

2

A heavy ion phenomenology primer

What macroscopic properties of matter emerge from the fundamental constituents and interactions of a non-Abelian gauge theory? The study of ultra-relativistic heavy ion collisions addresses this question for the theory of the strong interaction, Quantum Chromodynamics, in the regime of extreme energy density. To do this, heavy ion phenomenologists employ tools developed to identify and quantify collective phenomena in collisions that have many thousands of particles in their final states. Generically speaking, these tools quantify deviations with respect to benchmark measurements (for example in proton–proton and proton–nucleus collisions) in which collective effects are absent. In this chapter, we provide details for three cases of current interest: (i) the characterization of azimuthally anisotropic flow, which teaches us how soon after the collision matter moving collectively is formed and which allows us to constrain the value of the shear viscosity of this matter; (ii) the characterization of jet quenching, which teaches us how this matter affects and is affected by a high-velocity colored particle plowing through it; and (iii) the characterization of the suppression of quarkonium production, which has the potential to teach us about the temperature of the matter and of the degree to which it screens the interaction between colored particles.

2.1 General characteristics of heavy ion collisions

In a heavy ion collision experiment, large nuclei, such as gold (at RHIC) or lead (at the CERN SPS and LHC), are collided at an ultra-relativistic center of mass energy \sqrt{s} . The reason for using large nuclei is to create as large a volume as possible of matter at a high energy density, to have the best chance of discerning phenomena or properties that characterize macroscopic amounts of strongly interacting matter. In contrast, in energetic elementary collisions (say electron–positron collisions but to a good approximation also in proton–proton collisions) one may find many hadrons in the final state but these are understood to result from a few initial partons that

each fragment rather than from a macroscopic volume of interacting matter. Many years ago Phil Anderson coined the phrase “more is different” to emphasize that macroscopic volumes of (in his case condensed) matter manifest qualitatively new phenomena, distinct from those that can be discerned in interactions among few elementary constituents and requiring distinct theoretical methods and insights for their elucidation [53]. Heavy ion physicists do not have the luxury of studying systems containing a mole of quarks, but by using the heaviest ions that nature provides they go as far in this direction as is possible.

The purpose of building accelerators that achieve heavy ion collisions at higher and higher \sqrt{s} is simply to create matter at higher and higher energy density. A simple argument to see why this may be so arises upon noticing that in the center-of-mass frame we have the collision of two Lorentz-contracted nuclei, pancake-shaped, and increasing the collision energy makes these pancakes thinner. Thus, at $t = 0$ when these pancakes are coincident the entire energy of the two incident nuclei is found within a smaller volume for higher \sqrt{s} . This argument is overly simple, however, because not all of the energy of the collision is transformed into the creation of matter; much of it is carried by the debris of the two colliding nuclei that spray almost along the beam directions.

The question of how the initial state wave function of the colliding nuclei determines precisely how much matter, containing how much entropy, is produced soon after the collision, and consequently determines the number of particles in the final state, is a subject of intense theoretical interest. We shall not describe this branch of heavy ion phenomenology in any detail, but it is worth having a quantitative sense of just how many particles are produced in a typical heavy ion collision. In Fig. 2.1 we show the multiplicity of charged particles per unit pseudorapidity for RHIC collisions at four different values of \sqrt{s} . Recall that the pseudorapidity η is related to the polar angle θ measured with respect to the beam direction by $\eta = -\log \tan(\theta/2)$. Note also that, by convention, the incident ions in these collisions have a velocity such that individual nucleons colliding with that velocity would collide with a center of mass energy of \sqrt{s} . Since each gold nucleus has 197 nucleons and each Pb nucleus has 208 nucleons, the total center of mass energy in a heavy ion collision at the top RHIC energy is about 40 TeV and it rises to about 600 TeV at the current LHC energy. By integrating under the curve in Fig. 2.1, one finds that a heavy ion collision at top RHIC energy yields 5060 ± 250 charged particles [94, 95]. The multiplicity measurement is made by counting tracks, meaning that neutral particles (like π^0 s and the photons they decay into) are not counted. So, the total number of hadrons is greater than the total number of charged particles. If all the hadrons in the final state were pions, and if the small isospin breaking introduced by the different number of protons and neutrons in a gold nucleus can be neglected, there would be equal numbers of π^+ , π^- and π^0 meaning that the total




Figure 2.1 Charged particle multiplicity distributions for central nucleus–nucleus collisions (i.e. the 5% or 6% of collisions that have the smallest impact parameter) over more than two orders of magnitude in $\sqrt{s_{NN}}$. Data taken from Refs. [263] and [94].

multiplicity would be $3/2$ times the charged multiplicity. In reality, this factor turns out to be about 1.6 [96], meaning that heavy ion collisions at the top RHIC energy each produce about 8000 hadrons in the final state. At the LHC, the corresponding pseudorapidity distribution is known so far only in a range around mid-rapidity (see Fig. 2.1), with $dN_{ch}/d\eta = 1584 \pm 4(\text{stat}) \pm 76(\text{sys})$ at $\eta = 0$ in the 5% or 6% of collisions with $\sqrt{s} = 2.76$ TeV that have the smallest impact parameter [4]. We see from Fig. 2.1 that this multiplicity grows with increasing collision energy by a factor of close to 2.5 from the top RHIC energy to LHC at $\sqrt{s} = 2.76$ GeV. The multiplicity per unit pseudorapidity is largest in a range of angles centered around $\eta = 0$, meaning $\theta = \pi/2$. Moreover, the distribution extends with increasing center of mass energy to larger values of pseudorapidity, so that the total event multiplicity at LHC is estimated to be a factor ~ 5 larger than at RHIC, lying in the ballpark of $\sim 25\,000$ charged particles in central collisions. The illustrations in Fig. 2.2 provide an impression of what collisions with these multiplicities look like.

The large multiplicities in heavy ion collisions indicate large energy densities, since each of these particles carries a typical (mean) transverse momentum of several hundred MeV. There is a simple geometric method due to Bjorken [165], that can be used to estimate the energy density at a fiducial early time, conventionally




Figure 2.2 Event displays illustrating heavy ion collisions as seen by the STAR detector at RHIC (upper panel) and the ALICE detector at the LHC (lower panel). Nuclei (gold above; lead below) collided at the center of each image, and the resulting tracks made by those charged particles produced in the collision that pass through the STAR and ALICE time-projection chambers and the ALICE inner tracker are shown, projected onto the page in the upper image and in perspective in the lower image. Figures courtesy of Brookhaven National Laboratory (above) and the ALICE Collaboration and CERN (below).

chosen to be $\tau_0 = 1$ fm. The smallest reasonable choice of τ_0 would be the thickness of the Lorentz-contracted pancake-shaped nuclei, for instance $\sim (14 \text{ fm})/107$ at RHIC since gold nuclei have a radius of about 7 fm and the Lorentz factor is set by energy of the incident nucleons and their mass in the center-of-mass frame,

$\gamma \sim m_N/E$. But, at these early times ~ 0.1 fm the matter whose energy density one would be estimating would still be far from equilibrium. We shall see below that data on azimuthally anisotropic flow indicate that by ~ 1 fm after the collision, matter is flowing collectively like a fluid in local equilibrium. The geometric estimate of the energy density is agnostic about whether the matter in question is initial state partons that have not yet interacted and are far from equilibrium or matter in local equilibrium behaving collectively; because we are interested in the latter, we choose $\tau_0 = 1$ fm. Bjorken's geometric estimate can be written as

$$\varepsilon_{Bj} = \left. \frac{dE_T}{d\eta} \right|_{\eta=0} \frac{1}{\tau_0 \pi R^2}, \quad (2.1)$$

where $dE_T/d\eta$ is the transverse energy $\sqrt{m^2 + p_T^2}$ of all the particles per unit rapidity and $R \approx 7$ fm is the radius of the nuclei. The logic is simply that at time τ_0 the energy within a volume $2\tau_0$ in longitudinal extent between the two receding pancakes and πR^2 in transverse area must be at least $2dE_T/d\eta$, the total transverse energy between $\eta = -1$ and $\eta = +1$. At RHIC with $dE_T/d\eta \approx 800$ GeV [95], we obtain $\varepsilon_{Bj} \approx 5$ GeV/fm³. In choosing the volume in the denominator in the estimate (2.1) we neglected transverse expansion because $\tau_0 \ll R$. But, there is clearly an arbitrariness in the range of η used; if we had included particles produced at higher pseudorapidity (closer to the beam directions) we would have obtained a larger estimate of the energy density. Note also that there is another sense in which (2.1) is conservative. If there is an epoch after the time τ_0 during which the matter expands as a hydrodynamic fluid, and we shall later see evidence that this is so, then during this epoch its energy density drops more rapidly than $1/\tau$ because as it expands (particularly longitudinally) it is doing work. This means that by using $1/\tau$ to run the clock backwards from the measured final state transverse energy to that at τ_0 we have significantly underestimated the energy density at τ_0 . It is striking that even though we have deliberately been conservative in making this underestimate, we have found an energy density that is about five times larger than the QCD critical energy density $\varepsilon_c \approx 1$ GeV/fm³, where the crossover from hadronic matter to quark–gluon plasma occurs, according to lattice calculations of QCD thermodynamics [129].

As shown in Fig. 2.3, the spectrum in a nucleus–nucleus collision extends to very high momentum, much larger than the mean. However, the multiplicity of high-momentum particles drops very fast with momentum, as a large power of p_T . We may separate the spectrum into two sectors. In the soft sector, spectra drop exponentially with $\sqrt{m^2 + p_T^2}$ as in thermal equilibrium. In the hard sector, spectra drop like power laws in p_T as is the case for hard particles produced by high momentum-transfer parton–parton collisions at $\tau = 0$. The bulk of the particles




Figure 2.3 Charged particle spectrum as function of p_T in Pb+Pb collisions at LHC energy for nearly head-on (the 5% of collisions with the lowest impact parameter) and grazing collisions, compared to the corresponding spectrum in p+p collisions with an appropriately scaled normalization. Figure taken from Ref. [7].

have momenta in the soft sector; hard particles are rare in comparison. The separation between the hard and the soft sectors, which is by no means sharp, lies in the range of a few (say 3–6) GeV.

There are several lines of evidence that indicate that the soft particles in a heavy ion collision, which are the bulk of all the hadrons in the final state, have rescattered many times and come into local thermal equilibrium. The most direct approach comes via the analysis of the exponentially falling spectra of identified hadrons. Fitting a slope to these exponential spectra and then extracting an “effective temperature” for each species of hadron yields different “effective temperatures” for each species. This species dependence arises because the matter produced in a heavy ion collision expands radially in the directions transverse to the beam axis; perhaps explodes radially is a better phrase. This means that we should expect the p_T spectra to be a thermal distribution boosted by some radial velocity. If all hadrons are boosted by the same *velocity*, the heavier the hadron the more its




Figure 2.4 (a) Spectra for identified pions, kaons and protons as a function of p_T in head-on gold–gold collisions at top RHIC energy [25]. (b) Spectra for identified pions, kaons and protons as a function of p_T in (non-single-diffractive) proton–proton collisions at the same energy $\sqrt{s} = 200$ GeV [17].

momentum is increased by the radial boost. Indeed, what is found in data is that the effective temperature increases with the mass of the hadron species. This can be seen at a qualitative level in Fig. 2.4a: in the soft regime, the proton, kaon and pion spectra are ordered by mass, with the protons falling off most slowly with p_T , indicating that they have the highest effective temperature. Quantitatively, one uses the data for hadron species with varying masses to first extract the mass-dependence of the effective temperature, and thus the radial expansion velocity, and then to extrapolate the effective “temperatures” to the mass \rightarrow zero limit, and in this way obtain a measurement of the actual temperature of the final state hadrons. This “kinetic freezeout temperature” is the temperature at the (very late) time at which the gas of hadrons becomes so dilute that elastic collisions between the hadrons cease, and the momentum distributions therefore stop changing as the system expands.

In heavy ion collisions at the top RHIC energy, models of the kinetic freezeout account for the data with freezeout temperatures of ≈ 90 MeV and radial expansion velocities of $0.6c$ for collisions with the smallest impact parameters [16]. With increasing impact parameter, the radial velocity decreases and the freezeout temperature increases. This is consistent with the picture that a smaller system builds up less transverse flow and that during its expansion it cannot cool down as much as a bigger system, since it falls apart earlier.

The analysis just described is unique to heavy ion collisions: in elementary electron–positron or proton–(anti)proton collisions, spectra at low transverse momentum may also be fit by exponentials, but the “temperatures” extracted in this way do not have a systematic dependence on the hadron mass, see Fig. 2.4b. Simply seeing exponential spectra and fitting a “temperature” therefore does not in itself provide evidence for rescattering and equilibration. Making that case in the context of heavy ion collisions relies crucially on the existence of a collective radial expansion with a common velocity for all hadron species.

Demonstrating that the final state of a heavy ion collision at the time of kinetic freezeout is a gas of hadrons in local thermal equilibrium emboldens us to ask whether the material produced in these collisions reaches local thermal equilibrium at an earlier time, and thus at a higher temperature. The best evidence for an affirmative answer to this question comes from the analysis of “elliptic flow” in collisions with nonzero impact parameter. We shall discuss this at length in the next section.

We close this section with a simpler analysis that lays further groundwork by allowing us to see back to a somewhat earlier epoch than that of kinetic freezeout. If we think of a heavy ion collision as a “little bang”, replaying the history of the big bang in a small volume and with a vastly accelerated expansion rate, then kinetic freezeout is the analogue of the (late) cosmological time at which photons and electrons no longer scatter off each other. We now turn to the analogue of the (earlier) cosmological epoch of nucleosynthesis, namely the time at which the composition of the final state hadron gas stops changing. Experimentalists can measure the abundance of more than a dozen hadron species, and it turns out that all the ratios among these abundances can be fit upon assuming thermal distributions with some temperature T and some baryon number chemical potential μ_B , as shown in Fig. 2.5. This is a two parameter fit to about a dozen ratios. The temperature extracted in this way is called the chemical freezeout temperature, since one interpretation is that it is the temperature at which the hadronic matter becomes dilute enough that inelastic hadron–hadron collisions cease to modify the abundance ratios. The chemical freezeout temperature in heavy ion collisions at top RHIC energies is about 155–180 MeV [193, 56]. This is interesting for several reasons. First, it is not far below the QCD phase transition temperature, which




Figure 2.5 So-called thermal fit to different particle species. The relative abundance of different hadron species produced in RHIC collisions at $\sqrt{s} = 200$ GeV is well-described by a two-parameter grand canonical ensemble in terms of a temperature, T , and a chemical potential for baryon number, μ_B [56].

means that the appropriateness of a hadron gas description of this epoch may be questioned. Second, within error bars it is the same temperature that is extracted by doing a thermal model fit to hadron production in electron–positron collisions, in which final state rescattering, elastic or inelastic, can surely be neglected. So, by itself the success of the thermal fits to abundance ratios in heavy ion collisions could be interpreted as telling us about the statistical nature of hadronization, as must be the case in electron–positron collisions. However, given that we know that in heavy ion collisions (and not in electron–positron collisions) kinetic equilibrium is maintained down to a *lower* kinetic freezeout temperature, and given that as we shall see in the next section approximate local thermal equilibrium is achieved at a *higher* temperature, it does seem most natural to interpret the chemical freeze-out temperature in heavy ion collisions as reflecting the temperature of the matter produced at the time when species-changing processes cease.

