambers constructed at Yale (as well as a 4-GEM inner-readout chamber constructed by ALICE colleagues [68]). The test beam occurred in two phases. The first phase at the PS beam, with setup shown in Fig. A.1 (left), used a secondary beam of 1-3 GeV pions and electrons to measure the PID separation ability of the chambers. The result is shown in Fig. A.1 (right), and demonstrated successful PID separation, although during the course of the data analysis it was discovered that one of the four readout cards was excessively noisy, which inhibited the PID performance – so the true PID performance of the chamber is expected to be better than shown in Fig. A.1. The second phase at the SPS beam tested the sparking rate of the chambers under high-multiplicity LHC-like conditions [71]. We observed a discharge rate of 3.5×10^{-10} per incident particle, with the

discharges occurring primarily due to the MMG. This rate is higher than desired, and while it does not pose risk to the detector itself, additional studies were needed to improve the spark protection of the readout electronics. Both results were reported at the APS April Meeting 2015¹ and in [61, 71].




Figure A.1: Left: Test beam setup at the PS beam at CERN. A Cherenkov counter was used to identify electrons and pions, and scintillators were used for triggers. Right: Preliminary results for dE/dx separation of 3.56σ obtained from combined tracks of the two Yale 2-GEM/MMG chambers. It was later discovered that one of the chambers contained an excessively noisy readout card, limiting the separation we measured.

1. J. Mulligan, “Test Beam Results for ALICE TPC Upgrade Prototypes”.

A.2 Multi-layer extended gating grid

A novel idea to control ion back-flow in time projection chambers is to use a multi-layer extended gating grid to capture back-flowing ions at the expense of live time and electron transparency [72]. In comparison to a traditional gating grid, the extension of the grid with multiple layers allows a longer time for ions to drift through the gate, while still collecting the ions quickly. The operating principle is that the gate remains transparent to electrons until the ion drift time exceeds the grid length (divided by the ion drift velocity), at which point the gate is closed and the ions are collected. Enhanced IBF suppression comes at the sacrifice of live time and electron transparency; for a given IBF tolerance, the design goal is to increase the live time fraction A while maintaining sufficient electron transparency for reconstruction performance. Such a design could operate as a primary means of IBF suppression, or in cooperation with other elements such as Gas Electron Multipliers. Early work suggests that for a wire-plane gate, low-field regions between the wires prevent some ions from being captured quickly [72]. The detailed simulations presented below quantify this effect and serve as an initial study of the general feasibility of a multi-layer extended gating grid.

In this study [73], I perform simulations of a four-layer grid for the ALICE and STAR time projection chambers, using Ne – CO₂ (90 – 10) and Ar – CH₄ (90 – 10) gas mixtures, respectively. I report the live time and electron transparency for both 90% and 99% ion back-flow suppression. Additionally, for the ALICE configuration I study several effects: using a mesh vs. wire-plane grid, including a magnetic field, and varying the over-voltage distribution in the gating region. For 90% ion back-flow suppression, I achieve 75% live time with 86% electron transparency for ALICE, and 95% live time with 83% electron transparency for STAR.

A.2.1 Simulation Configurations

In order to study the performance of the grid in various TPC conditions, I simulate the gating region for two large gas TPCs: ALICE and STAR [110]. These TPCs use different gas mixtures (with significantly different ion mobilities) and different drift fields, which




Figure A.2: Schematics of the open (left) and closed (right) gating configurations. In the open configuration, electrons pass through with small losses, while positive ions back-drift through the gate. In the closed configuration, these back-drifted positive ions are collected on the first and third planes by E_c .

considerably impact the gating performance.

In both TPC configurations, I consider a four-layer grid (Fig. A.2), with the open field E_o parallel/anti-parallel to the closed field E_c . The spacing between layers is 3 mm, and the inter-layer wire spacing is 2 mm. I use a 3 mm drift volume above the grid, and a grounded plane 3 mm below the grid. The wire diameter is 100 μm . The ions are collected on the first and third planes from the gas amplification region, with $E_c \approx 2 \text{ kV/cm}$.

The fields are constructed using the finite element method software ANSYS [111]; the electron and ion drifts are simulated using Garfield++ [112]. Collision-level tracking is performed for electrons (“microscopic tracking”), and a more coarse-grained Monte Carlo tracking is used for ions. Diffusion is included for both electrons and ions. It should be noted that while Garfield++ can natively solve 2D fields, close examination revealed that the ANSYS solution is more accurate near the wires.

ALICE TPC

For the ALICE TPC, I use a gas mixture of Ne – CO₂ (90 – 10), as configured for LHC Run 1. The drift field is $\approx 0.4 \text{ kV/cm}$; a representative voltage switch required on the four gating planes (in volts) is: $(-600, 0, -600, 0) \leftrightarrow (-120, -240, -360, -480)$. The electron drift velocity in this mixture for the considered drift field is 2.73 cm/ μs [63]. Binary ion mobilities from the literature are linearly extrapolated to low fields and combined for the




Figure A.3: Visualizations of the multi-layer extended gating grids in the open configuration, with electron drift lines traveling downward. Plotted are 6×6 arrays of $2\text{mm} \times 2\text{mm} \times 15\text{mm}$ unit cells used for the simulation. Left: Wire-plane configuration. Right: Mesh configuration. Mesh spacing in-plane is 2 mm.

gas mixture using Blanc's Law [113–115]. The dominant ion in this mixture is CO_2^+ [116].

To study differences in ion collection time and electron transparency, I study separately a wire configuration and a mesh configuration (Fig. A.3). The final finite element meshes contain approximately $3 \cdot 10^5$ elements for the wire configuration unit cell, and $2 \cdot 10^6$ elements for the mesh configuration unit cell. These correspond to an average element size of $56 \mu\text{m}$ for the wire configuration, and $33 \mu\text{m}$ for the mesh configuration. However, adaptive meshing is employed, which creates finer elements near geometrical features. The final meshes were examined for quality, and an informal convergence study was performed, in which iteratively refined meshes were produced, and the maximum field near the wires showed convergence to $< 10\%$.

STAR TPC

For the STAR TPC, I use a gas mixture of $\text{Ar} - \text{CH}_4$ (90 – 10). The drift field is approximately 140 V/cm . The electron drift velocity in this mixture for the considered drift field is $5.45 \text{ cm}/\mu\text{s}$, and the ion mobility $1.6 \text{ cm}^2/\text{V} \cdot \text{s}$.