We have not yet talked about the baryon number chemical potential extracted from the thermal fit to abundance ratios. As illustrated in Fig. 2.6a, this μ_B decreases with increasing collision energy \sqrt{s} . This energy-dependence has two origins. The dominant effect is simply that at higher and higher collision energies more and more entropy is produced, while the total net baryon number in the collision is always 197+197. At top RHIC energies, these baryons are diluted among the 8000 or so hadrons in the final state, making the baryon chemical potential much smaller than it is in lower energy collisions where the final state multiplicity




Figure 2.6 (a) Chemical potential extracted from thermal fits at different center of mass energies [56]. (b) The number of protons minus number of antiprotons per unit rapidity for central heavy ion collisions [132]. This net proton number decreases with increasing center of mass energy from $\sqrt{s} = 5$ GeV (at the AGS collider at BNL), via $\sqrt{s} = 17$ GeV (at the SPS collider at CERN) to $\sqrt{s} = 200$ GeV (at RHIC). (For each collision energy, y_p indicates the rapidity of a hypothetical proton that has the same velocity after the collision as it did before.)

is much lower. The second effect is that, in the highest energy collisions, most of the net baryon number from the two incident nuclei stays at large pseudorapidity (meaning small angles near the incident beam directions). These two effects can be seen directly in the data shown in Fig. 2.1 and Fig. 2.6b: as the collision energy increases, the total number of hadrons in the final state grows while the net baryon

number at mid-rapidity drops.¹ This experimental fact that baryon number is not “fully stopped” teaches us about the dynamics of the earliest moments of a hadron–hadron collision. (In this respect, heavy ion collisions are not qualitatively different than proton–proton collisions.) In a high energy proton–proton collision, particle production at mid-rapidity is dominated by the partons in the initial state that carry a small fraction of the momentum of an individual nucleon – small Bjorken x . And, the small- x parton distribution functions that describe the initial state of the incident nucleons or nuclei are dominated by gluons and to a lesser extent by quark–antiquark pairs; the net baryon number is at larger x .

We shall not focus here on the many interesting questions related to the early-time dynamics in heavy ion collisions. Because QCD is asymptotically free, it is natural to expect that during the earliest moments of a sufficiently energetic heavy ion collision, the physics should not be thought of as strongly coupled. The relevant length scale at the moment of the collision between two highly Lorentz-contracted nuclei is the mean spacing between gluons in the transverse plane (the inverse of this length scale is called the saturation momentum) and in the high collision energy limit this length scale is short and the physics is weakly coupled. The analysis of this weak coupling, but strong field, regime is the subject of much active research that we shall not describe. One of the goals of this effort is to understand how rapidly local thermal equilibrium can be established. We shall see in Chapter 7 that calculations done via gauge/string duality have shed light on this particular question.

In the next section we turn to the evidence that local thermal equilibrium is established quickly, and therefore at a high temperature. This means that heavy ion collisions can teach us about properties of the high temperature phase of QCD, namely the quark–gluon plasma. And, we shall see later, so can calculations done via gauge/string duality. We shall henceforth always work at $\mu_B = 0$. This is a good approximation as long as $\mu_B/3$, the quark chemical potential, is much less than the

¹ The data in Fig. 2.6b are plotted versus rapidity

$$y \equiv \frac{1}{2} \ln \left(\frac{E + p_L}{E - p_L} \right), \quad (2.2)$$

where E and p_L are the energy and longitudinal momentum of a proton in the final state. Recall that rapidity and pseudorapidity $\eta \equiv -\ln \tan(\theta/2) = \frac{1}{2} \ln \left(\frac{p+p_L}{p-p_L} \right)$ (used in the plot in Fig. 2.1) become the same in the limit in which E and p_L are much greater than the proton mass and the three-momentum p approximates E . For smaller particle momenta, however, the transformation between η and y involves a non-trivial Jacobian. As a consequence, the pseudorapidity and rapidity distributions $dN_{\text{ch}}/d\eta$ and dN_{ch}/dy have different shapes. In ultra-relativistic heavy ion collisions, $dN_{\text{ch}}/d\eta$ looks somewhat trapezoidal, with an approximately flat plateau around $\eta \sim 0$ as in Fig. 2.1, while dN_{ch}/dy is closer to Gaussian in shape. In these high energy collisions, it is a reasonable rule of thumb that one can estimate dN_{ch}/dy at $y = 0$ by multiplying $dN_{\text{ch}}/d\eta$ at $\eta = 0$ by about 1.1. When one plots data for all charged hadrons, as in Fig. 2.1, only pseudorapidity can be defined since the rapidity of a hadron with a given polar angle θ depends on the hadron mass. When one plots data for identified protons, pseudorapidity can be converted into rapidity.

temperature T . The results in Fig. 2.6a show that this is a very good approximation at top RHIC energies and at the LHC.

2.2 Flow

2.2.1 Introduction and motivation

The word “flow” refers here to a suite of experimental observables in heavy ion physics that utilize the experimentalists’ ability to select events in which the impact parameter of the collision lies within some specified range and use these events to study how the matter produced in the collision flows collectively. The basic idea is simple. Suppose we select events in which the impact parameter is comparable to the nuclear radius. Now, imagine taking a beam’s eye view of one of these collisions. The two Lorentz-contracted nuclei (think circular “pancakes”) collide only in an “almond-shaped” region, see Fig. 2.7. The fragments of the nuclei outside the almond that did not collide (“spectator nucleons”) fly down the beam pipes. All the few thousand particles at mid-rapidity in the final state must have come from the few hundred nucleon–nucleon collisions that occurred within the almond. If these few thousand hadrons came instead from a few hundred independent nucleon–nucleon collisions, just by the central limit theorem the few thousand final state hadrons would be distributed uniformly in azimuthal angle ϕ (angle around the beam direction). This null hypothesis, which we shall make quantitative below, is ruled out by the data as we shall see. If, on the other hand, the collisions within the almond yield particles that interact, reach local equilibrium, and thus produce some kind of fluid, our expectations for the “shape” of the azimuthal distribution of the final state hadrons is quite different. The hypothesis that is logically the opposite extreme to pretending that the thousands of partons produced in the hundreds of nucleon–nucleon collisions do not see each other is to pretend that what is produced is a fluid that flows according to the laws of ideal, zero viscosity, hydrodynamics, since this extreme is achieved in the limit of zero mean free path. In hydrodynamics, the almond is thought of as a drop of fluid, with zero pressure at its edges and a high pressure at its center. This droplet of course explodes. And, since the pressure gradients are greater across the short extent of the almond than they are across its long direction, the explosion is azimuthally asymmetric. The first big news from the RHIC experimental program, now also seen at LHC energies, was the discovery that these azimuthal asymmetries can be large: the explosions can blast with summed transverse momenta of the hadrons that are twice as large in the short direction of the almond as they are in the long direction. Moreover, while the form of the nuclear overlap is almond-shaped if averaged over many events (left-hand side of Fig. 2.7), the initial distributions of individual collisions are expected to show event-wise fluctuations that deviate from an almond shape (right-hand side




Figure 2.7 Sketch of the collision of two nuclei, shown in the transverse plane perpendicular to the beam. Left: the event-averaged overlap of the two nuclei is limited to an interaction almond with $\phi \rightarrow \phi + \pi$ symmetry in the center of the transverse plane. Figure taken from Ref. [661]. Right: individual collisions show fluctuations around the event-averaged distribution. The yellow and orange circles depict spectator nucleons that do not participate in the collision. Participating nucleons are in violet color. Figure taken from Ref. [588].

of Fig. 2.7). As we shall see, it turns out that ideal hydrodynamics does a surprisingly good job of describing these asymmetric explosions of the matter produced in heavy ion collisions with nonzero impact parameter. Even the deviations from an almond-shaped distribution of final hadronic fragments are consistent with the fluid dynamic propagation of initial event-wise fluctuations like the ones shown in Fig. 2.7. This phenomenological success of fluid dynamics has implications which are sufficiently interesting that they motivate our describing this story in considerable detail over the course of this entire section. We close this introduction with a sketch of these implications.

First, the agreement between data and ideal hydrodynamics teaches us that the shear viscosity η of the fluid produced in heavy ion collisions must be low; η enters in the dimensionless ratio η/s , with s the entropy density, and it is η/s that is constrained to be small. A fluid that is close to the ideal hydrodynamic limit, with small η/s , requires strong coupling between the fluid constituents. Small η/s means that momentum is not easily transported over distances that are long compared to $\sim s^{-1/3}$, which means that there can be no well-defined quasiparticles with long mean free paths in a low viscosity fluid since, if they existed, they would transport momentum and damp out shear flows. No particles with long mean free paths means strongly coupled constituents. We shall return to this implication of the smallness of η/s at many points in this book, including in particular in Section 6.3.

Second, we learn that the strong coupling between partons that results in approximate local equilibrium and fluid flow close to that described by ideal hydrodynamics must set in very soon after the initial collision. If partons moved with significant mean free paths for many fm of time after the collision, delaying

equilibration for many fm, the almond would circularize to a significant degree during this initial period of time and the azimuthal momentum asymmetry generated by any later period of hydrodynamic behavior would be less than that observed. When this argument is made quantitative, the conclusion is that RHIC collisions produce strongly coupled fluid in approximate local thermal equilibrium within close to or even somewhat less than 1 fm after the collision [543]. Reaching approximate local thermal equilibrium and hence hydrodynamic behavior within less than 1 fm after a heavy ion collision has been thought of as “rapid equilibration”, since it is rapid compared to weak coupling estimates [102]. This observation has launched a large effort (that we shall not review) towards explaining equilibration as originating from weakly coupled processes that arise in the presence of the strong color fields that are present in the initial instants of a heavy ion collision. Recent calculations that we shall describe in Chapter 7 indicate, however, that the observed equilibration time may not be so rapid after all. We shall see in that chapter that when initially far-from-equilibrium matter thermalizes in a strongly coupled theory, for a very wide variety of initial states it does so on a time scale that is of order the inverse of the temperature in the final equilibrated state. Furthermore, we shall also see that in a strongly coupled field theory with a dual gravitational description, when two sheets of energy density with a finite thickness collide at the speed of light a hydrodynamic description of the plasma that results becomes reliable only ~ 3 sheet-thicknesses after the collision. And, a Lorentz-contracted incident gold nucleus at RHIC has a maximum thickness of only 0.14 fm. So, if the equilibration processes in heavy ion collisions could be thought of as strongly coupled throughout, perhaps local thermal equilibrium and hydrodynamic behavior would set in even more rapidly than is indicated by the data.

We can begin to see that the circle of ideas that emerge from the analysis of flow data is what makes heavy ion collisions of interest to the broader community of theoretical physicists. These analyses justify the conclusion that only 1 fm after the collision the matter produced can be described by using the language of thermodynamics and hydrodynamics. And, we have already seen that at this early time the energy density is well above the hadron–QGP crossover in QCD thermodynamics which is well-characterized in lattice calculations. This justifies the claim that heavy ion collisions produce quark–gluon plasma. Furthermore, the same analyses teach us that this quark–gluon plasma is a strongly coupled, low viscosity, fluid with no quasiparticles having any significant mean free path. Lattice calculations have recently begun to cast some light on these transport properties of quark–gluon plasma, but these lattice calculations that go beyond Euclidean thermodynamics are still in their pioneering epoch. Perturbative calculations of quark–gluon plasma properties are built upon the existence of quasiparticles. The analyses of

elliptic flow data thus cast doubt upon their utility. And, we are motivated to study the strongly coupled plasmas with similar properties that can be analyzed via gauge/gravity duality, since these calculational methods allow many questions that go beyond thermodynamics to be probed rigorously at strong coupling.

2.2.2 Relating flow observables to spatial asymmetries

We want to study the dependence of collective flow in heavy ion collisions on the size and anisotropy of the nuclear overlap in the transverse plane, as seen in the qualitative beam's eye view sketch in Fig. 2.7. To this end, it is obviously necessary to bin heavy ion collisions as a function of this impact parameter. This is possible in heavy ion collisions, since the number of hadrons produced in a heavy ion collision is anticorrelated with the impact parameter of the collision. For head-on collisions (conventionally referred to as “central collisions”) the multiplicity is high; the multiplicity is much lower in collisions with impact parameters comparable to the radii of the incident ions (often referred to as “semi-peripheral collisions”); the multiplicity is lower still in grazing (“peripheral”) collisions. Experimentalists therefore bin their events by multiplicity, using that as a proxy for impact parameter. The terminology used refers to the “0%–5% centrality bin” and the “5%–10% centrality bin and . . . , meaning the 5% of events with the highest multiplicities, the next 5% of events with the next highest multiplicity, The correlation between event multiplicity and impact parameter is described well by the so-called Glauber theory of multiple scattering [156], which we shall not review here. Suffice to say that even though the absolute value of the event multiplicities is the subject of much ongoing research, the question of what distribution of impact parameters corresponds to the 0%–5% centrality bin (namely the most head-on collisions) is well established. Although experimentalists cannot literally pick a class of events with a single value of the impact parameter, by binning their data in multiplicity they can select a class of events with a reasonably narrow distribution of impact parameters centered around any desired value. This is possible only because nuclei are big enough: in proton–proton collisions, which in principle have impact parameters since protons are not pointlike, there is no operational way to separate variations in impact parameter from event-by-event fluctuations in the multiplicity at a given impact parameter.

Suppose that we have selected a class of semi-peripheral collisions. Since these collisions have a nonzero impact parameter, the impact parameter vector together with the beam direction define a plane, conventionally called the reaction plane. The event-averaged almond-shaped nuclear overlap depicted in Fig. 2.7 is then often characterized roughly in terms of averages of the initial transverse energy density $\rho(x, y)$

$$\epsilon_2 e^{2i\Psi_2} \equiv -\frac{\{r^2 e^{2i\phi}\}}{\{r^2\}}, \quad \{\dots\} \equiv \frac{\int dx dy \rho(x, y) \dots}{\int dx dy \rho(x, y)}. \quad (2.3)$$

Here, ϵ_2 and Ψ_2 denote the standard participant eccentricity and the participant plane, respectively. They have a straightforward interpretation as characterizing the azimuthal orientation and eccentricity $\epsilon_2 = \frac{\{y^2\} - \{x^2\}}{\{y^2\} + \{x^2\}}$ of the ellipsoid that best fits the initial transverse energy density distribution (where x' , y' denote transverse coordinates along the main axes of the ellipsoid). However, since event-wise fluctuations can lead to significant deviations from an elliptic shape, the elliptic eccentricity ϵ_2 and second order reaction plane Ψ_2 are in general not sufficient. For a more complete characterization of spatial eccentricities from fluctuations in the initial state, one defines

$$\epsilon_n e^{in\Psi_n} \equiv -\frac{\{r^n e^{in\phi}\}}{\{r^n\}}. \quad (2.4)$$

For a large class of semi-peripheral collisions, the elliptic coefficient ϵ_2 will naturally characterize the dominant spatial asymmetry, as it captures the main features of the almond-like shape of the event-averaged nuclear overlap. However, in more central, almost head-on, collisions, when the event-averaged nuclear overlap shows only small azimuthal asymmetries, higher order terms characterizing fluctuations, and in particular ϵ_3 , can be of the same magnitude if not larger than ϵ_2 . We note that spatial eccentricities are typically defined in a coordinate system that is shifted in the transverse plane such that (2.4) vanishes for $n = 1$. This does not imply that the distribution $\rho(x, y)$ cannot have non-vanishing first moments. It just indicates that the ansatz (2.4) is too limited to characterize them. A complete ansatz could be based for instance on the two-parameter set of moments $\epsilon_{n,m} \equiv -\{r^n e^{im\phi}\}/\{r^n\}$ that contains the subset $\epsilon_n \equiv \epsilon_{n,n}$ of (2.4). In this framework, the components $\epsilon_{n,1}$, $n \neq 1$ would characterize first harmonics of the spatial distribution, see e.g. Ref. [778]. A discussion of such refined characterizations of $\rho(x, y)$ lies beyond our scope.

The central question is now how the dynamics of relativistic heavy ion collisions propagates the spatial eccentricities of the initial energy density distribution into the observable momentum spectra. More specifically, as the azimuthal directions within the transverse plane of Fig. 2.7 are not equivalent, we can ask for example to what extent the multiplicity and momentum of hadrons flying across the short direction of the collision almond (in the reaction plane) differs from that of the hadrons flying along the long direction of the collision almond (perpendicular to the reaction plane). And if the initial nuclear overlap shows a significant triangularity ϵ_3 , or higher moments ϵ_4 , ϵ_5 , ϵ_6 , \dots , we can ask which imprints these have on the measured spectra.

To address this question, we characterize now the dependence on the reaction plane for the case of the single inclusive particle spectrum $dN/d^3\mathbf{p}$ of a particular species of hadron. The three-momentum \mathbf{p} of a particle of mass m is parametrized conveniently in terms of its transverse momentum p_T , its azimuthal angle ϕ , and its rapidity y which specifies its longitudinal momentum. Specifically,

$$\mathbf{p} = \left(p_T \cos \phi, p_T \sin \phi, \sqrt{p_T^2 + m^2} \sinh y \right). \quad (2.5)$$

The energy of the particle is $E = \sqrt{p_T^2 + m^2} \cosh y$. The single particle spectrum can then be written as

$$\frac{dN}{d^2\mathbf{p}_t dy} = \frac{1}{2\pi p_T} \frac{dN}{dp_T dy} [1 + 2v_1 \cos(\phi - \Psi_1) + 2v_2 \cos 2(\phi - \Psi_2) + \dots], \quad (2.6)$$

where the Ψ_n denote explicitly the azimuthal orientations of the corresponding flow component in the transverse plane. Thus, the azimuthal dependence of particle production is characterized by the harmonic coefficients

$$v_n \equiv \langle \exp [i n (\phi - \Psi_n)] \rangle = \frac{\int \frac{dN}{d^3\mathbf{p}} e^{i n (\phi - \Psi_n)} d^3 p}{\int \frac{dN}{d^3\mathbf{p}} d^3 p}. \quad (2.7)$$

The coefficients v_n are referred to generically as n th order flow. In particular, v_1 is referred to as “directed flow”, v_2 as “elliptic flow”, and v_3 as “triangular flow”. In general, the v_n can depend on the transverse momentum p_T , the rapidity y , the impact parameter of the collision, and they can differ for different particle species.

We can now make our question about the relation between flow observables and spatial asymmetries of the initial energy density more precise. We ask how the flow harmonics v_n depend on the spatial eccentricities ϵ_n of the initial transverse energy density distribution. We have two principal reasons to limit this discussion to the moments $n \geq 2$. First, as mentioned already, defining first moments of the spatial distribution would require going beyond the ansatz (2.4). Second, the measured v_1 is known to be sensitive not only to medium response, but also to global constraints from energy–momentum conservation. For instance, if the total momentum of all particles in some rapidity window (in some p_T range) points along $\phi = 0$ and defines a positive v_1 , energy–momentum conservation implies that it must point in some other rapidity window (in some other p_T range) along $\phi = \pi$, corresponding to a negative v_1 . In short, the relation between first moments of the spatial eccentricities of the initial energy density distribution and the observable momentum spectra is complicated by confounding factors. In principle, these can be analyzed and controlled, but that requires a more extended analysis than we present that is not yet standard in comparisons between measurements of v_1 and fluid dynamic




Figure 2.8 Transverse momentum dependence of the elliptic flow $v_2(p_T)$ for different centrality bins. Measurements made by the ALICE Collaboration at the LHC (colored points) are compared with parametrized data from the STAR Collaboration at RHIC (gray shaded bands). We see v_2 increasing as one goes from nearly head-on collisions to semi-peripheral collisions. Figure taken from Ref. [5].

simulations of the type we shall discuss below. We shall therefore only discuss the dynamical understanding of how the ϵ_n are related to the v_n for the moments with $n \geq 2$. We shall first consider an event-averaged almond-shaped nuclear overlap zone (left-hand side of Fig. 2.7), before we turn to a discussion of the novel opportunities arising from a study of event-by-event fluctuations (like those illustrated on the right-hand side of Fig. 2.7).

A Discussion for event-averaged spatial asymmetries

In Fig. 2.8, we show data for the transverse momentum dependence of the elliptic flow $v_2(p_T)$ measured for different centrality classes in Au+Au collisions at RHIC and in Pb+Pb collisions at the LHC. It is striking that the $v_2(p_T)$ measured at $\sqrt{s} = 2.76$ TeV by ALICE in three different impact parameter bins agrees within error bars at all values of p_T with that measured at $\sqrt{s} = 200$ GeV by the STAR collaboration at RHIC out to beyond 4 GeV in p_T . On a qualitative level, this indicates that the quark-gluon plasma produced at the LHC is comparably strongly coupled, with comparably small η/s , to that produced and studied at RHIC.

Heavy ion collisions at both RHIC and the LHC feature large azimuthal asymmetries. To appreciate the size of the measured elliptic flow signal, we read from (2.6) that the ratio of $dN/d^3\mathbf{p}$ in whatever azimuthal direction it is largest to $dN/d^3\mathbf{p}$ ninety degrees in azimuth away is $(1 + 2v_2)/(1 - 2v_2)$, which is a factor

of 2 for $v_2 = 1/6$. Thus, a v_2 of the order of magnitude seen in semi-peripheral collisions at RHIC and LHC for $p_T \sim 2$ GeV, as illustrated in Fig. 2.8, corresponds to collisions that are azimuthally asymmetric by more than a factor of 2. In addition to being large, this flow signal displays a characteristic centrality dependence, as we discuss now. The azimuthal asymmetry v_2 of the final state single inclusive hadron spectrum is maximal in semi-peripheral collisions. v_2 is less for more central collisions. Therefore, the measured elliptic flow v_2 traces the event-averaged spatial eccentricity of the initial condition at least qualitatively: the initial event-averaged geometric asymmetry is less for more central collisions since the almond-shaped collision region becomes closer to circular as the impact parameter is reduced.

One can make the relation between spatial ellipticity and measured elliptic flow more quantitative by modeling the elliptic eccentricity ϵ_2 of the spatial energy density distributions, sketched for example in Fig. 2.7. While v_2 is measured directly, the value of ϵ_2 will have some model uncertainty. It turns out, however, that this uncertainty is relatively small, and ϵ_2 is determined predominantly by the impact parameter of the collision which in turn is constrained by the event multiplicity which is directly measurable. As a consequence, one finds strong support for a model-independent picture according to which the p_T -averaged elliptic flow v_2 traces the initial elliptic eccentricity ϵ_2 . For two different models of ϵ_2 , this is shown in the upper panel of Fig. 2.9. Elliptic flow and initial elliptic eccentricity show an approximately linear relation for different centrality classes

$$v_2 \propto \epsilon_2. \quad (2.8)$$

Here, “approximate” means that the proportionality factor differs by less than a factor 2 as a function of centrality and model-dependent uncertainties. In the idealized case of zero impact parameter and vanishing initial state fluctuations, v_2 due to collective effects should vanish. The reason that v_2 is not even smaller in the sample of the 5% most head-on collisions is that this sample includes events with a distribution of impact parameters in the range $0 < b < 3.5$ fm. Moreover, event-by-event fluctuations can introduce ellipticity even in the most central, head-on, collisions. We turn in the next subsection to experimental information about these fluctuations.

We mention as an aside that there are azimuthal asymmetries in particle production that are not related to flow. For instance, a jet produced at azimuthal angle ϕ will often recoil against a jet at angle $\sim (\phi + \pi)$. Such a dijet event introduces an azimuthal asymmetry that results obviously from energy–momentum conservation in the transverse plane, and is not related to collective dynamics. There are several techniques to disentangle such sources of asymmetry from the signals of collective dynamics that one is interested in. One option is to use data at high rapidity to determine the Ψ_n s, and then to measure the v_n by applying (2.7) to mid-rapidity data.




Figure 2.9 Top: the centrality dependence of second and third order flow traces approximately the centrality dependence of the corresponding initial spatial eccentricities ϵ_2, ϵ_3 although the precise value of ϵ_n depends on details of modeling (two models ϵ_n^{CGC} and ϵ_n^W are shown). Bottom: the centrality-dependence of p_T -averaged flow harmonics v_2, v_3, v_4 . Triangular flow is finite if measured with respect to the third order reaction plane but vanishes if measured with respect to the second order reaction plane. Figures taken from Ref. [6].

This eliminates the contribution of all statistical fluctuations uncorrelated with the reaction plane unless these fluctuations introduce correlations between particles in the mid-rapidity and high rapidity regions of the detector. An alternative method is to use the fact that particle correlations resulting from a dijet or some other source of microscopic dynamics affect only a small subset of all particles in a collision, while an asymmetry resulting from the response to an initial spatial asymmetry can affect all particles in an event. Therefore, these effects will scale differently with

event multiplicity in the measured particle correlations. By a suitable so-called cumulant analysis of 2-, 4-, and 6- particle correlations, it is then possible to disentangle these two effects [177]. The values $v_2\{2\}$, $v_2\{4\}$ and $v_3\{2\}$, $v_3\{4\}$ shown in Figs. 2.8 and 2.9 refer to such an analysis based on 2- and 4-particle correlations, respectively. For practical purposes, results of second order cumulants can be regarded as being contaminated by a significant ($\sim 20\%$) contribution from effects that do not arise from collective dynamics (“non-flow” effects). That measurements based on fourth order cumulants, like $v_2\{4\}$ and $v_3\{4\}$ are corrected for all non-flow effects can be seen from the fact that these values then agree with those obtained by other complementary techniques for measuring v_2 .

In summary, the empirical observation that $v_2 \propto \epsilon_2$ shows that the azimuthal momentum anisotropies have a geometrical origin. The large value of v_2 provides evidence that the underlying collective dynamics is very efficient in translating initial spatial anisotropies into momentum anisotropies. And, as we shall discuss further below, a strongly coupled fluid that flows with little dissipation is needed to explain such an efficient translation of spatial anisotropies into observed momentum anisotropies.

B Spatial asymmetries including initial event-by-event fluctuations

In the absence of initial event-by-event fluctuations, the collision region of identical nuclei at mid-rapidity is symmetric under $\phi \rightarrow \phi + \pi$ and all odd spatial eccentricities $\epsilon_1, \epsilon_3, \dots$ must vanish. Since dynamics cannot break a $\phi \rightarrow \phi + \pi$ asymmetry of the initial state, all odd flow harmonics should vanish in this case. Measuring odd flow harmonics at mid-rapidity is therefore direct evidence for initial state fluctuations. As depicted on the right-hand side of Fig. 2.7, fluctuations in the initial conditions of individual nucleus–nucleus collisions can break the $\phi \rightarrow \phi + \pi$ symmetry of the event-averaged almond-like nuclear overlap. Recently, significant experimental evidence has accumulated that these fluctuations around the event-averaged distribution are themselves propagated into event-by-event fluctuations of the final-state momentum anisotropies as the fluid produced in heavy ion collisions expands. Here, we discuss these data and the promise that they represent, namely the promise of further refining our understanding of hot QCD matter.

As seen in Fig. 2.9, higher harmonics of the flow (v_3, v_4, \dots) are indeed nonzero in heavy ion collisions at mid-rapidity. For the most central Pb+Pb collisions at the LHC, v_2 and v_3 are of comparable strength and non-vanishing higher harmonics v_4, v_5, \dots are needed to account for the measured azimuthal distribution of produced particles. Consistent with the picture that v_3 depends purely on initial state fluctuations, the azimuthal orientation Ψ_3 of the measured triangularity is not correlated with the orientation of the event-averaged almond-like nuclear overlap and therefore is not correlated to the orientation Ψ_2 of the elliptic flow. Third order flow

therefore vanishes if reconstructed with respect to the standard participant plane $\Phi_2 = \Phi_{pp}$, see lower panel of Fig. 2.9. Also, the centrality dependence of higher flow harmonics is much flatter than that of v_2 , exactly because higher harmonics have a stronger dependence on fluctuations and a weaker dependence on changes in the shape of the almond-like event-averaged nuclear overlap. From model studies of the event-by-event fluctuations in the initial energy density distribution, one concludes that although the origins of v_3 and v_2 are different, v_3 is nevertheless related approximately linearly to the third order eccentricity, $v_3 \propto \epsilon_3$, as seen in the upper panel of Fig. 2.9. This provides a strong indication that collective dynamics is also very efficient in translating higher order *spatial* eccentricities into *momentum* anisotropies.

While we shall not give a detailed account of the dynamical propagation of higher order eccentricities in the following, we would like to emphasize here the generic interest in these studies. In general, any asymmetry in the initial spatial density distribution translates into pressure gradients that will propagate as perturbations. Whether such perturbations are damped out or propagate unattenuated in a heavy ion collision will depend on the dissipative properties of the QCD medium through which they propagate. It is widely known from studies of the cosmic microwave background that the analysis of fluctuations that are propagated fluid dynamically gives access to measures of the matter content of the Universe. In close analogy, one expects that in the coming years the analysis of the event-by-event fluctuations seen in ultra-relativistic heavy ion collisions will provide stringent and complementary tests of the paradigm that the QCD matter produced in ultra-relativistic nucleus–nucleus collisions is a strongly coupled almost ideal liquid and will tighten the determination of the parameter η/s that characterizes the (small) amount of dissipation that arises as it flows.

2.2.3 Calculating elliptic flow using (ideal) hydrodynamics

We have now seen that the azimuthal asymmetry *in space* present at the start of the little explosions created by heavy ion collisions with nonzero impact parameter, is subsequently converted by collective dynamics into an asymmetry *in momentum space*. This conversion of initial spatial anisotropy into final momentum anisotropy is characteristic of any explosion – one shapes the explosive in order to design a charge that blasts with greater force in some directions than others. Hydrodynamics provides the natural language for describing such processes: the initial spatial anisotropy corresponds to anisotropic pressure gradients. Let us assume first that event-by-event fluctuations in the initial state are absent and that the nuclear overlap is almond-shaped. The pressure is then maximal at the center of the nuclear overlap and zero at its edge, the gradient is greater across the

almond (in the reaction plane) than along it (perpendicular to the reaction plane). The elliptic flow v_2 measures the extent to which these pressure gradients lead to an anisotropic explosion with greater momentum flow in the reaction plane; it characterizes the efficiency of translating initial pressure gradients into collective flow. In the presence of initial state fluctuations, the same logic applies to higher order flow harmonics v_n and their relation to higher order initial spatial eccentricities ϵ_n . For concreteness, we focus in the following only on elliptic flow, which was also historically the first example of azimuthal anisotropy that was analyzed. By doing hydrodynamic calculations and comparing the calculated v_2 to that in the data (and also comparing the final state radial flow velocity to that determined from the single-particle spectra as described in the previous section) one can constrain the input quantities that go into a hydrodynamic description.

The starting point in any hydrodynamic analysis is to consider the limit of ideal, zero viscosity, hydrodynamics. In this limit, the hydrodynamic description is specified entirely by an equation of state, which relates the pressure and the energy density, and by the initial spatial distribution of energy density and fluid velocity. In particular, in ideal hydrodynamics one is setting all dissipative coefficients (shear viscosity, bulk viscosity, and their many higher order cousins) to zero. If the equation of state is held fixed and viscosity is turned on, v_2 must decrease: turning on viscosity introduces dissipation that has the effect of turning some of the initial anisotropy in pressure gradients into entropy production, rather than into directed collective flow. So, upon making some assumption for the equation of state and for the initial energy density distribution, setting the viscosities to zero yields an upper bound on the v_2 in the final state. Ideal, inviscid, hydrodynamics has therefore long been used as a calculational benchmark in heavy ion physics. As we shall see below, in heavy ion collisions at RHIC energies ideal hydrodynamics does a good job of describing $v_2(p_T)$ for pions, kaons and protons for transverse momenta p_T below about 1–2 GeV. This motivates an ongoing research program in which one begins by comparing data to the limiting case of ideal, inviscid, hydrodynamics and then turns to a characterization of dissipative effects, asking how large a viscosity will spoil the agreement with data. In this subsection, we sketch how practitioners determine the equation of state and initial energy density profile, and we recall the basic principles behind the hydrodynamic calculations on which all these studies rest. In the next subsection, we summarize the current constraints on the shear viscosity that are obtained by comparing to v_2 data.