Only a wire-plane configuration is simulated. The finite element mesh contains approx-

imately $3 \cdot 10^5$ elements for the unit cell, as for the ALICE wire-plane mesh.

A.2.2 Simulation Results – ALICE TPC

Electron Transparency

I measure electron transparency by randomly placing electrons at the top of the drift region, and measuring the fraction that pass through the grid in the open configuration. At each layer in the grid, I increment the field by a value ΔE in order to boost the transparency; increasing ΔE amounts to putting negative charge on the wires, which repels drifting electrons. I use fixed over-voltages corresponding to $\Delta E = 0, 10, 20, 30$ V/cm across each plane, yielding average open gating fields of $E_o = 400, 425, 450, 475$ V/cm. Figure A.4 shows the results for both the wire-plane and mesh configurations.

Additionally, I repeated the electron transparency measurement in the mesh configuration with a magnetic field $B = 0.5$ T parallel to the electric field. This results in a slight increase in transparency (Fig. A.4), which may be due to reduced transverse diffusion (from the B -field) outweighing $E \times B$ effects (which may deviate drifting electrons from the electric field lines into a wire); the cause remains to be investigated.

Next, I examined the over-voltage distribution to determine if there is an optimal way to distribute ΔE over different planes, rather than fixing it to be constant across each layer. A comparison of fixed ΔE over each plane against having nonzero ΔE only across the first plane shows little difference (Fig. A.5). In the latter case, fewer electrons are captured on the first layer, but more are captured in subsequent layers. This suggests that if more layers are added to the grid, the fixed ΔE configuration is better.

Live Time and Ion Collection

Following [1], the live time fraction A of the gating grid can be written

$$A = \frac{T_{active}}{T_{cycle}} = \frac{T_o - T_e}{T_o + T_c},$$

where T_o is the open time, T_c is the closed time, and T_e is the time for an electron to drift the length of the chamber. For a gating grid of N planes, layer separation Δh , ion mobility




Figure A.4: Electron transparency as a function of field incrementation ΔE at each layer of the grid, in the ALICE configuration. The error bars estimate the statistical uncertainty. For the wire-plane configuration, each point corresponds to $2.5 \cdot 10^4$ electrons. For the mesh configuration, each point corresponds to 10^4 electrons.

K_I , closed field E_c , average drift field within grid E_o , drift length L_e , and electron drift velocity v_e , these times can be estimated as:

$$\begin{aligned}
 T_o &= \frac{N\Delta h}{K_I E_o}, \\
 T_c &= \alpha \frac{\Delta h}{K_I E_c}, \\
 T_e &= \frac{L_e}{v_e}.
 \end{aligned}$$

The factor α in the collection time accounts for the fact that the field is not from parallel plates, but rather has low-field regions in between the wires due to saddle points in the potential. Therefore α depends on the IBF threshold imposed. From the above expressions, the live time can be written:

$$A = \frac{1 - \frac{E_o K_I L_e}{N \Delta h v_e}}{1 + \alpha \frac{E_o}{N E_c}}. \quad (\text{A.1})$$




Figure A.5: Electron transparency as a function of average ΔE across each layer of the grid, for two different overvoltage distributions, in the ALICE configuration. Each point corresponds to 10^4 electrons; the error bars estimate the statistical uncertainty.

This makes clear the dependence of the live time on various detector parameters. The present simulations involve the following parameter values, determined with the ALICE TPC in mind:

Param	Estimated value	$A \uparrow$ if	Physical reason	Constrained by
N	4	$\uparrow$	Longer T_o	Transparency
Δh	3 mm	$\uparrow$	Longer T_o, T_c ; fixed T_e	Voltage; transparency
E_o	400 – 475 V/cm	$\downarrow$	Longer T_o	Transparency
E_c	2000 V/cm	$\uparrow$	Faster collection	Voltage
K_I	4.8 cm ² /V · s	$\downarrow$	Longer T_o, T_c ; fixed T_e	Gas choice
v_e	2.73 cm/ μ s	$\uparrow$	Smaller T_e	Gas choice
L_e	250 cm	$\downarrow$	Smaller T_e	Detector size
α	1 – 4	$\downarrow$	Longer collection time	IBF tolerance
w	2 mm	$\downarrow$	Smaller saddle area	Transparency

Ion Collection

To estimate the live time of the simulated configuration, the parameter α must be measured in simulation, or equivalently T_c . I measure the ion collection time by randomly (uniformly) placing ions in the gating region, as one would expect for backdrifting ions, and counting the time it takes to collect the ions in the closed configuration. Figure A.6 shows the results.

Note that a constant plateau out to $t = \frac{\Delta h}{K_I E_c} \approx 31 \mu s$ exists for both cases (as expected from a parallel plate solution), while the wire configuration has a significantly longer tail, due to more low-field regions.

Additionally, I introduced a magnetic field $B = 0.5$ T in the mesh configuration, and repeated the ion collection. This causes no change in collection times, as expected since the magnetic force on ions is negligible due to their slow drift velocities (and additionally, the magnetic field would perturb ion trajectories not only into the low-field regions, but out of them as well).

Live Time Estimates

The measured collection times, in conjunction with the above table of parameters, yield the following live times, reported for 90% and 99% IBF, and for transparencies corresponding to $\Delta E = 0, 20$ V/cm, for both the wire-plane and mesh configurations:

	Wire Configuration	Mesh Configuration
$E_o = 400$ V/cm	(80% transparency)	(62% transparency)
99% IBF	73%	77%
90% IBF	78%	80%
$E_o = 450$ V/cm	(86% transparency)	(68% transparency)
99% IBF	70%	74%
90% IBF	75%	78%





Figure A.6: Histograms of $5 \cdot 10^4$ ion collection times in the ALICE configuration. The 90% and 99% IBF thresholds are illustrated. Top: Wire-plane configuration. Bottom: Mesh configuration.