The equation of state relates the pressure P to the energy density ϵ . P is a thermodynamic quantity, and therefore can be calculated by using the methods of lattice quantum field theory, as we describe in Chapter 3. Lattice calculations (or fits to them) of $P(\epsilon)$ in the quark–gluon plasma and in the crossover

regime between QGP and hadron gas are often used as inputs to hydrodynamic calculations. At lower energy densities, practitioners either use a hadron resonance gas model equation of state, or match the hydrodynamic calculation onto a hadron cascade model. One of the advantages of focusing on the v_2 observable is that it is insensitive to the late time epoch of the collision, when all the details of these choices matter. This insensitivity is easy to understand. v_2 describes the conversion of a spatial anisotropy into anisotropic collective flow. As this conversion begins, the initial almond-shaped collision explodes with greater momentum across the short direction of the almond, and therefore circularizes. Once it has circularized, no further v_2 can develop. Thus, v_2 is generated early in the collision. By the late times when a hadron gas description is needed, v_2 has already been generated. In contrast, the final state radial flow velocity reflects a time integral over the pressure built up during all epochs of the collision.

The discussion above reminds us of a second sense in which the ideal hydrodynamic calculation of v_2 is a benchmark: ideal hydrodynamics requires local equilibrium. It therefore cannot be valid from time $t = 0$. By using an ideal hydrodynamic description beginning at $t = 0$ we must again be overestimating v_2 , and so we can ask how long an initial phase during which partons stream freely without starting to circularize the almond-shaped region and generating any momentum anisotropy, can be tolerated without spoiling the agreement between calculations and data.

After choosing an equation of state, an initialization time, and viscosities (zero in the benchmark calculation), the only thing that remains to be specified is the distribution of energy density as a function of position in the almond-shaped collision region. (The transverse velocities are assumed to be zero initially.) In the simplest approach, called the Glauber model, this energy density is proportional to the product of the thickness of the two nuclei at a given point in the transverse plane. It is thus zero at the edge of the almond, where the thickness of one nucleus goes to zero, and maximum at the center of the almond. The proportionality constant is determined by fitting to data other than v_2 , see e.g. Ref. [542]. The assumptions behind this Glauber approach to estimating how much energy density is created at a given location as a function of the nuclear thickness at that location are assumptions about physics of the collision at $t = 0$. There are alternative model parametrizations. Here, we mention a second one for which the energy density rises towards the center of the almond more rapidly than the product of the nuclear thicknesses. This parametrization is referred to as the CGC initial condition, since it was first motivated by ideas of parton saturation (called “color glass condensate”) [456]. The Glauber and CGC models for the initial energy density distribution are often used as benchmarks in the hope that they bracket nature’s choice.

A Hydrodynamics – generalities

We turn now to the formulation of the hydrodynamic equations of motion. Hydrodynamics is an effective theory which describes the small frequency and long wavelength limit of an underlying interacting dynamical theory [355]. It can be used to describe motions of the fluid that occur on macroscopic length scales and time scales associated with how the fluid is “stirred”, scales that are long compared to any microscopic scales characterizing the fluid itself. It is a classical field theory, where the fields can be understood as the expectation values of certain quantum operators in the underlying theory. In the hydrodynamic limit, since the length scales under consideration are longer than any correlation length in the underlying theory, by virtue of the central limit theorem all n -point correlators of the underlying theory can be factorized into one point functions. The fluctuations on these average values are small, and a description in terms of expectation values is meaningful. If the underlying theory admits a (quasi)particle description, this statement is equivalent to saying that the hydrodynamic description involves averages over many of these fundamental degrees of freedom and is valid only on length scales that are long enough for this to be an appropriate procedure.

The hydrodynamic degrees of freedom include the expectation values of conserved currents such as the stress tensor $T^{\mu\nu}$ or the currents of conserved charges J_B , which fulfill the conservation equations

$$d_\mu T^{\mu\nu} = 0, \quad (2.9)$$

$$d_\mu J_B^\mu = 0, \quad (2.10)$$

where d_μ is the covariant derivative. As a consequence of these conservation laws, long wavelength excitations of these fields can only relax on long timescales, since their relaxation must involve moving stress–energy or charges over distances of order the wavelength of the excitation. As a consequence, these conservation laws lead to excitations whose lifetime diverges with their wavelength. Such excitations are called hydrodynamic modes.

It is worth pausing to explain why we have introduced a covariant derivative, even though we will only ever be interested in heavy ion collisions – and thus hydrodynamics – occurring in flat spacetime. It is nevertheless often convenient to use curvilinear coordinates with a non-trivial metric. For example, the longitudinal dynamics is more conveniently described using proper time $\tau = \sqrt{t^2 - z^2}$ and spacetime rapidity $\xi = \text{arctanh}(z/t)$ as coordinates rather than t and z . In these “Milne coordinates”, the metric is given by $g_{\mu\nu} = \text{diag}(g_{\tau\tau}, g_{xx}, g_{yy}, g_{\xi\xi}) = (-1, 1, 1, \tau^2)$. These coordinates are useful because boost invariance simply translates into the requirement that ε , as well as the fluid velocity u^μ and $\Pi^{\mu\nu}$, the contribution to the stress tensor from gradients, must all be independent of ξ ,

depending on τ only. In particular, if the initial conditions are boost invariant then the fluid dynamic evolution will preserve this boost invariance, and the numerical calculation reduces in Milne coordinates to a $2 + 1$ -dimensional problem. Boost-invariant initial conditions have often been used as a simplifying assumption for hydrodynamics ever since they were introduced in this context in Ref. [165], but fully $3 + 1$ -dimensional calculations that do not assume boost invariance can also be found in the literature [455, 424, 654, 710].

If the only long-lived modes are those from conserved currents, then hydrodynamics describes a normal fluid. However, there can be other degrees of freedom that lead to long-lived modes in the long wavelength limit. For example, in a phase of matter in which some global symmetry is spontaneously broken, the Goldstone boson(s) is (are) also hydrodynamic modes [355]. The classic example of this is a superfluid, in which a global $U(1)$ symmetry is spontaneously broken. Chiral symmetry is spontaneously broken in QCD, but there are two reasons why we can neglect the potential hydrodynamic modes associated with the chiral order parameter [746]. First, explicit chiral symmetry breaking gives these modes a mass (the pion mass) and we are interested in the hydrodynamic description of physics on length scales longer than the inverse pion mass. Second, we are interested in temperatures above the QCD crossover, at which the chiral order parameter is disordered, the symmetry is restored, and this question does not arise. So, we need only consider normal fluid hydrodynamics. Furthermore, as we have discussed in Section 2.1, the matter produced in ultra-relativistic heavy ion collisions has only a very small baryon number density, and it is a good approximation to neglect J_B^μ . The only hydrodynamic degrees of freedom are therefore those described by $T^{\mu\nu}$.

At the length scales at which the hydrodynamic approximation is valid, each point of space can be regarded as a macroscopic *fluid cell*, characterized by its energy density ε , pressure P , and a velocity u^μ . The velocity field can be defined by the energy flow together with the constraint $u^2 = -1$. In the so-called Landau frame, the four equations

$$u_\mu T^{\mu\nu} = -\varepsilon u^\nu. \quad (2.11)$$

determine ε and u from the stress tensor.

Hydrodynamics can be viewed as a gradient expansion of the stress tensor (and any other hydrodynamic fields). In general, the stress tensor can be separated into a term with no gradients (*ideal*) and a term which contains all the gradients:

$$T^{\mu\nu} = T_{\text{ideal}}^{\mu\nu} + \Pi^{\mu\nu}. \quad (2.12)$$

In the rest frame of each fluid cell ($u^i = 0$), the ideal piece is diagonal and isotropic $T_{\text{ideal}}^{\mu\nu} = \text{diag}(\varepsilon, P, P, P)$. Thus, in any frame,

$$T_{\text{ideal}}^{\mu\nu} = (\varepsilon + P)u^\mu u^\nu + P g^{\mu\nu}, \quad (2.13)$$

where $g_{\mu\nu}$ is the spacetime metric.

If there were a nonzero density of some conserved charge n , the velocity field could either be defined in the Landau frame as above or may instead be defined in the so-called Eckart frame, with $J^\mu = nu^\mu$. In the Landau frame, the definition (2.11) of u^μ implies $\Pi^{\mu\nu}u_\nu = 0$ (transversality). Hence, there is no heat flow but there can be currents of the conserved charge. In the Eckart frame, the velocity field is comoving with the conserved charge, but there can be heat flow.

Ideal hydrodynamics is the limit in which all gradient terms in $T^{\mu\nu}$ are neglected. Corrections to ideal hydrodynamics – namely the gradient terms in $\Pi^{\mu\nu}$ that we shall discuss shortly – introduce internal length and time scales, including time scales for relaxation of perturbations away from local thermal equilibrium, and length scales associated with mean free paths. Hydrodynamics works on longer length scales than these. Introducing the gradient terms that correct ideal hydrodynamics also introduces dissipation and introduces the possibility of hydrodynamic flows in which the pressure is not isotropic. At long enough time scales, however, gradients become unimportant, hydrodynamics becomes ideal, the pressure P in the rest frame of each fluid cell becomes isotropic, and ε and P are related by the equilibrium equation of state. This equation of state can be determined by studying a homogeneous system at rest with no gradients, for example via a lattice calculation.

The range of applicability of hydrodynamics can be characterized in terms of the isotropization scale τ_{iso} and the hydrodynamization scale τ_{hydro} . The isotropization scale measures the characteristic time over which an initially anisotropic stress tensor becomes isotropic in the local fluid rest frame, to within some criterion that must be defined. The hydrodynamization scale measures the characteristic time after which the flow of the fluid is well described (again to within some criterion that must be defined) by the equations of (possibly viscous) hydrodynamics. In different contexts, the two time scales τ_{iso} and τ_{hydro} can be ordered in either way. If $\tau_{\text{iso}} < \tau_{\text{hydro}}$, as may be the case for a sufficiently weakly coupled plasma [79], there is a period of time when the plasma is isotropic but is not yet described by ideal hydrodynamics with P and ε related by the equilibrium equation of state. In some circumstances [79], ideal hydrodynamics may nevertheless be used during this period of time, as long as $P(\varepsilon)$ is replaced by some non-equilibrium “equation of state” that will depend on exactly how the system is out of thermal equilibrium. It is also possible that during this period of time the production of entropy may not yet have ceased. If $\tau_{\text{hydro}} < \tau_{\text{iso}}$, on the other hand, there is a period of time when the way in which the plasma flows is described well by viscous hydrodynamics even though gradients in the flow remain important, entropy is still being produced, and the pressure in the local fluid rest frame is not isotropic. We shall return to these considerations at length in Chapter 7 where we shall see that calculations done

via gauge/string duality indicate that when a strongly coupled plasma is produced it hydrodynamizes first and isotropizes later. This conclusion has been reached for hydrodynamization starting from a wide variety of far-from-equilibrium states. Although this could in principle have been considered as a possibility beforehand, in fact it was the analysis using gauge/string duality of many explicit examples in which this conclusion was manifest that brought it to the fore and that has yielded the insight that hydrodynamization before isotropization may be a generic feature of the production of strongly coupled plasma.

B First order dissipative fluid dynamics

Going beyond the infinite wavelength limit requires the introduction of viscosities. To first order in gradients, the requirement that $\Pi^{\mu\nu}$ be transverse means that it must take the form

$$\Pi^{\mu\nu} = -\eta(\varepsilon)\sigma^{\mu\nu} - \zeta(\varepsilon)\Delta^{\mu\nu}\nabla\cdot u, \quad (2.14)$$

where η and ζ are the shear and bulk viscosities, $\nabla^\mu = \Delta^{\mu\nu}d_\nu$, with d_ν the covariant derivative and

$$\Delta^{\mu\nu} = g^{\mu\nu} + u^\mu u^\nu, \quad (2.15)$$

$$\sigma^{\mu\nu} = \Delta^{\mu\alpha}\Delta^{\nu\beta}(\nabla_\alpha u_\beta + \nabla_\beta u_\alpha) - \frac{2}{3}\Delta_{\alpha\beta}\nabla\cdot u. \quad (2.16)$$

The operator $\Delta^{\mu\nu}$ is the projector onto the space components of the fluid rest frame. Note that in this frame the only time derivatives or spatial gradients that appear in Eq. (2.14) are spatial gradients of the velocity fields. By symmetry, time derivatives of the velocity fields and spatial gradients of ε cannot arise in $\Pi^{\mu\nu}$ to first order in gradients. The reason that time derivatives of ε do not appear is that they can be eliminated in the first order equations by using the zeroth order equation of motion

$$D\varepsilon = -(\varepsilon + P)\nabla_\mu u^\mu, \quad (2.17)$$

where $D = u^\mu d_\mu$ is the time derivative in the fluid rest frame. (Similarly, time derivatives of the energy density can be eliminated in the second order equations that we shall give below using the first order equations of motion.)

It is often convenient to phrase the hydrodynamic equations in terms of the entropy density s . In the absence of conserved charges, i.e. with baryon chemical potential $\mu_B = 0$, the entropy density is $s = (\varepsilon + P)/T$. Using this and another fundamental thermodynamic relation, $DE = TDS - PDV$ (where $E/V \equiv \varepsilon$) the zeroth order equation of motion (2.17) becomes exactly the equation of entropy flow for an ideal isentropic fluid

$$Ds = -s\nabla_\mu u^\mu. \quad (2.18)$$

Repeating this analysis at first order, including the viscous terms, one easily derives from $u_\mu \nabla_\nu T^{\mu\nu} = 0$ that

$$\frac{Ds}{s} = -\nabla_\mu u^\mu - \frac{1}{sT} \Pi^{\mu\nu} \nabla_\nu u_\mu. \quad (2.19)$$

A similar analysis of the other three hydrodynamic equations then shows that they take the form

$$Du_\alpha = -\frac{1}{Ts} \Delta_{\alpha\nu} \left(\nabla^\nu P + \nabla_\mu \Pi^{\mu\nu} \right). \quad (2.20)$$

It then follows from the structure of the shear tensor $\Pi^{\mu\nu}$ (2.14) that shear viscosity and bulk viscosity always appear in the hydrodynamic equations of motion in the dimensionless combinations η/s and ζ/s . The net entropy increase is proportional to these dimensionless quantities. Gradients of the velocity field are measured in units of $1/T$.

In a conformal theory, $\zeta = 0$ since $\Pi^{\mu\nu}$ must be traceless. There are a number of indications from lattice calculations that as the temperature is increased above $(1.5-2)T_c$, with T_c the crossover temperature, the quark–gluon plasma becomes more and more conformal. The equation of state approaches $P = \frac{1}{3}\varepsilon$ [277, 179]. The bulk viscosity drops rapidly [615]. So, we shall set $\zeta = 0$ throughout the following, in so doing neglecting temperatures close to T_c . One of the things that makes heavy ion collisions at the LHC interesting is that in these collisions the plasma that is created is expected to be better approximated as conformal than is the case at RHIC, where the temperature at $\tau = 1$ fm is thought to be between $1.5T_c$ and $2T_c$.

Just like the equation of state $P(\varepsilon)$, the shear viscosity $\eta(\varepsilon)$ is an input to the hydrodynamic description that must be obtained either from experiment or from the underlying microscopic theory. We shall discuss in Section 3.2 how transport coefficients like η are obtained from correlation functions of the underlying microscopic theory via Kubo formulae.

C Second order dissipative hydrodynamics

Even though hydrodynamics is a controlled expansion in gradients, the first order expression for the tensor $\Pi^{\mu\nu}$, Eq. (2.14), is unsuitable for numerical computations. The problem is that the set of equations (2.9) with the approximations (2.14) leads to acausal propagation. Even though this problem only arises for modes outside of the region of validity of hydrodynamics (namely high momentum modes with short wavelengths of the order of the microscopic length scale defined by η), the numerical evaluation of the first order equations of motion is sensitive to the acausality in these hard modes. This problem is solved by going to one higher order in the gradient expansion. This is known as second order hydrodynamics.