A.2.3 Simulation Results – STAR TPC

Recalling equation (1), the corresponding table of values for STAR is estimated to be:

Param	Estimated value
N	4
Δh	3 mm
E_o	190 V/cm
E_c	2000 V/cm
K_I	1.6 cm ² /V · s
v_e	5.45 cm/ μ s
L_e	209 cm
α	2
w	2 mm

I take $E_o = 190$ V/cm, corresponding to a 140 V/cm drift field plus 20 V/cm per plane overvoltage. I then measure the electron transparency with this overvoltage via simulation of 25,000 electrons in the wire-plane configuration to be:

$$82.6\%.$$

The statistical error is $\approx \sqrt{np(1-p)}/n \approx 0.2\%$, but the dominant uncertainty is expected to come from the field map or one of many other possible sources of error, which haven't been quantified. The transparency could be boosted at the expense of live time by increasing the average gate field E_o . Also, small improvements $\sim 1\%$ were observed in ALICE simulations when a magnetic field was included, although this was not done here (the electrons in STAR are hotter than in ALICE, so the diffusion and $E \times B$ effects are both probably larger, and one would need to verify the outcome of their balance).

I measure the parameter α by simulating the ion collection time for 90% and 99% of backdrifting ions (Fig. A.7). Recall that $\alpha = 2$ means that the closed time is equal to twice that of a perfect parallel plate.

I find

$$\alpha_{90\%} = 1.6, \quad \alpha_{99\%} = 2.9.$$

The live time for this set of parameters, with 82.6% electron transparency, is then:

$$A_{90\%} = 95\%, \quad A_{99\%} = 93\%.$$

Relative to ALICE, the smaller drift field and the smaller ion mobility allow the open gate to be open longer (since ions drift back more slowly), and additionally the larger electron drift velocity reduces the electron drift time per cycle T_e . Note also that in this scheme, the detector could operate continuously for up to $T_o \sim 4$ ms before the gate needs to be closed.




Figure A.7: Histogram of $5 \cdot 10^4$ ion collection times in the STAR wire-plane configuration. The 90% and 99% IBF thresholds are illustrated.

A.2.4 Discussion

The results above provide an early quantitative look at possible modified gating grid configurations. The specific results for live time and electron transparency should not be viewed as expected limits, but rather starting points from which optimization could begin.

One concrete conclusion, however, is that the mesh grid appears untenable. The idea of the mesh configuration is to increase the live time A by decreasing T_c , at the expense of transparency. However, the simulations suggest that the transparency cost is large for only a small improvement in live time. Further, it should be noted that the wire configuration has an additional advantage in that it preserves momentum information along the direction of the wires, whereas the mesh distorts momentum information in both directions. If one is determined to reduce T_c , an additional avenue to pursue is a dynamically switched gating cycle, in which the saddle point ions are swept out of the low-field region. This could be accomplished by the closed time consisting of two periods of $E_c \parallel E_o$ interspersed with a period of $E_c \perp E_o$.

A handful of additional parameters directly exhibit a tradeoff between live time and electron transparency: N , Δh , and E_o . The idea in designing a detector is to favor those variables that give maximal live time boost with minimal transparency loss.

If electron transparency is a concern, one should increase E_o as much as possible. To boost the live time, equation (A.1) suggests it is better to try to increase T_o rather than decrease T_c . Further study of varying N , Δh , and E_o should be undertaken. For example, in the ALICE configuration, plots of the final positions of electrons show that transparency would decrease by approximately 3 – 5% if another layer is added to the grid. This extra layer will cause $T_o \rightarrow \frac{5}{4}T_o$, yielding live time improvements of approximately 5%. Adding yet another layer would have an even smaller effect on transparency, and yield a further boost in live time. Similar arguments can be made for increasing Δh , at the expense of longer collection time, perhaps worse transparency, and larger voltage switches. This option may be particularly attractive for situations in which 90% IBF suppression is acceptable. The possibility of significantly increasing E_o in concert with these approaches may be particularly appealing, and should be tested.

To estimate the live time for configurations other than ALICE or STAR, one can use equation (A.1), upon choosing an α comparable to those for the ALICE and STAR results (for a given IBF suppression). Electron transparency estimates are more difficult, and require detailed simulation.

Overall, the presented simulations suggest that a multi-layer extended gating grid may be a feasible option for reducing IBF in TPCs, depending on acceptable losses of live time and electron transparency, and the TPC configuration. There remains significant room for optimization, and it is expected that results will continue to improve as they are adapted for particular applications. Experimental tests are also being pursued.

Appendix B

Inclusive jet analysis

B.1 Analysis code

This analysis employs the ALICE EMCAL-Jet framework. The bulk of the analysis was carried out using the task `AliAnalysisTaskEmcalJetPerformance`. A significant amount of post-processing was then done using several PyRoot-based macros (not in `AliPhysics`), available on request. The analysis was processed on the ALICE LEGO train system, using the train `Jets EMC PbPb`. Relevant train numbers are:

- Main results:
 - Measured data: 3819, 3822
 - Background scale factors: 3361
 - Embedding response: 3971, 3975-3993
- Cell threshold studies:
 - Measured data: 3192
 - Background scale factors: 3057
 - Embedding response: 3152-3152, 3169-3186

B.2 Combining $p_{T,\text{hard}}$ bins

B.2.1 Scaling $p_{T,\text{hard}}$ bins

In order to construct a minimum bias equivalent sample from a set of $p_{T,\text{hard}}$ bins, we must properly scale each $p_{T,\text{hard}}$ bin according to the cross-section of that $p_{T,\text{hard}}$, the average number of trials to produce an accepted event in our detector acceptance, and the relative number of events in each $p_{T,\text{hard}}$ bin. For a generic counting histogram $h(x)$, we combine the $p_{T,\text{hard}}$ bins k as:

$$h(x) = \sum_k c_k h_k(x),$$

where

$$c_k = s_k \times \frac{\langle N_{event}^{p_{T,\text{hard}}} \rangle}{N_{event}^k}, \quad s_k = \frac{\sigma_{p_{T,\text{hard}}} \text{ per event}}{N_{trials} \text{ per event}}.$$

The factors s_k for LHC16j5 are listed in Table B.1. The relative weights can be verified to be correct by plotting the re-weighted $p_{T,\text{hard}}$ distribution.

In order to produce a cross-section (as is reported in Section 3.6), we must further scale by the average number of events per $p_{T,\text{hard}}$ bin:

$$\frac{d\sigma_{jet}}{dp_T}(p_T) = \frac{1}{\langle N_{event}^{p_{T,\text{hard}}} \rangle} \sum_k c_k \frac{dN_{jet}^k}{dp_T}(p_T).$$

B.2.2 Removing outliers

There is a generic pathology of constructing a weighted sum of finitely sampled statistical distributions: There is large sensitivity to the statistical fluctuations in those distributions with large weight. In our case, the problem is that the low $p_{T,\text{hard}}$ bins may fluctuate to have a count when the expected value is smaller than one count. If this is not dealt with, significant spikes are observed in the merged distributions.