There is a phenomenological approach to second order hydrodynamics due to Müller, Israel and Stewart aimed at explicitly removing the acausal propagation [631, 485, 486]. In this approach, the tensor $\Pi^{\mu\nu}$ is treated as a new hydrodynamic variable and a new dynamical equation is introduced. In its simplest form this equation is

$$\tau_{\Pi} D\Pi^{\mu\nu} = -\Pi^{\mu\nu} - \eta\sigma^{\mu\nu}, \quad (2.21)$$

where τ_{Π} is a new (second order) coefficient. Note that as $\tau_{\Pi} \rightarrow 0$, Eq. (2.21) coincides with Eq. (2.14) with the bulk viscosity ζ set to zero. Eq. (2.21) is such that $\Pi^{\mu\nu}$ relaxes to its first order form in a (proper) time τ_{Π} . There are several variants of this equation in the literature, all of which follow the same philosophy. They all introduce the relaxation time as the characteristic time in which the tensor $\Pi^{\mu\nu}$ relaxes to its first order value. The variations arise from different ways of fixing some pathologies of Eq. (2.21), since as written Eq. (2.21) does not lead to a transverse stress tensor (although this is a higher order effect) and is not conformally invariant. Since in this approach the relaxation time is introduced *ad hoc*, it may not be possible to give a prescription for extracting it from the underlying microscopic theory.

The systematic extraction of second order coefficients demands a similar analysis of the second order gradients as was done at first order. The strategy is, once again, to write all possible terms with two derivatives which are transverse and consistent with the symmetries of the theory. As before, only spatial gradients (in the fluid rest frame) are considered, since time gradients can be related to the former via the zeroth order equations of motion.

In a conformal theory, second order hydrodynamics simplifies. First, only terms such that $\Pi_{\mu}^{\mu} = 0$ are allowed. Furthermore, the theory must be invariant under Weyl transformations

$$g_{\mu\nu} \rightarrow e^{-2\omega(x)} g_{\mu\nu} \quad (2.22)$$

which implies

$$T \rightarrow e^{\omega(x)} T, \quad u^{\mu} \rightarrow e^{\omega(x)} u^{\mu}, \quad T^{\mu\nu} \rightarrow e^{(d+2)\omega(x)} T^{\mu\nu}, \quad (2.23)$$

where T is the temperature and d the number of spacetime dimensions. The Weyl transformation of the stress tensor can be derived from its definition in terms of the action S , which is Weyl invariant: $T^{\mu\nu} = (2/\sqrt{g}) \delta S/\delta g_{\mu\nu}$. The normalization of the velocity field, $u_{\mu} u^{\mu} = -1$, fixes its Weyl transformation. Finally, the transformation of the stress tensor together with the relation (2.13) and the fact that in a conformal theory $\varepsilon \sim T^d$ yield the Weyl transformation of T .

It turns out that there are only five operators that respect these constraints [107]. The second order contributions to the tensor $\Pi^{\mu\nu}$ are linear combinations of these operators, and can be cast in the form [107, 155]

$$\begin{aligned}
\Pi^{\mu\nu} = & -\eta\sigma^{\mu\nu} - \tau_{\Pi} \left[\langle D\Pi^{\mu\nu} \rangle + \frac{d}{d-1}\Pi^{\mu\nu}(\nabla\cdot u) \right] \\
& + \kappa \left[R^{\langle\mu\nu\rangle} - (d-2)u_{\alpha}R^{\alpha\langle\mu\nu\rangle\beta}u_{\beta} \right] \\
& + \frac{\lambda_1}{\eta^2}\Pi^{\langle\mu}_{\lambda}\Pi^{\nu\rangle\lambda} - \frac{\lambda_2}{\eta}\Pi^{\langle\mu}_{\lambda}\Omega^{\nu\rangle\lambda} + \lambda_3\Omega^{\langle\mu}_{\lambda}\Omega^{\nu\rangle\lambda}. \quad (2.24)
\end{aligned}$$

Here, $R^{\mu\nu}$ is the Ricci tensor, the indices in brackets are the symmetrized traceless projectors onto the space components in the fluid rest frame, namely

$$\langle A^{\mu\nu} \rangle \equiv \frac{1}{2}\Delta^{\mu\alpha}\Delta^{\nu\beta}(A_{\alpha\beta} + A_{\beta\alpha}) - \frac{1}{d-1}\Delta^{\mu\nu}\Delta^{\alpha\beta}A_{\alpha\beta} \equiv A^{\langle\mu\nu\rangle}, \quad (2.25)$$

and the vorticity tensor is defined as

$$\Omega^{\mu\nu} \equiv \frac{1}{2}\Delta^{\mu\alpha}\Delta^{\nu\beta}(\nabla_{\alpha}u_{\beta} - \nabla_{\beta}u_{\alpha}). \quad (2.26)$$

In deriving (2.24), we have replaced $\eta\sigma^{\mu\nu}$ by $\Pi^{\mu\nu}$ on the right-hand side in places where doing so makes no change at second order. We see from (2.24) that five new coefficients τ_{Π} , κ , λ_1 , λ_2 , and λ_3 arise at second order in the hydrodynamic description of a conformal fluid, in addition to η and the equation of state which arise at first and zeroth order respectively. The coefficient κ is not relevant for hydrodynamics in flat spacetime. The λ_i coefficients involve nonlinear combinations of fields in the rest frame and, thus, are invisible in linearized hydrodynamics. Thus, these three coefficients cannot be extracted from linear response. Of these three, only λ_1 is relevant in the absence of vorticity, as in the numerical simulations that we will describe in the next subsection. These simulations have also shown that, for physically motivated choices of λ_1 , the results are insensitive to its precise value, leaving τ_{π} as the only phenomenologically relevant second order parameter in the hydrodynamic description of a conformal fluid. In a generic, nonconformal fluid, there are nine additional transport coefficients [715].

For more in-depth discussions of second order viscous hydrodynamics and its applications to heavy ion collisions, see e.g. Refs. [632, 776, 633, 634, 106, 105, 716, 754, 333, 752, 589, 753, 625, 755, 590, 756, 757, 442, 714, 779, 735, 758].

2.2.4 Comparing elliptic flow in heavy ion collisions and hydrodynamic calculations

For the case of ideal hydrodynamics, the hydrodynamic equations of motion are fully specified once the equation of state $P = P(\varepsilon)$ is given. A second order dissipative hydrodynamic calculation also requires knowledge of the transport coefficients $\eta(\varepsilon)$ and $\zeta(\varepsilon)$ (although in practice the latter is typically set to zero) and the relaxation time τ_{Π} and the second order coefficient λ_1 entering Eq. (2.24). (Note that κ would enter in curved spacetime, and λ_2 and λ_3 would enter in the

presence of vorticity.) All these parameters are well-defined in terms of correlation functions in the underlying quantum field theory. In this sense, the hydrodynamic evolution equations are model-independent.

The output of any hydrodynamic calculation depends on more than the evolution equations. One must make model assumptions about the initial energy density distribution. As we have discussed in Section 2.2.2, there are two benchmark models for the energy distribution across the event-averaged almond-shaped collision region. More recent analyses also include initial event-by-event fluctuations around these event-averaged density profiles [458]. These model variations give us a sense of the degree to which results are sensitive to our lack of knowledge of the details of this initial profile. Often, the initial transverse velocity fields are set to zero and boost invariance is assumed for the longitudinal velocity field and the evolution. For dissipative hydrodynamic simulations, the off-diagonal elements of the energy–momentum tensor are additional hydrodynamic fields which must be initialized. The initialization time τ_0 , at which these initial conditions are fixed, is an additional model parameter. It can be viewed as characterizing the isotropization time, at which hydrodynamics starts to apply but collective flow has not yet developed. In addition to initial state sensitivity, results depend on assumptions made about how the system stops behaving hydrodynamically and *freezes out*. In practice, freezeout is often assumed to happen as a rapid decoupling: when a specified criterion is satisfied (e.g. when a fluid cell drops below a critical energy or entropy density) then the hydrodynamic fields in the unit cell are mapped onto hadronic equilibrium Bose/Fermi distributions. This treatment assumes that hydrodynamics is valid all the way down to the kinetic freezeout temperature, below which one has noninteracting hadrons. Alternatively, at a higher temperature close to the crossover where hadrons are formed, one can map the hydrodynamic fields onto a hadron cascade which accounts for the effects of rescattering in the interacting hadronic phase without assuming that its behavior is hydrodynamic [121, 775]. Indeed, recent work suggests that hadronization may be triggered by cavitation induced by the large bulk viscosity in the vicinity of the crossover temperature [702]. As we have discussed, v_2 is insensitive to details of how the late-time evolution is treated because v_2 is generated during the epoch when the collision region is azimuthally anisotropic. Nevertheless, these late-time issues do matter when one does a global fit to v_2 and the single-particle spectra, since the latter are affected by the radial flow which is built up over the entire history of the collision. Finally, the validity of results from any hydrodynamic calculation depends on the assumption that a hydrodynamic description is applicable. This assumption can be checked at late times by checking the sensitivity to how freezeout is modeled and can be checked at early times by confirming the insensitivity of results to the values of the second order hydrodynamic coefficients and to the initialization of the higher order




Figure 2.10 The elliptic flow v_2 versus p_T for a large number of identified hadrons (pions, kaons, protons, Λ s) showing the comparison between an *ideal* hydrodynamic calculation to data from RHIC. Figure taken from Ref. [440].

off-diagonal elements of the energy–momentum tensor: if hydrodynamics is valid, the gradients must be small enough at all times that second order effects are small compared to first order effects.

In practice, the dependence of physics conclusions on all these model assumptions has to be established by systematically varying the initial conditions and freezeout prescriptions within a wide physically motivated parameter range, and comparing to data on both the single-particle spectra (i.e. the radial velocity) and the azimuthal flow anisotropy coefficients. At the current time, several generic observations have emerged from pursuing this program in comparison to data from RHIC and LHC.

- (1) *Perfect fluid dynamics approximately reproduces the size and centrality dependence of v_2*

RHIC and LHC data on single inclusive hadronic spectra $d^3N/p_T dp_T dy$ and their leading azimuthal dependence $v_2(p_T)$ can be reproduced approximately in magnitude and shape by ideal hydrodynamic calculations, for particles with $p_T < (1-2)$ GeV, see Fig. 2.10. The hydrodynamic picture is expected to break down for sufficiently small wavelength, i.e. high momenta, consistent with the observation that significant deviations occur for $p_T > 2$ GeV, see Fig. 2.10 again. The initialization time for these calculations is $\tau_0 = 0.6-1$ fm.




Figure 2.11 Centrality dependence of the p_T -averaged elliptic (v_2) and triangular (v_3) flow. The plot compares fluid dynamic simulations to the ratio of data on v_n over eccentricity ϵ_n , the latter being calculated in two different models of fluctuating initial conditions (MC-KLN and MC-Glb) that formulate opposite extreme assumptions about the radial dependence of the initial transverse energy density. The figures show that data on v_2 alone cannot differentiate between these two different assumptions about the initial energy density profile, since each can fit the v_2 data comparably well upon making very different choices for η/s . The analysis indicates that using data on v_3 in addition can result in separate constraints on both η/s and initial conditions. Figure taken from Ref. [701].

If τ_0 is chosen larger, the agreement between ideal hydrodynamics and data is spoiled. This gives significant support to a picture in which thermalization is achieved within 1 fm after the collision. Historically, the agreement between ideal hydrodynamic calculations and experimental measurements of v_2 provided the first indication that the shear viscosity of the fluid produced at RHIC must be small.

(2) The mass ordering of identified hadron spectra

The p_T -differential azimuthal asymmetry $v_2(p_T)$ of identified single inclusive hadron spectra shows a characteristic mass ordering in the range of $p_T < 2$ GeV: at small p_T , the azimuthal asymmetry of light hadrons is significantly more pronounced than that of heavier hadrons, see e.g. Fig. 2.10. This qualitative agreement between hydrodynamic simulations and experimental data supports the picture that all hadron species emerge from a single fluid moving with a common flow-velocity field.

(3) Data on v_2 and v_3 support small dissipative coefficients such as shear viscosity

Above the crossover temperature, the largest dissipative correction is expected to arise from shear viscosity η , which enters the equations of motion of second order dissipative hydrodynamics in the combination η/s , where s is the entropy density. Figure 2.11 shows a comparison of data on v_2 and v_3 with model simulations that include the major physics effects considered so far, including initial event-by-event fluctuations with different initial density

profiles, viscous hydrodynamic evolution, and the interface between hydrodynamic evolution and hadronic freezeout at late times. Viscous corrections generically decrease the observed v_2 or v_3 for a given initial eccentricity ϵ_2 or ϵ_3 , because dissipation results in the production of more heat and less collective flow. So, for example, v_2 decreases with increasing η/s . Hydrodynamic analyses of heavy ion collision data have put increasingly tight constraints on η/s [589, 475, 756, 757] and most recent analyses favor nonzero but small values in the range $1/(4\pi) < \eta/s < 2/(4\pi)$ [701]. Analyses at the current frontier seek to use data on several v_n s to constrain how different harmonics are sourced differently by initial state fluctuations and damped differently by the effects of η/s . It is anticipated that these analyses will further tighten constraints on η/s in coming years, while at the same time yielding experimental insights into the initial fluctuations. Also, comparison of the analyses of heavy ion collisions at RHIC and the LHC may begin to teach us about the temperature dependence of η/s . The smallness of η/s is remarkable, since almost all other known liquids have $\eta/s > 1$ and most have $\eta/s \gg 1$. The one liquid that is comparably close to ideal is an ultracold gas of strongly coupled fermionic atoms, whose η/s is also well below 1 and may be comparably small to that of the quark–gluon plasma produced at RHIC [235]. Both these fluids are much better described by ideal hydrodynamics than water is. Both have η/s comparable to the value $1/4\pi$, as we shall comment on below.

We close this section by noting that while hydrodynamic calculations reproduce elliptic flow, a treatment in which the Boltzmann equation for quark and gluon (quasi)particles is solved, including all $2 \rightarrow 2$ scattering processes with the cross-sections as calculated in perturbative QCD, fails dramatically. It results in values of v_2 that are much smaller than in the data. Agreement with data can only be achieved if the parton scattering cross-sections are increased *ad hoc* by more than a factor of 10 [624]. With such large cross-sections, a Boltzmann description cannot be reliable since the mean free path of the particles becomes comparable to or smaller than the interparticle spacing. Another way of reaching the same conclusion is to note that if a perturbative description of the QGP as a gas of interacting quasiparticles is valid, the effective QCD coupling α_s describing the interaction among these quasiparticles must be small, and for small α_s perturbative calculations of η/s are controlled and yield parametrically large values [75]

$$\left. \frac{\eta}{s} \right|_{\text{perturbatively}} \propto \frac{1}{\alpha_s^2 \ln [1/\alpha_s]}. \quad (2.27)$$

It is not possible to get as small a value of η/s as the data requires from the perturbative calculation without increasing α_s to the point that the calculation is invalid.

In contrast, as we shall see in Section 6.2, any gauge theory with a gravity dual must have $\eta/s = 1/(4\pi)$ in the large- N_c and strong coupling limit and, furthermore, the plasma fluids described by these theories in this limit do not have any well-defined quasiparticles. This calculational framework thus seems to do a much better job of capturing the qualitative features needed for a successful phenomenology of collective flow in heavy ion collisions.

2.3 Jet quenching

Having learned that heavy ion collisions produce a low viscosity, strongly coupled, fluid we now turn to experimental observables with which we may study properties of the fluid beyond just how it flows. There are many such observables available. In this section and the next we shall describe two classes of observables, selected because in both cases there is (the promise of) a substantive interplay between data from RHIC and the LHC and qualitative insights gained from the analysis of strongly coupled plasmas with dual gravity descriptions.

Jet quenching refers to a suite of experimental observables that together reveal what happens when a very energetic quark or gluon (with momentum much greater than the temperature) plows through the strongly coupled plasma. Some measurements focus on how rapidly the energetic parton loses its energy; other measurements give access to how the strongly coupled fluid responds to the energetic parton passing through it. These energetic partons are not external probes; they are produced within the same collision that produces the strongly coupled plasma itself.