We implement two strategies to mitigate these effects:

- We reject events with $p_T^{\text{jet}} > 4 \times p_{T,\text{hard}}$ (such as due to ISR or overlapping jets)

Bin	$p_{\text{T,hard}}$ range (GeV/c)	s_k
1	5-7	0.3471
2	7-9	0.2978
3	9-12	0.2888
4	12-16	0.1894
5	16-21	9.596e-2
6	21-28	4.613e-2
7	28-36	1.619e-2
8	36-45	5.920e-3
9	45-57	2.547e-3
10	57-70	8.795e-4
11	70-85	3.517e-4
12	85-99	1.272e-4
13	99-115	6.276e-5
14	115-132	2.956e-5
15	132-150	1.452e-5
16	150-169	7.395e-6
17	169-190	3.955e-6
18	190-212	2.111e-6
19	212-235	1.141e-6
20	>235	1.472e-06

Table B.1: $p_{\text{T,hard}}$ scale factors s_k for LHC16j5.

- We truncate the response matrix in each $p_{\text{T,hard}}$ bin when the $p_{\text{T,gen}}^{\text{jet}}$ projection has a 4-bin moving average below 2 counts.

These strategies slightly overcorrect the yield, but are expected to have no significant effect on the merged distribution.

In addition to outliers due to fluctuations in the pp spectrum, an additional source of outliers may arise in embedding from a limited sampling of the Pb–Pb background for a given $p_{\text{T,hard}}$ event. We do not address this possibility.

B.3 EMCal cell threshold studies, additional details

Below are listed the background scale factors for three sets of EMCal cell thresholds. The parameterized scale factor is obtained by fitting the mean scale factor with a second-order polynomial over the range $C \in [0, 50]$:

- $E_{\text{cell}} = 50 \text{ MeV}$, $E_{\text{seed}} = 100 \text{ MeV}$: $s(C) = 1.785 - 0.00815C + 0.000064C^2$.
- $E_{\text{cell}} = 100 \text{ MeV}$, $E_{\text{seed}} = 300 \text{ MeV}$: $s(C) = 1.443 - 0.00490C + 0.000054C^2$.
- $E_{\text{cell}} = 150 \text{ MeV}$, $E_{\text{seed}} = 300 \text{ MeV}$: $s(C) = 1.397 - 0.00355C + 0.000038C^2$.




Figure B.1: Mean background scale factor computed in the EMCal, for various EMCal cell thresholds. The errors plotted on the right-hand plot are the errors in the mean of the scale factor. Top Left: $E_{\text{cell}} = 50 \text{ MeV}$, $E_{\text{seed}} = 100 \text{ MeV}$. Top Right: $E_{\text{cell}} = 100 \text{ MeV}$, $E_{\text{seed}} = 300 \text{ MeV}$. Bottom: $E_{\text{cell}} = 150 \text{ MeV}$, $E_{\text{seed}} = 300 \text{ MeV}$.

The δ_{p_T} distributions are also provided below, for reference. Details of the unfolding procedure are as follows. Note that the unfolding ranges and other considerations were not treated as strictly as for the actual analysis, since here we are interested only in the relative performance. A fine-binned response matrix is generated over the range $p_{T,\text{gen}}^{\text{jet}}, p_{T,\text{det}}^{\text{jet}} \in [0, 200] \text{ GeV}/c$, which is then truncated over the range $p_{T,\text{det}}^{\text{jet}} \in [36, 120] \text{ GeV}/c$, and re-binned into the response matrix used for the actual unfolding. The first three $p_{T,\text{hard}}$ bins (5-12 GeV/c) were excluded from this process, since at the time a low- $p_{T,\text{gen}}^{\text{jet}}$ band in the

response matrix was not understood. The kinematic efficiency is also shown below, and was corrected for in the unfolding process. The d-vector is shown for each of the considered cell thresholds. From these, $k = 4$ appears to be a reasonable choice.




Figure B.2: Background fluctuations δ_{p_T} in the EMCAL, for $R = 0.2$. Upper left four plots: $E_{\text{cell}} = 50$ MeV, $E_{\text{seed}} = 100$ MeV. Upper right four plots: $E_{\text{cell}} = 100$ MeV, $E_{\text{seed}} = 300$ MeV. Bottom four plots: $E_{\text{cell}} = 150$ MeV, $E_{\text{seed}} = 300$ MeV.




Figure B.3: Background fluctuations δ_{p_T} in the EMCAL, for $R = 0.3$. Upper left four plots: $E_{\text{cell}} = 50$ MeV, $E_{\text{seed}} = 100$ MeV. Upper right four plots: $E_{\text{cell}} = 100$ MeV, $E_{\text{seed}} = 300$ MeV. Bottom four plots: $E_{\text{cell}} = 150$ MeV, $E_{\text{seed}} = 300$ MeV.

The correlation coefficients are plotted below. The unfolded spectra for various k is shown below, for the case $E_{\text{cell}} = 50$ MeV, as well as the ratio of nearby k to $k = 4$.

Refolding and closure tests are shown below for the case $k = 4$. Similar structures in the ratios appear for $k = 3$ and $k = 5$. Note that two independent datasets were used each the refolding and closure tests to generate two statistically independent response matrices, “Response1” and “Response2”. The re-folding test unfolds the measured Pb–Pb with “Response1”, and then folds the result with “Response2” and compares to the original measured spectrum. The closure tests takes the MC truth spectrum (generated with full statistics, but smeared with a Gaussian according to the statistical precision of the measured spectrum) and folds it using “Response1”, and then unfolds the result with “Response2”.




Figure B.4: Background fluctuations δp_T in the EMCAL, for $R = 0.4$. Upper left four plots: $E_{\text{cell}} = 50$ MeV, $E_{\text{seed}} = 100$ MeV. Upper right four plots: $E_{\text{cell}} = 100$ MeV, $E_{\text{seed}} = 300$ MeV. Bottom four plots: $E_{\text{cell}} = 150$ MeV, $E_{\text{seed}} = 300$ MeV.




Figure B.5: Response matrices for the case $E_{\text{cell}} = 50$ MeV. Left: Fine-binned response matrix. Right: Response matrix used in unfolding.