In a small fraction of *proton–proton* collisions, partons from the incident protons scatter with a large momentum transfer, producing back-to-back partons in the final state with transverse momenta of the order of ten or a few tens of GeV. These “hard” processes are rare, but data samples are large enough that they are nevertheless well studied. The high transverse momentum partons in the final state manifest themselves in the detector as jets. Individual high p_T hadrons in the final state come from such hard processes and are typically found within jets. In addition to copious data from proton-(anti)proton collisions, there is a highly developed quantitatively controlled calculational framework built upon perturbative QCD that is used to calculate the rates for hard processes in high energy hadron–hadron collisions. These calculations are built upon factorization theorems. Consider as an example the single inclusive charged hadron spectrum at high p_T , see Fig. 2.3. That is, consider the production cross-section for a single charged hadron with a given high transverse momentum p_T , regardless of what else is produced in the hadron–hadron collision. This quantity is calculated as a convolution of separate (factorized) functions that describe different aspects of the process: (i) the process-independent parton

distribution function gives the probability of finding partons with a given momentum fraction in the incident hadrons; (ii) the process-dependent hard scattering cross-section gives the probability that those partons scatter into final state partons with specified momenta; and (iii) the process-independent parton fragmentation functions that describe the probability that a final state parton fragments into a jet that includes a charged hadron with transverse momentum p_T . Functions (i) and (iii) are well measured and at high transverse momentum function (ii) is both systematically calculated and well measured. This body of knowledge provides a firm foundation, a well-defined baseline with respect to which we can measure changes if such a hard scattering process occurs instead in an ultra-relativistic heavy ion collision.

In hard scattering processes in which the momentum transfer Q is high enough, the partonic hard scattering cross-section (function (ii) above) is expected to be the same in an ultrarelativistic heavy ion collision as in a proton–proton collision. This is so because the hard interaction occurs on a timescale and length scale $\propto 1/Q$ which is too short to resolve any aspects of the hot and dense strongly interacting medium that is created in the same collision. The parton distribution functions (function (i) above) are different in nuclei than in nucleons, but they may be measured in proton–nucleus, deuteron–nucleus, and electron–nucleus collisions. The key phenomenon that is unique to ultra-relativistic nucleus–nucleus collisions is that after a very energetic parton is produced, unless it is produced at the edge of the fireball heading outwards it must propagate through as much as 5–10 fm of the hot and dense medium produced in the collision. These hard partons therefore serve as well-calibrated probes of the strongly coupled plasma whose properties we are interested in. The presence of the medium results in the hard parton losing energy and changing the direction of its momentum. The change in the direction of its momentum is often referred to as “transverse momentum broadening”, a phrase which needs explanation. “Transverse” here means perpendicular to the original direction of the hard parton. (This is different from p_T , the component of the (original) momentum of the parton that is perpendicular to the beam direction.) “Broadening” refers to the effect on a jet when the directions of the momenta of many hard partons within it are kicked; averaged over many partons in one jet, or perhaps in an ensemble of jets, there is no change in the mean momentum but the spread of the momenta of the individual partons broadens.

Because the rates for hard scattering processes drop rapidly with increasing p_T , energy loss translates into a reduction in the number of partons produced with a given p_T . (Partons with the given p_T must have been produced with a higher p_T , and are therefore rarer than they would be in proton–proton collisions; as a consequence, the yield of high p_T hadrons is rarer since it results from hadronization of highly energetic partons.) Transverse momentum broadening, on the other hand,




Figure 2.12 CMS data showing a highly unbalanced dijet event in a Pb+Pb collision at $\sqrt{s_{NN}} = 2.76$ TeV. Tower heights denote the sum of transverse energy deposited in the electromagnetic and hadron calorimeters in a particular segment of azimuthal angle ϕ and pseudo rapidity η . The reconstructed jets, in red, are labeled with their corrected jet transverse momentum. Figure taken from Ref. [264].

carries part of the jet energy away from the jet axis and thus also leads to a reduction in the rate of jets observed at a given jet energy. Furthermore, the hard parton dumps energy into the medium, which motivates the use of observables involving correlations between soft final state hadrons and a high momentum hadron. Most generally, “jet quenching” refers to the whole suite of medium-induced modifications of high p_T processes in heavy ion collisions and modifications of the medium in heavy ion collisions in which a high p_T process occurs, all of which have their origin in the propagation of a highly energetic parton through the strongly coupled plasma.

As we discuss in the following, one of the most detailed experimental sets of information about jet quenching is provided by the medium-induced suppression of single inclusive hadron spectra first discovered at RHIC. A more recent, and arguably more pictorial, manifestation of jet quenching in heavy ion collisions is provided by the CMS event display shown in Fig. 2.12. This Pb+Pb event was selected by triggering on a “leading jet” (i.e. a highly collimated spray of energetic particles that may be thought of as arising from the fragmentation of a single highly energetic parent parton). By momentum conservation, the total transverse momentum of this leading jet must be balanced by recoil in the opposite azimuthal hemisphere. However, the subleading jet seen in Fig. 2.12 in the opposite azimuthal

hemisphere at $\Delta\phi \approx \pi$ balances only approximately one third of the momentum of the leading jet. Since no other jet structure is visible, a total recoil transverse momentum of $205 - 70 = 135$ GeV must have been lost in this particular event by the recoiling parent parton and must reside in many soft fragments that are, in the present plot, indistinguishable from the background in a high multiplicity heavy ion collision. Therefore, qualitatively, Fig. 2.12 illustrates the case of a Pb+Pb collision at the LHC in which one parton escaped relatively unscathed while its back-to-back partner was very significantly degraded by the presence of the medium.

We shall limit our presentation to several generic features of jet quenching that have been established at the LHC and at RHIC.

(1) *Characteristic strong centrality dependence of dijet asymmetry A_J*

The imbalance between the transverse energy of the leading jet E_{T1} and that of the recoil jet E_{T2} can be characterized by measuring the normalized difference $A_J = \frac{E_{T1} - E_{T2}}{E_{T1} + E_{T2}}$. We caution the reader that Fig. 2.12 does *not* imply that the recoiling jet has lost $205 - 70 = 135$ GeV by interactions with the surrounding medium. Even in the absence of medium effects, one finds in “elementary” p+p or p+ \bar{p} collisions that dijet events are broadly distributed in A_J . This can be understood in perturbative QCD as a consequence of perturbative parton branching processes, due to which the recoil is taken by more than one jet or due to which energy is put outside the recoiling jet cone. The main finding of first measurements at the LHC is, however, that the effects of this perturbative fragmentation are by far not sufficient to understand the distribution of dijet asymmetries measured in central Pb+Pb collisions at the LHC. More precisely, by varying the centrality of a heavy ion collision, one changes the typical in-medium path length over which hard partons produced in these collisions must propagate through the dense matter. For the most central head-on collisions, corresponding to the longest in-medium path lengths for the hard partons, the dijet asymmetry distribution in Pb+Pb collisions is significantly broader than in the baseline p+p collisions. In contrast, the dijet asymmetry distribution has a comparable width in p+p and peripheral Pb+Pb collisions [2, 264]. This establishes that there is jet quenching: a significant fraction of the recoiling jet energy must be transported outside the jet cone by effects due to the presence of the medium produced in heavy ion collisions.

(2) *Absence of azimuthal decorrelation of dijets*

In p+p collisions, momentum conservation dictates that in dijet events the jets recoil against each other with an angular distribution that peaks at azimuthal angle $\Delta\phi = \pi$. Deviations from this back-to-back correlation can arise in p+p collisions, for instance from the presence of three-jet events. Since the medium transfers momentum to the dijet system, it is in principle conceivable

that this dijet angular correlation broadens, but such an effect is not observed in Pb+Pb collisions. Indeed, the $\Delta\phi$ -distribution in Pb+Pb collisions is independent of centrality and comparable in width and shape to the one seen in p+p collisions [2, 264]. This provides important constraints on the dynamics of jet quenching. For instance, the recoiling jet cannot lose its energy by radiating a single high energy particle outside the jet cone, since the recoil from such radiation would necessarily broaden the $\Delta\phi$ -distribution. The additional lost jet energy must be distributed amongst many soft fragments.

(3) *Soft jet fragments transported to large angles outside jet cone*

By analyzing the soft background in wide phase-space regions around the sub-leading jet in dijet events, the CMS collaboration has found the apparently missing energy that is lost from the recoiling jet [264]. As expected from the absence of azimuthal decorrelation of dijets, this energy is indeed distributed over many low energy particles. Furthermore, it is broadly distributed in η and ϕ over the azimuthal hemisphere ($\pi/2 < \Delta\phi < 3\pi/2$) opposite to the leading jet. This means, in particular, that the recoiling jet cannot lose its energy by radiating particles that stay almost collinear with it.

These generic findings show that jet quenching occurs via a mechanism that degrades the energy of the hardest jet fragments significantly and that transports this energy into soft fragments moving at large angles relative to the direction in which the initial parton was propagating through the medium. The quenching of calorimetrically reconstructed jets is a rapidly progressing subject of ongoing research in which many further characteristics of medium-modified jet fragmentation are just becoming available. Further qualitative advances are also expected in the coming years from the study of jets recoiling against isolated high energy photons or Z -bosons. In such events, the initial energy (and direction) of the recoiling jet must be the same as (opposite to) that of the photon or Z -boson – which cannot be affected by the presence of the strongly coupled plasma since photons and Z -bosons interact only via the electromagnetic and weak interactions. So, once statistically significant samples of such events become available they will yield samples of jets whose initial energy and direction are known with much higher precision than at present. A more detailed account of these developments lies outside the scope of this book. In the following, we limit our discussion of jet quenching mainly to the study of leading hadron spectra and their modeling.

2.3.1 *Single inclusive high p_T spectra and “jet” measurements*

The RHIC and LHC heavy ion programs have established that the measurement of single inclusive hadronic spectra yields a generic and quantitative manifestation of

jet quenching. Because the spectra in hadron–hadron collisions are steeply falling functions of p_T , if the hard partons produced in a heavy ion collision lose energy as they propagate through the strongly coupled plasma shifting the spectra leftward – to lower energy – is equivalent to depressing them. This effect is quantified via the measurement of the nuclear modification factor R_{AB}^h , which characterizes how the number of hadrons h produced in a collision between nucleus A and nucleus B differs from the number produced in an equivalent number of proton–proton collisions:

$$R_{AB}^h(p_T, \eta, \text{centrality}) = \frac{\frac{dN_{\text{medium}}^{AB \rightarrow h}}{dp_T d\eta}}{\langle N_{\text{coll}}^{AB} \rangle \frac{dN_{\text{vacuum}}^{pp \rightarrow h}}{dp_T d\eta}}. \quad (2.28)$$

Here, $\langle N_{\text{coll}}^{AB} \rangle$ is the average number of inelastic nucleon–nucleon collisions in $A+B$ collisions within a specified range of centralities. This number is typically determined by inferring the transverse density distribution of nucleons in a nucleus from the known radial density profile of nuclei, and then calculating the average number of collisions with the help of the inelastic nucleon–nucleon cross-section. This so-called Glauber calculation can be checked experimentally by independent means, for instance via the measurement of the nuclear modification factor for photons or Z -bosons discussed below.

The nuclear modification factor depends in general on the transverse momentum p_T and pseudorapidity η of the particle, the particle identity h , the centrality of the collision and the orientation of the particle trajectory with respect to the reaction plane (which is often averaged over). If R_{AB} deviates from 1 this reflects either medium effects or initial state effects – the parton distributions in A and B need not be simply related to those in correspondingly many protons. Measurements of R_{pA} in proton– A collisions (or R_{dA} in deuteron– A collisions which is a good proxy for R_{pA}) are used to determine whether an observed deviation of R_{AA} from 1 is due to initial state effects or the effects of parton energy loss in medium.

At mid-rapidity, RHIC data on R_{AA} show the following generic features.

(1) *Characteristic strong centrality dependence of R_{AA}*

By varying the centrality of a heavy ion collision, one changes the typical in-medium path length over which hard partons produced in these collisions must propagate through the dense matter. For the most central head-on collisions (e.g. 0%–10% centrality), the average L is large. For a peripheral collision (e.g. 80%–92% centrality), the average L is small. RHIC and LHC data (see Fig. 2.13) for charged hadrons show that for the most peripheral centrality bin, the nuclear modification factors are consistent with the absence of medium effects, while R_{AA} decreases monotonically with increasing centrality and reaches a suppression of about 0.2 (0.13) at RHIC (LHC) for $p_T \sim 5\text{--}10$ GeV




Figure 2.13 R_{AA} for charged hadrons as a function of p_T . Top: comparison of results for central and peripheral collisions at the LHC. Bottom: comparison of results from RHIC and LHC. Figures taken from [7].




Figure 2.14 The nuclear modification factor R_{AA} in the range up to transverse momenta $m_T = \sqrt{m^2 + p_T^2}$ of 100 GeV for the 10% most central Pb+Pb collisions at the LHC. Data are shown for charged hadrons, b -quarks identified via secondary J/ψ -decays, as well as for photons and the electroweak gauge bosons W and Z . The latter do not interact strongly with the medium and can hence emerge from heavy ion collisions unsuppressed and without energy loss. Data were compiled by the CMS collaboration from Refs. [265, 271, 267, 270, 272].

in the most central collisions. The suppression increases mildly with transverse momentum and persists up to the highest p_T experimentally measured so far, see Fig. 2.14. Figures 2.13 and 2.14 illustrate a direct manifestation of jet quenching: for $R_{AA} = 0.2$, 80% of the energetic hadrons that would be seen in the absence of a medium are gone.

(2) *Jet quenching is not observed in R_{dAu} and R_{pPb}*

In deuteron–gold collisions at RHIC, R_{dAu} is consistent with or greater than 1 for all centralities and all transverse momenta. Jet quenching is not observed. Very first data for R_{pPb} at the LHC support this conclusion [12]. In fact, the centrality dependence measured at RHIC is opposite to that seen in gold–gold collisions, with R_{dAu} reaching maximal values of around 1.5 for $p_T = 3\text{--}5$ GeV/c in the most central collisions [23, 15]. The high p_T hadrons are measured at or near mid-rapidity, meaning that they are well separated from the fragments of the struck gold nucleus. And, d-Au collisions produce at best

a much smaller volume of hot matter in the final state. In these collisions, therefore, the partons produced in hard scattering processes and tallied in R_{dAu} do not have to propagate any significant distance through matter after they are produced. The fact that R_{dAu} is consistent with or greater than 1 in these collisions therefore demonstrates that the jet quenching measured in R_{AuAu} is attributable to the propagation of the hard partons produced in heavy ion collisions through the medium that is present only in those collisions.

(3) *Photons, Z- and W-bosons are not quenched*

For single inclusive photon spectra in heavy ion collisions at RHIC, the nuclear modification factor shows only mild deviations from $R_{AuAu}^{\gamma} \approx 1$ [484]. Within errors, these are consistent with perturbative predictions that take into account the nuclear modifications of parton distribution functions (mainly the isospin difference between protons and nuclei) [64]. These statements apply also to photons produced in heavy ion collisions at LHC energies and to the electroweak gauge bosons produced in those collisions also, see Fig. 2.14. Since photons and electroweak gauge bosons, unlike partons or hadrons, do not interact strongly with the medium, this gives independent support that the jet quenching observed in heavy ion collisions is a final state effect. And, it provides experimental evidence in support of the Glauber-type calculation of the factor $\langle N_{\text{coll}}^{AA} \rangle$ in (2.28) discussed above. (That is, it provides experimental confirmation that the p+p data in Fig. 2.3 have indeed been scaled appropriately in order to use these data as a reference for nucleus–nucleus collisions with varying impact parameter.)

(4) *Species-independent suppression of R_{AA} at high p_T*

R_{AuAu}^h is independent of the species of the hadron h [24]. This eliminates the possibility that hadrons are formed within the medium and then lose energy upon propagating through the medium, since different hadrons would have different cross-sections for interaction with the medium. These data support the picture that the origin of the observed suppression is energy loss by a parton propagating through the medium prior to its hadronization.