Figure B.6: Kinematic efficiency for the case $E_{\text{cell}} = 50$ MeV.




Figure B.7: Plot of the d-vector for the cases $E_{\text{cell}} = 50$ MeV, $E_{\text{cell}} = 100$ MeV, $E_{\text{cell}} = 150$ MeV, respectively.




Figure B.8: Plot of the correlation coefficients for the case $E_{\text{cell}} = 50$ MeV for $k = 2, 3, 4, 5$, respectively.




Figure B.9: Unfolded spectra for the case $E_{\text{cell}} = 50$ MeV for various k .




Figure B.10: Left: Refolding test for the case $E_{\text{cell}} = 50$ MeV for $k = 4$. Right: Closure test for the case $E_{\text{cell}} = 50$ MeV for $k = 4$.

B.4 $R = 0.4$ unfolding details




Figure B.11: Left: Fine-binned response matrix for $R = 0.4$, over the full range of $p_{T,\text{det}}^{\text{jet}}$ and $p_{T,\text{gen}}^{\text{jet}}$. Right: Re-binned response matrix for $R = 0.4$, used in the unfolding, with selections imposed on $p_{T,\text{det}}^{\text{jet}}$ and $p_{T,\text{gen}}^{\text{jet}}$ as described in the text.




Figure B.12: Left: Unfolded result for $R = 0.4$ jets as a function of k . Right: D-vector for $R = 0.4$ jets.




Figure B.13: Pearson correlation coefficients for $R = 0.4$ jets for $k = 2, 3, 4, 5$ from left to right.




Figure B.14: Left: Re-folding test for $R = 0.4$ jets, performed for the main result $k = 4$. Right: “Self-closure” test comparing Pb-Pb embedded det-level data (statistically smeared according to uncertainties in real data) with Pythia truth.

Bibliography

- [1] T. Aoyama, M. Hayakawa, T. Kinoshita, M. Nio, Tenth-order electron anomalous magnetic moment: Contribution of diagrams without closed lepton loops, Phys. Rev. D 91 (2015) 033006. doi:10.1103/PhysRevD.91.033006.
- [2] G. Dissertori, G. Salam, Quantum chromodynamics.
URL <http://pdg.lbl.gov/2010/reviews/rpp2010-rev-qcd.pdf>
- [3] M. Peskin, D. Schroeder, An introduction to quantum field theory, Addison-Wesley, 1995.
- [4] A. Zee, Quantum field theory in a nutshell, Princeton University Press, 2003.
- [5] S. Bethke, The 2009 world average of α_s , The European Physical Journal C 64 (4) (2009) 689–703. doi:10.1140/epjc/s10052-009-1173-1.
- [6] H. Fukaya, Can axial U(1) anomaly disappear at high temperature?, EPJ Web Conf. 175 (2018) 01012. doi:10.1051/epjconf/201817501012.
- [7] B. Foster, A. Martin, M. Vinster, Structure functions.
URL <http://pdg.lbl.gov/2011/reviews/rpp2011-rev-structure-functions.pdf>
- [8] A. Accardi, et al., EIC white paper. Electron ion collider: The next QCD frontier, BNL-98815-2012-JA.
- [9] The QCD phase diagram.
URL http://theorie.ikp.physik.tu-darmstadt.de/~tripolt/conferences/Tripolt_QCD_Phase_Diagram_2012.pdf

- [10] A. Bazavov, T. Bhattacharya, C. DeTar, H.-T. Ding, S. Gottlieb, R. Gupta, P. Hegde, U. M. Heller, F. Karsch, E. Laermann, L. Levkova, S. Mukherjee, P. Petreczky, C. Schmidt, C. Schroeder, R. A. Soltz, W. Soeldner, R. Sugar, M. Wagner, P. Vranas, Equation of state in $(2 + 1)$ -flavor qcd, Phys. Rev. D 90 (2014) 094503. doi:10.1103/PhysRevD.90.094503.
- [11] T. Schafer, D. Teaney, Nearly perfect fluidity: from cold atomic gases to hot quark gluon plasmas, Reports on Progress in Physics 72 (12) (2009) 126001.
URL <http://stacks.iop.org/0034-4885/72/i=12/a=126001>
- [12] P. K. Kovtun, D. T. Son, A. O. Starinets, Viscosity in strongly interacting quantum field theories from black hole physics, Phys. Rev. Lett. 94 (2005) 111601. doi:10.1103/PhysRevLett.94.111601.
- [13] Entropy production, viscosity bounds and bumpy black holes, Journal of High Energy Physics 2016 (3) (2016) 170. doi:10.1007/JHEP03(2016)170.
- [14] W. Busza, K. Rajagopal, W. van der Schee, Heavy ion collisions: The big picture, and the big questions.
URL <https://arxiv.org/abs/1802.04801>
- [15] New state of matter created at CERN (2000).
URL <https://cds.cern.ch/record/716634?ln=en>
- [16] T. K. NAYAK, Heavy ions: Results from the large hadron collider, Pramana 79 (4) (2012) 719–735. doi:10.1007/s12043-012-0373-7.
- [17] M. L. Miller, K. Reygers, S. J. Sanders, P. Steinberg, Glauber modeling in high-energy nuclear collisions, Annual Review of Nuclear and Particle Science 57 (1) (2007) 205–243. doi:10.1146/annurev.nucl.57.090506.123020.
- [18] B. Schenke, P. Tribedy, R. Venugopalan, Fluctuating glasma initial conditions and flow in heavy ion collisions, Phys. Rev. Lett. 108 (2012) 252301. doi:10.1103/PhysRevLett.108.252301.