(5) *R_{AA} for heavy-flavored and light-flavored hadrons is comparable.*

On general grounds in QCD, one expects that light-flavored partons lose more energy in the medium than heavy quarks [327]. At the time of this writing, there is no unambiguous experimental evidence for this mass hierarchy of parton energy loss. It is a matter of ongoing discussion to what extent the uncertainties in the existing data on the parton-mass dependence of jet quenching observables are already small enough to put interesting constraints on models of parton energy loss. More progress can be expected in the near future, once detector upgrades at RHIC and measurements at the LHC allow for differentiation of bottom quarks and charm quarks.

In short, these observations support a picture in which highly energetic partons are produced in high momentum transfer processes in heavy ion collisions as if they were produced in vacuum, but instead they find themselves propagating through a strongly coupled medium which causes them to lose a significant fraction of their initial energy. Jet quenching is a partonic final state effect that depends on the length of the medium through which the parton must propagate. It is expected to have many consequences in addition to the strong suppression of single inclusive hadron spectra, which tend to be dominated by the most energetic hadronic fragments of parent partons. As discussed at the beginning of this section, the entire parton fragmentation process is expected to be modified, with consequences for observables including multi-particle jet-like correlations and for calorimetric jet measurements including the dijet imbalance shown in Fig. 2.12.

2.3.2 Analyzing jet quenching

For concreteness, we shall focus in this section on those aspects of the analysis of jet quenching that bear upon the calculation of the nuclear modification factor R_{AA} defined in (2.28). We shall describe other aspects of the analysis of jet quenching more briefly, as needed, in subsequent sections. The single inclusive hadron spectra which define R_{AA} are typically calculated upon assuming that the modification of the spectra in nucleus–nucleus collisions relative to that in proton–proton collisions arises due to parton energy loss. This assumption is well supported by data, as we have described above. But, from a theoretical point of view it is an assumption, not backed up by any formal factorization theorem. Upon making this assumption, we write

$$d\sigma_{(\text{med})}^{AA \rightarrow h + \text{rest}} = \sum_f d\sigma_{(\text{vac})}^{AA \rightarrow f + X} \otimes P_f(\Delta E, L, \hat{q}, \dots) \otimes D_{f \rightarrow h}^{(\text{vac})}(z, \mu_F^2). \quad (2.29)$$

Here, \otimes denotes convolution in the energy fraction of the parton f and

$$d\sigma_{(\text{vac})}^{AA \rightarrow f + X} = \sum_{ijk} f_{i/A}(x_1, Q^2) \otimes f_{j/A}(x_2, Q^2) \otimes \hat{\sigma}_{ij \rightarrow f+k}, \quad (2.30)$$

where $f_{i/A}(x, Q^2)$ are the nuclear parton distribution functions and $\sigma_{ij \rightarrow f+k}$ are the perturbatively calculable partonic cross-sections. The medium dependence enters via the function $P_f(\Delta E, L, \hat{q}, \dots)$, which characterizes the probability that a parton f produced with cross-section $\sigma_{ij \rightarrow f+k}$ loses energy ΔE while propagating over a path length L in a medium. This probability depends of course on properties of the medium, which are represented schematically in this formula

by the symbol \hat{q} , the jet quenching parameter. We shall see below that in the high parton energy limit, the properties of the medium enter P_f only through one parameter, and in that limit \hat{q} can be defined precisely. At non-asymptotic parton energies, \hat{q} in (2.29) is a place-holder, representing all relevant attributes of the medium. It is often conventional to refer to the combination of P_f and $D_{f \rightarrow h}^{(\text{vac})}$ together as a modified fragmentation function. It is only in the limit of high parton energy where one can be sure that the parton emerges from the medium before fragmenting into hadrons in vacuum that these two functions can be cleanly separated as we have done in (2.29). This aspect of the ansatz (2.29) is supported by the data: as we have described above, all hadrons exhibit the same suppression factor indicating that R_{AA} is due to partonic energy loss, before hadronization.

The dynamics of how parton energy is lost to the medium is specified in terms of the probability $P_f(\Delta E, L, \hat{q}, \dots)$. In the high parton energy limit, the parton loses energy dominantly by inelastic processes that are the QCD analog of bremsstrahlung: the parton radiates gluons as it interacts with the medium. It is a familiar fact from electromagnetism that bremsstrahlung dominates the loss of energy of an electron moving through matter in the high energy limit. The same is true in calculations of QCD parton energy loss in the high energy limit, as established first in Refs. [421, 98, 817]. The hard parton undergoes multiple inelastic interactions with the spatially extended medium, and this induces gluon bremsstrahlung. Here and throughout, by the high parton energy limit we mean the combined set of limits that can be summarized as:

$$E \gg \omega \gg |\mathbf{k}|, |\mathbf{q}| \equiv \left| \sum_i \mathbf{q}_i \right| \gg T, \Lambda_{\text{QCD}}, \quad (2.31)$$

where E is the energy of the high energy projectile parton, where ω and \mathbf{k} are the typical energy and momentum of the gluons radiated in the elementary radiative processes $q \rightarrow qg$ or $g \rightarrow gg$, and where \mathbf{q} is the transverse momentum (transverse to its initial direction) accumulated by the projectile parton due to many radiative interactions in the medium, and where T and Λ_{QCD} represent any energy scales that characterize the properties of the medium itself. This set of approximations underlies all the pioneering analytical calculations of radiative parton energy loss [98, 817, 797, 420, 414, 794]. The premise of the analysis is the assumption that QCD at scales of order $|\mathbf{k}|$ and $|\mathbf{q}|$ is weakly coupled, even if the medium (with its lower characteristic energy scales of order T and Λ_{QCD}) is strongly coupled. We shall spend most of this section on the analysis valid in this high parton energy limit, in which case all we need to ask of analyses of strongly coupled gauge theories with gravity duals is insight into those properties of the strongly coupled

medium that enter into the calculation of jet quenching in QCD. However, the analysis based upon the limits (2.31) may not be under quantitative control when applied to data, since gluon radiation outside the eikonal region $E \gg \omega \gg |\mathbf{k}|$ may be relevant for parton energy loss even at very high parton energies. Moreover, in Section 8.1 we shall see that for partons with low enough energy E physics at all scales in the problem up to E is strongly coupled and new approaches are needed. (This applies for any E in a conformal theory like $\mathcal{N} = 4$ SYM.)

A Gluon production in the eikonal limit and \hat{q}

In the discussion above, we have argued that in the eikonal limit of asymptotic parton energies, a single jet quenching parameter \hat{q} may characterize the medium-modification of gluon radiation. We now discuss an illustrative calculation in which this relation can be made explicit. We start by considering a high energy parton. In the rest frame of this parton, the target that is spatially extended but of finite thickness appears Lorentz contracted, so in the projectile rest frame the parton propagates through the target in a short period of time and the transverse position of the projectile does not change during the propagation. So, at ultra-relativistic energies, the main effect of the target on the projectile is a “rotation” of the parton’s color due to the color field of the target. These rotation phases are given by Wilson lines along the (straight line) trajectories of the propagating projectile:

$$W(\mathbf{x}) = \mathcal{P} \exp\{i \int dz^- T^a A_a^+(\mathbf{x}, z^-)\}. \quad (2.32)$$

Here, \mathbf{x} is the *transverse* position of the projectile – which does not change as the parton propagates at the speed of light along the $z^- \equiv (z - t)/\sqrt{2}$ lightlike direction. A^+ is the large component of the target color field and T^a is the generator of $SU(N)$ in the representation corresponding to the given projectile – fundamental if the hard parton is a quark and adjoint if it is a gluon. The eikonal approach to scattering treats the (unphysical, in the case of colored projectiles) setting in which the projectile impinges on the target from outside, after propagating for an arbitrarily long time and building up a fully developed coherent Coulomb cloud $\sim g \frac{\mathbf{x}_i}{x_i^2}$ of gluons dressing the bare projectile. (This cloud is often referred to as a non-Abelian Weizsäcker–Williams field). The interaction of this dressed projectile with the target results in an eikonal phase (Wilson line) for the projectile itself and for each gluon in the cloud. Gluon radiation then corresponds to the decoherence of components of the dressed projectile that pick up different phases. Analysis of this problem yields a calculation of $N_{\text{prod}}(\mathbf{k})$, the number of radiated gluons with momentum \mathbf{k} , with the result:

$$\begin{aligned}
N_{\text{prod}}(\mathbf{k}) = & \\
& \frac{\alpha_s C_F}{2\pi} \int d\mathbf{x} d\mathbf{y} e^{i\mathbf{k}\cdot(\mathbf{x}-\mathbf{y})} \frac{\mathbf{x}\cdot\mathbf{y}}{\mathbf{x}^2 \mathbf{y}^2} \left[1 - \frac{1}{N^2 - 1} \langle \text{Tr} [W^{A\dagger}(\mathbf{x}) W^A(\mathbf{0})] \rangle \right. \\
& \quad - \frac{1}{N^2 - 1} \langle \text{Tr} [W^{A\dagger}(\mathbf{y}) W^A(\mathbf{0})] \rangle \\
& \quad \left. + \frac{1}{N^2 - 1} \langle \text{Tr} [W^{A\dagger}(\mathbf{y}) W^A(\mathbf{x})] \rangle \right], \quad (2.33)
\end{aligned}$$

where the C_F prefactor is for the case where the projectile is a quark in the fundamental representation, where the projectile is located at transverse position $\mathbf{0}$, and where the $\langle \dots \rangle$ denotes averaging over the gluon fields of the target. If the target is in thermal equilibrium, these are thermal averages. (See Refs. [548, 550] for details.)

Although the simple result (2.33) is not applicable to the physically relevant case, as we shall describe in detail below, we can nevertheless glean insights from it that will prove relevant. We note that the entire medium-dependence of the gluon number spectrum (2.33) is determined by target expectation values of the form $\langle \text{Tr} [W^{A\dagger}(\mathbf{x}) W^A(\mathbf{y})] \rangle$ of two eikonal Wilson lines. The jet quenching parameter \hat{q} that will appear below defines the fall-off properties of this correlation function in the transverse direction $L \equiv |\mathbf{x} - \mathbf{y}|$:

$$\langle \text{Tr} [W^A(\mathcal{C})] \rangle \approx \exp \left[-\frac{1}{4\sqrt{2}} \hat{q} L^- L^2 \right] \quad (2.34)$$

in the limit of small L , with L^- (the extent of the target along the z^- direction) assumed large but finite [582, 584]. Here, the contour \mathcal{C} traverses a distance L^- along the light cone at transverse position \mathbf{x} , and it returns at transverse position \mathbf{y} . These two long straight light-like lines are connected by short transverse segments located at $z^- = \pm L^-/2$, far outside the target. We see from the form of (2.33) that $|\mathbf{k}|$ and L are conjugate: the radiation of gluons with momentum $|\mathbf{k}|$ is determined by Wilson loops with transverse extent $L \sim 1/|\mathbf{k}|$. This means that in the limit (2.31), the only property of the medium that enters (2.33) is \hat{q} . Furthermore, inserting (2.34) into (2.33) yields the result that the gluons that are produced have a typical \mathbf{k}^2 that is of order $\hat{q} L^-$. This suggests that \hat{q} can be interpreted as the transverse momentum squared picked up by the hard parton per distance L^- that it travels, an interpretation that can be validated more rigorously via other calculations [251, 317].

B Parton energy loss in a finite medium

The reason that the eikonal formalism cannot be applied verbatim to the problem of parton energy loss in heavy ion collisions is that the high energy partons we

wish to study do not impinge on the target from some distant production site. They are produced within the same collision that produces the medium whose properties they subsequently probe. As a consequence, they are produced with significant virtuality. This means that even if there were no medium present, they would radiate copiously. They would fragment in what is known in QCD as a parton shower. The analysis of medium-induced parton energy loss then requires understanding the interference between radiation in vacuum and the medium-induced bremsstrahlung radiation. It turns out that the resulting interference resolves longitudinal distances in the target [98, 817, 797, 420], meaning that its description goes beyond the eikonal approximation. The analysis of parton energy loss in the high energy limit (2.31) must include terms that are subleading in $1/E$, and therefore not present in the eikonal approximation, that describe the leading interference effects. To keep these $\mathcal{O}(1/E)$ effects, one must replace eikonal Wilson lines by retarded Green's functions that describe the propagation of a particle with energy E from position z_1^-, \mathbf{x}_1 to position z_2^-, \mathbf{x}_2 without assuming $\mathbf{x}_1 = \mathbf{x}_2$ [817, 544, 798]. (In the $E \rightarrow \infty$ limit, $\mathbf{x}_1 = \mathbf{x}_2$ and the eikonal Wilson line is recovered.) It nevertheless turns out that even after Wilson lines are replaced by Green's functions the only attribute of the medium that arises in the analysis, in the limit (2.31), is the jet quenching parameter \hat{q} defined in (2.34) that already arose in the eikonal approximation [817, 797].

We shall not present the derivation, but it is worth giving the complete (albeit somewhat formal) result for the distribution of gluons with energy ω and transverse momentum \mathbf{k} that a high energy parton produced within a medium radiates:

$$\omega \frac{dI}{d\omega d\mathbf{k}} = \frac{\alpha_s C_R}{(2\pi)^2 \omega^2} 2\text{Re} \int_{\xi_0}^{\infty} dy_l \int_{y_l}^{\infty} d\bar{y}_l \int d\mathbf{u} e^{-i\mathbf{k}\cdot\mathbf{u}} \exp \left[-\frac{1}{4} \int_{\bar{y}_l}^{\infty} d\xi \hat{q}(\xi) \mathbf{u}^2 \right] \\ \times \frac{\partial}{\partial \mathbf{x}} \cdot \frac{\partial}{\partial \mathbf{u}} \int_{\mathbf{x}=\mathbf{r}(y_l)=\mathbf{0}}^{\mathbf{u}=\mathbf{r}(\bar{y}_l)} \mathcal{D}\mathbf{r} \exp \left[\int_{y_l}^{\bar{y}_l} d\xi \left(\frac{i\omega}{2} \dot{\mathbf{r}}^2 - \frac{1}{4} \hat{q}(\xi) \mathbf{r}^2 \right) \right]. \quad (2.35)$$

We now walk through the notation in this expression. The Casimir operator C_R is in the representation of the projectile parton. The integration variables ξ , y_l and \bar{y}_l are all positions along the z^- lightcone direction. ξ_0 is the z^- at which the projectile parton was created in a hard scattering process. Since we are not taking this to $-\infty$, the projectile is not assumed on shell. The projectile parton was created at the transverse position $\mathbf{x} = \mathbf{0}$. The integration variable \mathbf{u} is also a transverse position variable, conjugate to \mathbf{k} . The path integral is over all possible paths $\mathbf{r}(\xi)$ going from $\mathbf{r}(y_l) = \mathbf{0}$ to $\mathbf{r}(\bar{y}_l) = \mathbf{u}$. The derivation of (2.35) proceeds by writing $dI/d\omega d\mathbf{k}$ in terms of a pair of retarded Green's functions in their path-integral representations, one of which describes the radiated gluon in the amplitude, radiated at y_l , and the other of which describes the radiated gluon in the conjugate amplitude, radiated at

\bar{y}_l . The expression (2.35) then follows after a lengthy but purely technical calculation. The properties of the medium enter (2.35) only through the jet quenching parameter $\hat{q}(\xi)$. There are many, closely related formulations of parton energy loss. The first works on this subject are by Baier, Dokshitzer, Mueller, Peigné, Schiff (BDMPS) [98] and independently by Zakharov (Z) [817]. The expression (2.35) was derived in the so-called path-integral approach in Refs. [817, 797]. This and related approaches to parton energy loss in QCD have been developed by many authors [101, 414, 795, 76, 97, 551, 422, 796, 710, 821, 592, 700, 818, 122, 68, 13, 790, 463, 276, 591, 236]. Recent reviews include [251, 799].