- [19] J. S. Moreland, J. E. Bernhard, S. A. Bass, Alternative ansatz to wounded nucleon and binary collision scaling in high-energy nuclear collisions, *Phys. Rev. C* 92 (2015) 011901. doi:10.1103/PhysRevC.92.011901.
- [20] B. Schenke, S. Jeon, C. Gale, Elliptic and triangular flow in event-by-event $D = 3 + 1$ viscous hydrodynamics, *Phys. Rev. Lett.* 106 (2011) 042301. doi:10.1103/PhysRevLett.106.042301.
- [21] C. Shen, Z. Qiu, H. Song, J. Bernhard, S. Bass, U. Heinz, The iEBE-VISHNU code package for relativistic heavy-ion collisions, *Computer Physics Communications* 199 (2016) 61 – 85. doi:10.1016/j.cpc.2015.08.039.
- [22] L.-G. Pang, H. Petersen, X.-N. Wang, Pseudorapidity distribution and decorrelation of anisotropic flow within the open-computing-language implementation CLVisc hydrodynamics, *Phys. Rev. C* 97 (2018) 064918. doi:10.1103/PhysRevC.97.064918.
- [23] Z.-W. Lin, C. M. Ko, B.-A. Li, B. Zhang, S. Pal, Multiphase transport model for relativistic heavy ion collisions, *Phys. Rev. C* 72 (2005) 064901. doi:10.1103/PhysRevC.72.064901.
- [24] J. Uphoff, O. Fochler, Z. Xu, C. Greiner, Heavy-quark production in ultrarelativistic heavy-ion collisions within a partonic transport model, *Phys. Rev. C* 82 (2010) 044906. doi:10.1103/PhysRevC.82.044906.
- [25] T. Sjostrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, P. Z. Skands, An introduction to PYTHIA 8.2, *Computer Physics Communications* 191 (2015) 159 – 177. doi:10.1016/j.cpc.2015.01.024.
- [26] M. Bähr, S. Gieseke, M. A. Gigg, D. Grellscheid, K. Hamilton, O. Latunde-Dada, S. Plätzer, P. Richardson, M. H. Seymour, A. Sherstnev, B. R. Webber, Herwig++ physics and manual, *The European Physical Journal C* 58 (4) (2008) 639–707. doi:10.1140/epjc/s10052-008-0798-9.
- [27] F. Cooper, G. Frye, Single-particle distribution in the hydrodynamic and statistical

- thermodynamic models of multiparticle production, *Phys. Rev. D* 10 (1974) 186–189. doi:10.1103/PhysRevD.10.186.
- [28] S. Bass, M. Belkacem, M. Bleicher, M. Brandstetter, L. Bravina, C. Ernst, L. Gerland, M. Hofmann, S. Hofmann, J. Konopka, G. Mao, L. Neise, S. Soff, C. Spieles, H. Weber, L. Winckelmann, H. StÄ¶cker, W. Greiner, C. Hartnack, J. Aichelin, N. Amelin, Microscopic models for ultrarelativistic heavy ion collisions, *Progress in Particle and Nuclear Physics* 41 (1998) 255 – 369. doi:10.1016/S0146-6410(98)00058-1.
- [29] J. Weil, V. Steinberg, J. Staudenmaier, L. G. Pang, D. Oliinychenko, J. Mohs, M. Kretz, T. Kehrenberg, A. Goldschmidt, B. BÄuchle, J. Auvinen, M. Attems, H. Petersen, Particle production and equilibrium properties within a new hadron transport approach for heavy-ion collisions, *Phys. Rev. C* 94 (2016) 054905. doi:10.1103/PhysRevC.94.054905.
- [30] ALICE Collaboration, Centrality determination of Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV with ALICE, *Phys. Rev. C* 88 (2013) 044909. doi:10.1103/PhysRevC.88.044909.
- [31] U. Heinz, The strongly coupled quark-gluon plasma created at RHIC, *Journal of Physics A: Mathematical and Theoretical* 42 (21) (2009) 214003.
URL <http://stacks.iop.org/1751-8121/42/i=21/a=214003>
- [32] [link].
URL <http://ars.els-cdn.com/content/image/1-s2.0-S0146641008001336-gr1.jpg>
- [33] C. Gale, S. Jeon, B. Schenke, P. Tribedy, R. Venugopalan, Event-by-event anisotropic flow in heavy-ion collisions from combined yang-mills and viscous fluid dynamics, *Phys. Rev. Lett.* 110 (2013) 012302. doi:10.1103/PhysRevLett.110.012302.
- [34] T. Schornder-Sadenius, et al., *The Large Hadron Collider: Harvest of Run 1*, Springer, 2015.

- [35] ALICE Collaboration, Anisotropic flow of identified particles in Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV.
URL <https://arxiv.org/abs/1805.04390>
- [36] ALICE Collaboration, J/ψ suppression at forward rapidity in Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV, Physics Letters B 766 (2017) 212–224. doi:10.1016/j.physletb.2016.12.064.
- [37] B. Müller, J. Schukraft, B. Wyslouch, First results from pb+pb collisions at the lhc, Annual Review of Nuclear and Particle Science 62 (1) (2012) 361–386. doi:10.1146/annurev-nucl-102711-094910.
- [38] The CMS collaboration, Observation of long-range, near-side angular correlations in proton-proton collisions at the lhc, Journal of High Energy Physics 2010 (9) (2010) 91. doi:10.1007/JHEP09(2010)091.
- [39] ALICE Collaboration, Enhanced production of multi-strange hadrons in high-multiplicity proton-proton collisions, Nat Phys 13 (6) (2017) 535–539. doi:10.1038/nphys4111.
- [40] J. Casalderrey-Solana, Z. Hulcher, G. Milhano, D. Pablos, K. Rajagopal, A simultaneous description of hadron and jet suppression in heavy ion collisions.
URL <https://arxiv.org/abs/1808.07386>
- [41] M. Cacciari, G. Salam, G. Soyez, The catchment area of jets, JHEP 005 (0804).
- [42] M. Cacciari, G. Salam, G. Soyez, The anti- k_T jet cluster algorithm, JHEP 063 (0804).
- [43] D. d’Enterria, Jet quenching, Springer-Verlag. doi:10.1007/978-3-642-01539-7_16.
- [44] K. M. Burke, A. Buzzatti, N. Chang, C. Gale, M. Gyulassy, U. Heinz, S. Jeon, A. Majumder, B. Müller, G.-Y. Qin, B. Schenke, C. Shen, X.-N. Wang, J. Xu, C. Young, H. Zhang, Extracting the jet transport coefficient from jet quenching in high-energy heavy-ion collisions, Phys. Rev. C 90 (2014) 014909. doi:10.1103/PhysRevC.90.014909.

- [45] M. Spousta, B. Cole, Interpreting single jet measurements in Pb–Pb collisions at the LHC, *The European Physical Journal C* 76 (2) (2016) 50. doi:10.1140/epjc/s10052-016-3896-0.
- [46] ALICE Collaboration, Constraints on jet quenching in p–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV measured by the event-activity dependence of semi-inclusive hadron-jet distributions, *Physics Letters B* 783 (2018) 95 – 113. doi:10.1016/j.physletb.2018.05.059.
- [47] CMS Collaboration, Observation and Studies of Jet Quenching in Pb–Pb Collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV, *Phys. Rev. C* 84 (024906).
- [48] ATLAS Collaboration, Measurement of jet fragmentation in Pb–Pb and pp collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV with the ATLAS detector, *Phys. Rev. C* 98 (2018) 024908. doi:10.1103/PhysRevC.98.024908.
- [49] ALICE Collaboration, Measurement of jet quenching with semi-inclusive hadron-jet distributions in central Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV, *Journal of High Energy Physics* 2015 (9) (2015) 170. doi:10.1007/JHEP09(2015)170.
- [50] ALICE preliminary results, Inclusive measurement of z_g in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV ($R_g < 0.1$).
URL <http://alice-figure.web.cern.ch/node/13214>
- [51] ALICE preliminary results, Inclusive measurement of z_g in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV ($R_g > 0.2$).
URL <http://alice-figure.web.cern.ch/node/13213>
- [52] M. Dasgupta, A. Fregoso, S. Marzani, G. P. Salam, Towards an understanding of jet substructure, *Journal of High Energy Physics* 2013 (9) (2013) 29. doi:10.1007/JHEP09(2013)029.
- [53] A. J. Larkoski, S. Marzani, G. Soyez, J. Thaler, Soft drop, *Journal of High Energy Physics* 2014 (5) (2014) 146. doi:10.1007/JHEP05(2014)146.

- [54] The ALICE collaboration, Medium modification of the shape of small-radius jets in central Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV, Journal of High Energy Physics 2018 (10) (2018) 139. doi:10.1007/JHEP10(2018)139.
- [55] The ALICE collaboration, First measurement of jet mass in pb–pb and p–pb collisions at the lhc, Physics Letters B 776 (2018) 249 – 264. doi:<https://doi.org/10.1016/j.physletb.2017.11.044>.
- [56] ALICE Collaboration, The ALICE experiment at the CERN LHC, J. Instrum. 3 (S08002).
- [57] ALICE Collaboration.
URL <http://alice-collaboration.web.cern.ch>
- [58] ALICE Collaboration, Performance of the ALICE Experiment at the CERN LHC, Int. J. Mod. Phys. A29 (1430044).
- [59] LHC Machine, The CERN large hadron collider: accelerator and experiments, J. Instrum. 3 (S08001).
- [60] CERN accelerator complex.
URL <http://public-archive.web.cern.ch/public-archive/en/research/AccelComplex-en.html>
- [61] ALICE Collaboration, Addendum to the Technical Design Report for the Upgrade of the ALICE Time Projection Chamber. (2015).
URL <http://cds.cern.ch/record/1984329>
- [62] ALICE Collaboration, ALICE electromagnetic calorimeter technical design report CERN-ALICE-TDR-014, CERN-LHCC-2008-014.
URL <http://inspirehep.net/record/794183>
- [63] ALICE Collaboration, Technical Design Report for the Upgrade of the ALICE Time Projection Chamber ALICE-TDR-016, CERN-LHCC-2013-020.
URL <http://cds.cern.ch/record/1622286>

- [64] ALICE Collaboration, Technical Design Report for the Upgrade of the ALICE Inner Tracking System ALICE-TDR-017, CERN-LHCC-2013-024.
URL <https://cds.cern.ch/record/1625842>
- [65] ALICE Collaboration, Technical design report for the upgrade of the online-offline computing system ALICE-TDR-019, CERN-LHCC-2015-006.
URL <http://cds.cern.ch/record/2011297>
- [66] ALICE Collaboration, Technical Design Report for the Muon Forward Tracker ALICE-TDR-018, CERN-LHCC-2015-001.
URL <https://cds.cern.ch/record/1981898>
- [67] ALICE Collaboration, Technical Design Report for the Upgrade of the ALICE Read-out and Trigger System ALICE-TDR-015, CERN-LHCC-2013-019.
URL <http://cds.cern.ch/record/1603472>
- [68] ALICE TPC Collaboration, Particle identification studies with a full-size 4-GEM prototype for the ALICE TPC upgrade, Nucl. Inst. Meth. A 903.
- [69] F. Sauli, GEM: a new concept for electron amplification in gas detectors, Nucl. Inst. Meth. A 386.
- [70] Y. Giomataris, P. Rebourgeard, J. Robert, G. Charpak, Micromegas: a high-granularity position-sensitive gaseous detector for high particle-flux environments, Nucl. Inst. Meth. A 376.
- [71] S. Aiola, R.J. Ehlers, S. Gu, J.W. Harris, R. Majka, J.D. Mulligan, M. Oliver, J. Schambach, N. Smirnov, Combination of two gas electron multipliers and a micromegas as gain elements for a time projection chamber, Nucl. Inst. Meth. A 834. doi:10.1016/j.nima.2016.08.007.
- [72] H. Wieman, Gating grid concept for ALICE upgrade (2014).
URL https://wiki.bnl.gov/eic/upload/Alice_upgrade_gating_grid_idea.pdf
- [73] J. Mulligan, Simulations of a multi-layer extended gating grid.
URL <https://arxiv.org/abs/1603.05648>

- [74] ALICE Collaboration, Measurement of jet suppression in central Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV, Physics Letters B 1 (746).
URL [10.1016/j.physletb.2015.04.039](https://arxiv.org/abs/10.1016/j.physletb.2015.04.039)
- [75] ATLAS Collaboration, Measurement of the jet radius and transverse momentum dependence of inclusive jet suppression in lead-lead collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV with the ATLAS detector, Physics Letters B 719 (4) (2013) 220 – 241. doi:
[10.1016/j.physletb.2013.01.024](https://arxiv.org/abs/10.1016/j.physletb.2013.01.024).
- [76] CMS Collaboration, Measurement of inclusive jet cross sections in pp and Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV, Phys. Rev. C 96 (2017) 015202. doi:[10.1103/PhysRevC.96.015202](https://arxiv.org/abs/10.1103/PhysRevC.96.015202).
- [77] ATLAS Collaboration, Measurement of the nuclear modification factor for inclusive jets in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV with the ATLAS detector.
URL <https://arxiv.org/abs/1805.05635>
- [78] ALICE Collaboration, Measurement of charged jet suppression in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV, JHEP 013 (03).
URL [10.1007/JHEP03\(2014\)013](https://arxiv.org/abs/10.1007/JHEP03(2014)013)
- [79] ALICE preliminary results, Measurement of charged jet spectra in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV with ALICE at LHC (update including high interaction Pb-Pb runs).
URL <https://alice-figure.web.cern.ch/node/9760>
- [80] ALICE Collaboration, Measurement of event background fluctuations for charged particle jet reconstruction in Pb–Pb collisions at 2.76 TeV, Journal of High Energy Physics 2012 (3) (2012) 53. doi:[10.1007/JHEP03\(2012\)053](https://arxiv.org/abs/10.1007/JHEP03(2012)053).
- [81] R. Ma, Jets measurements in pp and Pb–Pb collisions in ALICE, Ph.D. thesis.
URL http://rhig.physics.yale.edu/Theses/Ma_Thesis.pdf
- [82] A. Hocker, V. Kartvelishvili, SVD approach to data unfolding, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors

- and Associated Equipment 372 (3) (1996) 469 – 481. doi:10.1016/0168-9002(95)01478-0.
- [83] ALICE Collaboration, The ALICE definition of primary particles.
URL <https://cds.cern.ch/record/2270008/files/cds.pdf>
- [84] ALICE Collaboration, Transverse momentum spectra and nuclear modification factors of charged particles in pp, p-Pb and Pb-Pb collisions at the LHC.arXiv:1802.09145.
- [85] [link].
URL <http://hepunx.rl.ac.uk/~adye/software/unfold/RooUnfold.html>
- [86] G. D’Agostini, A multidimensional unfolding method based on bayes’ theorem, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 362 (2) (1995) 487 – 498. doi:10.1016/0168-9002(95)00274-X.
- [87] ALICE Collaboration, ALICE luminosity determination for pp collisions at $\sqrt{s_{\text{NN}}} = 5$ TeV.
URL <https://cds.cern.ch/record/2202638/>
- [88] S. Alioli, K. Hamilton, P. Nason, C. Oleari, E. Re, Jet pair production in POWHEG, Journal of High Energy Physics 2011 (4) (2011) 81. doi:10.1007/JHEP04(2011)081.
- [89] The POWHEG BOX user manual: common features.
URL <http://th-www.if.uj.edu.pl/~erichter/POWHEG-BOX-V2/Docs/manual-BOX.pdf>
- [90] The POWHEG BOX V2 framework.
URL <http://th-www.if.uj.edu.pl/~erichter/POWHEG-BOX-V2/Docs/V2-paper.pdf>
- [91] POWHEG Merging.
URL <http://home.thep.lu.se/~torbjorn/pythia82html/POWHEGMerging.html>

- [92] POWHEG-Pythia matching scheme effects in NLO simulation of dijet events.
URL <http://inspirehep.net/record/1480875?ln=en>
- [93] K. Zapp, F. Krauss, U. Wiedemann, A perturbative framework for jet quenching, JHEP 2013 (3).
- [94] R. Elayavalli, K. Zapp, Medium response in JEWEL and its impact on jet shape observables in heavy ion collisions, JHEP 2017 (7).
- [95] R. Hosokawa, private communication.
- [96] R. Elayavalli, K. Zapp, Simulating V+jet processes in heavy ion collisions with JEWEL, EPJ C 76 (12).
- [97] Y. He, T. Luo, X.-N. Wang, Y. Zhu, Linear boltzmann transport for jet propagation in the quark-gluon plasma: Elastic processes and medium recoil, Phys. Rev. C 91 (054908).
- [98] Y. He, S. Cao, W. Chen, T. Luo, L. Pang, X.-N. Wang, Conspiracy behind single inclusive jet suppression in heavy-ion collisions.
URL <https://arxiv.org/abs/1809.02525>
- [99] Z.-B. Kang, F. Ringer, I. Vitev, Inclusive production of small radius jets in heavy-ion collisions, Phys. Lett. B 769 (242).
- [100] H. T. Li, I. Vitev, Inverting the mass hierarchy of jet quenching effects with prompt b-jet substructure.
URL <https://arxiv.org/abs/1801.00008>
- [101] I. Vitev, private communication.
- [102] J. Casalderrey-Solana, D. C. Gulhan, J. G. Milhano, D. Pablos, K. Rajagopal, A hybrid strong/weak coupling approach to jet quenching, JHEP 019 (1410).
- [103] J. Casalderrey-Solana, D. C. Gulhan, J. G. Milhano, D. Pablos, K. Rajagopal, Predictions for boson-jet observables and fragmentation function ratios from a hybrid strong/weak coupling model for jet quenching, JHEP 053 (1603).

- [104] J. Casalderrey-Solana, D. C. Gulhan, J. G. Milhano, D. Pablos, K. Rajagopal, Angular structure of jet quenching within a hybrid strong/weak coupling model, JHEP 135 (1703).
- [105] Z. Hulcher, D. Pablos, K. Rajagopal, Resolution effects in the hybrid strong/weak coupling model, JHEP 010 (1803).
- [106] D. Pablos, private communication.
- [107] [link].
URL <http://jetscape.wayne.edu/jetscape/>
- [108] F. D’Eramo, K. Rajagopal, Y. Yin, Molière scattering in quark-gluon plasma: Finding point-like scatterers in a liquid.
URL <https://arxiv.org/abs/1808.03250>
- [109] H. Andrews, et al., Novel tools and observables for jet physics in heavy-ion collisions.
URL <https://arxiv.org/abs/1808.03689>
- [110] STAR.
URL <https://www.star.bnl.gov>
- [111] ANSYS.
URL <http://www.ansys.com>
- [112] Garfield++.
URL <http://garfieldpp.web.cern.ch>
- [113] L. Viehland, E. Mason, Transport properties of gaseous ions over a wide energy range, part iv, Atomic Data and Nuclear Data Tables 60 (37-95) (1995) p. 71.
- [114] Y. Kaneko, T. Koizumi, N. Kobayashi, J. Phys. Soc. Japan 43 (1817).
- [115] W. Blum, W. Riegler, L. Rolandi, Particle Detection With Drift Chambers, Springer-Verlag, 2008.
- [116] R. Veenhof, Calculations for the ALICE TPC read-out.
URL http://rjd.home.cern.ch/rjd/Alice/neon_mobility.html