The result (2.35) is both formal and complicated. However, its central qualitative consequences can be characterized almost by dimensional analysis. For simplicity, we consider first the case that the jet quenching parameter does not depend on the position in the medium, $\hat{q} = \hat{q}(\xi)$ (for a generalization, see the next subsection). All dimensionful quantities can be scaled out of (2.35) if ω is measured in units of the so-called characteristic gluon energy

$$\omega_c \equiv \hat{q}(L^-)^2, \quad (2.36)$$

and the transverse momentum \mathbf{k}^2 in units of $\hat{q}L^-$ [724]. In a numerical analysis of (2.35), one finds that the transverse momentum distribution of radiated gluons scales indeed with $\hat{q}L^-$, as expected for the transverse momentum due to the Brownian motion in momentum space that is induced by multiple small angle scatterings. If one integrates the gluon distribution (2.35) over transverse momentum and takes the upper limit of the \mathbf{k} -integration to infinity, one recovers [724] an analytical expression first derived by Baier, Dokshitzer, Mueller, Peigné and Schiff [98]:

$$\omega \frac{dI_{\text{BDMPS}}}{d\omega} = \frac{2\alpha_s C_R}{\pi} \ln \left| \cos \left[(1+i) \sqrt{\frac{\omega_c}{2\omega}} \right] \right|, \quad (2.37)$$

which yields the limiting cases

$$\omega \frac{dI_{\text{BDMPS}}}{d\omega} \simeq \frac{2\alpha_s C_R}{\pi} \begin{cases} \sqrt{\frac{\omega_c}{2\omega}} & \text{for } \omega \ll \omega_c, \\ \frac{1}{12} \left(\frac{\omega_c}{\omega}\right)^2 & \text{for } \omega \gg \omega_c, \end{cases} \quad (2.38)$$

for small and large gluon energies. In the soft gluon limit, the BDMPS spectrum (2.37) displays the characteristic $1/\sqrt{\omega}$ dependence, which persists up to a gluon energy of the order of the characteristic gluon energy (2.36). Hence, ω_c can be viewed as an effective energy cut-off, above which the contribution of medium-induced gluon radiation is negligible. These analytical limits provide a rather accurate characterization of the full numerical result. In particular, one expects from the above expressions that the average parton energy loss $\langle \Delta E \rangle$, obtained by

integrating (2.35) over \mathbf{k} and ω , is proportional to $\propto \int_0^{\omega_c} d\omega \sqrt{\omega_c}/\omega \propto \omega_c$. One finds indeed

$$\langle \Delta E \rangle_{\text{BDMPS}} \equiv \int_0^\infty d\omega \omega \frac{dI_{\text{BDMPS}}}{d\omega} = \frac{\alpha_s C_R}{2} \omega_c = \frac{\alpha_s C_R}{2} \hat{q} (L^-)^2. \quad (2.39)$$

This is the well-known $(L^-)^2$ -dependence of the average radiative parton energy loss [99, 98, 817]. In summary, the main qualitative properties of the medium-induced gluon energy distribution (2.35) are the scaling of \mathbf{k}^2 with $\hat{q} L^-$ dictated by Brownian motion in transverse momentum space, the $1/\sqrt{\omega}$ dependence of the \mathbf{k} -integrated distribution characteristic of the non-Abelian Landau–Pomeranchuk–Migdal (LPM) effect, and the resulting $(L^-)^2$ -dependence of the average parton energy loss.

C From medium-induced gluon radiation to jet quenching models

We now discuss how to relate calculations of medium-induced gluon radiation to data on jet quenching in heavy ion collisions. To this end, we recall first how in QCD in the vacuum, partons produced with high transverse energy in hadronic collisions evolve into hadronic fragments. Such highly energetic partons typically undergo a so-called parton shower, that is a series of partonic $1 \rightarrow 2$ splittings in which they degrade their high initial virtuality. It is only at the end of this parton shower that hadronization, i.e. the transition from partonic to hadronic degrees of freedom, sets in. This vacuum parton shower is calculable in QCD perturbation theory and it is theoretically well-understood. It determines, for instance, the so-called scale dependence of fragmentation functions (i.e. the μ_F^2 dependence in Eq. (2.29)), as well as many characteristics of the distribution of the hadronic fragments of jets. Most generally, the phenomenology of jet quenching aims at modeling how the passage through dense matter affects this QCD parton shower and what it reveals about the properties of the matter through which it passes. The theorist's task is therefore to formulate a medium-modified parton shower that is consistent with QCD-based calculations of parton energy loss. Since the theoretical understanding of jet quenching is still incomplete, this task requires elements of phenomenological modeling to relate QCD-based calculations to jet quenching data. Here, we discuss one particularly simple and widely used jet quenching model in some detail. We then comment on open challenges and further developments.

As explained above, the basic building block of a QCD parton shower in the vacuum is the elementary partonic $1 \rightarrow 2$ splitting function. The splittings of a quark into a quark and a gluon ($q \rightarrow qg$) and of a gluon into two gluons ($g \rightarrow gg$) dominate kinematically. Both processes can be viewed as gluon radiation of a parent parton in the vacuum. One particularly simple way of formulating a medium-modified parton shower is then to replace the $1 \rightarrow 2$ splitting in the

vacuum by a calculation of medium-induced $1 \rightarrow 2$ splitting. In particular, our derivation (2.35) of medium-induced gluon radiation of a quark or gluon can be used in this set-up as a medium-induced $1 \rightarrow 2$ splitting. This idea has been implemented in several Monte Carlo programs that simulate the entire parton shower, see e.g. [818, 67]. A discussion of the complexity of Monte Carlo programs for final state parton showers and the technical and conceptual differences in existing Monte Carlo implementations lies outside the scope of this book. Instead, we restrict our discussion to a version of such jet quenching models that is limited to describing single inclusive hadron spectra in nucleus–nucleus collisions. As seen from expression (2.29), the quenching of such spectra can be described by the probability $P(\Delta E)$ that the initial parton loses a fraction $\Delta E/E$ of its total energy via medium effects. We now sketch how the resulting $P(\Delta E)$ in (2.29) can be estimated. If gluons are emitted independently, $P(\Delta E)$ is the normalized sum of the emission probabilities for an arbitrary number of n gluons which carry away the total energy ΔE [103]:

$$P(\Delta E) = \exp \left[- \int_0^\infty d\omega \frac{dI}{d\omega} \right] \sum_{n=0}^{\infty} \frac{1}{n!} \left[\prod_{i=1}^n \int d\omega_i \frac{dI(\omega_i)}{d\omega} \right] \delta \left(\Delta E - \sum_{i=1}^n \omega_i \right). \quad (2.40)$$

Here, the factor $\exp \left[- \int_0^\infty d\omega \frac{dI}{d\omega} \right]$ denotes the probability that no energy loss occurs. This factor ensures that $P(\Delta E)$ is properly normalized, namely $\int d\Delta E P(\Delta E) = 1$. Equation (2.40) thus resums the effects that arbitrarily many independent medium-induced gluon radiations (2.35) have on the energy of the most energetic parton in the shower. On average, this parton will suffer an additional medium-induced mean energy loss

$$\langle \Delta E \rangle = \int d\Delta E (\Delta E) P(\Delta E) = \int d\omega \omega \frac{dI}{d\omega}, \quad (2.41)$$

which is consistent with (2.39) above.

The phenomenological strategy for constraining the jet quenching parameter \hat{q} is then based on comparing the single inclusive hadron spectrum (2.29) in nucleus–nucleus collisions to data. Here, the jet quenching parameter \hat{q} enters via the probability $P(\Delta E)$ that the initial parton loses a fraction $\Delta E/E$ of its total energy via medium effects. We note, however, that such a program of determining \hat{q} has to control various complications.

In particular, as seen from Fig. 2.3, single inclusive hadron spectra are distributions which fall steeply with p_T . Since $P(\Delta E)$ is a very broad distribution, not peaked around its mean, the modifications which parton energy loss induce on spectra cannot be characterized by an average energy loss. Rather, what matters for a steeply falling distribution is not how much energy a parton loses on average, but

which fraction of all the partons escapes with much less than the average energy loss [103]. This so-called trigger-bias effect is quantitatively very important, and can be accounted for by the probability distribution (2.40). However, this trigger-bias effect is also a surface-bias effect: those partons that escape with the smallest in-medium path length have the highest probability of contributing to the single inclusive hadron spectrum. As a consequence, it is important that jet quenching models embed the propagation of the highly energetic parton in a realistic spatial and temporal structure of a heavy ion collision. This includes a suitable probability distribution of the production points of the hard partons in the transverse plane, and a resulting realistic distribution of the in-medium path lengths L^- over which the parton propagates through the medium.

Another important aspect is that as a consequence of longitudinal and transverse flow, the density of the medium degrades significantly during this time period L^- , and approximating $\hat{q}(\xi)$ by a constant value \hat{q} is not a good approximation. In general, this motivates the formulation of jet quenching models for which the probability of interactions between the medium and the jet decreases with time. For the medium-induced gluon radiation (2.35), analytical solutions in the saddle point approximation are known if one approximates the ξ dependence of the jet quenching parameter as $\hat{q}(\xi) = \hat{q}_0 (\xi_0/\xi)^\alpha$ [100] with α between 1 and 3. This range of α s scans the range of phenomenologically relevant cases between one-dimensional longitudinal expansion ($\alpha = 1$; ‘‘Bjorken expansion’’) and scenarios which also account for the transverse expansion $1 < \alpha < 3$. Remarkably, one finds that irrespective of the value of α , for fixed in-medium path length $L^-/\sqrt{2}$ the transverse momentum integrated gluon energy distribution (2.35) has the same ω -dependence if $\hat{q}(\xi)$ is simply replaced by a constant given by the linear line-averaged transport coefficient [723]

$$\langle \hat{q} \rangle \equiv \frac{1}{2L^{-2}} \int_{\xi_0}^{\xi_0+L^-} d\xi (\xi - \xi_0) \hat{q}(\xi). \quad (2.42)$$

In practice, this means that comparisons of different parton energy loss calculations to data can be performed as if the medium were static. The line-averaged transport coefficient $\langle \hat{q} \rangle$ determined in this way can then be related via (2.42) to the transport coefficient at a given time, once a model for the expansion of the medium is specified. Hence, we can continue our discussion for the case $\hat{q}(\xi) = \hat{q}$ without loss of generality.

Historically, the first class of jet quenching models proceeded by implementing $P(\Delta E)$ in a model in which hard scattering events are distributed with suitable probability at locations in the transverse plane, and are then propagated through the medium. To this end, the expanding and cooling plasma was given either by a simple parametrization, or it was modeled hydrodynamically. The jet quenching

parameter was assumed for instance to be given in terms of the time-dependent energy density ε ,

$$\hat{q} = 2K\varepsilon^{3/4}, \quad (2.43)$$

and the parameter K was then obtained by comparing jet quenching calculations to data. For a weakly coupled quark–gluon plasma, Baier had argued that $K \approx 1$ [97]. In contrast, in fitting to data from PHENIX and STAR at RHIC, several studies obtained significantly larger values. For instance, the PHENIX collaboration [20] use the jet quenching model from Ref. [305] and quote a jet quenching parameter which is constrained by the experimental data to lie within the range $13.2_{-3.2}^{+2.1}$ or $13.2_{-5.2}^{+6.3}$ GeV²/fm at the one or two standard deviation levels, respectively. This translates into the estimate

$$K = 4.1 \pm 0.6, \quad (2.44)$$

at one standard deviation. It is important to realize that the quoted errors arise only from the experimental uncertainties and do not incorporate the “systematic uncertainty” arising from the choices made in the formulation of the theoretical model to which the PHENIX authors compare their data. We will compare the result (2.44) to calculations done for strongly coupled plasmas in gauge theories with dual gravitational descriptions in Section 8.5.

As mentioned already, we have limited our presentation in this section to one particularly simple jet quenching model in order to showcase an example of an explicit connection between a property of dense QCD matter, namely \hat{q} , and the measurement of quenched hadron spectra. The large values of \hat{q} extracted from several model comparisons with data motivate the need to turn to strong coupling techniques for describing jet–medium interactions. Theoretical work is currently under way to improve on significant assumptions of the simple model approach discussed here, for example the assumption that a calculation of parton energy loss in the eikonal limit (2.31) can provide sufficiently accurate results for the phenomenologically relevant kinematics [70] or the assumption that effects of multiple gluon emission can be taken into account probabilistically via (2.40) without accounting for (destructive) quantum interference in the emission of more than one gluon and without tracing the energy loss of the projectile after a gluon emission. Many of the resulting model uncertainties can only be controlled by going beyond the kinematical limit (2.31). Going beyond this limit is also required if one is to assess the possible role of parton energy loss via elastic interactions with the medium. Another current challenge is to go beyond single inclusive hadron spectra and to formulate models that can account for the medium-modified fragmentation of entire jets as seen in Fig. 2.12. All these open questions are subjects of ongoing research and they lie outside the scope of the present book. As described

in this section, the basic problem is to understand how a QCD final state parton shower is modified in the presence of a hot and dense plasma. In the coming years, progress on this question can be expected from a tight interplay between theory and experiment that includes a more and more detailed characterization of parton fragmentation via a suite of upcoming further measurements including jet fragmentation functions, γ - or Z -triggered jet distributions, heavy flavored single inclusive hadron spectra and b -tagged jets.

2.4 Quarkonia in hot matter

One way of thinking about the operational meaning of the statement that quark–gluon plasma is deconfined is to ask what prevents the formation of a meson within quark–gluon plasma. The answer is that the attractive force between a quark and an antiquark which are separated by a distance of order the size of a meson is screened by the presence of the quark–gluon plasma between them. This poses a quantitative question: how close together do the quark and antiquark have to be in order for their attraction not to be screened? How close together do they have to be in order for them to feel the same attraction that they would feel if they were in vacuum? It was first suggested by Matsui and Satz [609] in 1986 that measurements of how many quarkonia – mesons made of a heavy quark–antiquark pair – are produced in heavy ion collisions could be used as a tool with which to answer this question, because they are significantly smaller than typical mesons or baryons.

The generic term quarkonium refers to the charm–anticharm or charmonium, mesons (J/ψ , ψ' , χ_c , ...) and the bottom–antibottom, or bottomonium, mesons (Υ , Υ' , ...). The first quarkonium state that was discovered was the $1s$ state of the $c\bar{c}$ bound system, the J/ψ . It is roughly half the size of a typical meson like the ρ . The bottomonium $1s$ state, the Υ , is smaller again by roughly another factor of two. It is therefore expected that if one can study quark–gluon plasma in a series of experiments with steadily increasing temperature, J/ψ mesons survive as bound states in the quark–gluon plasma up to some dissociation temperature that is higher than the crossover temperature (at which generic mesons and baryons made of light quarks fall apart) and Υ mesons survive as bound states up to some even higher temperature. More realistically, what Matsui and Satz suggested is that if high energy heavy ion collisions create deconfined quark–gluon plasma that is hot enough, then color screening would prevent charm and anticharm quarks from binding to each other in the deconfined interior of the droplet of matter produced in the collision, and as a result the number of J/ψ mesons produced in the collisions would be suppressed. However, bottomonium mesons in the Υ $1s$ state should be able to bind, and the rate of production of these mesons should therefore not




Figure 2.15 Schematic picture of the dissociation of a $Q\bar{Q}$ -pair in hot QCD matter due to color screening. Figure taken from Ref. [728]. The straight black lines attached to the heavy Q and \bar{Q} indicate, that these quarks are external probes, in contrast to the dynamical quarks within the quark–gluon plasma. Figure taken from Ref. [728].

be suppressed, until such high temperatures are achieved that the quark–antiquark attraction is screened even on the short length scale corresponding to the size of the Υ meson in its $1s$ state.

To study this effect, Matsui and Satz suggested comparing the temperature dependence of the screening length for the quark–antiquark force, which can be obtained from lattice QCD calculations, with the J/ψ meson radius calculated in charmonium models. They then discussed the feasibility to detect this effect clearly in the mass spectrum of e^+e^- dilepton pairs. Between 1986, when Matsui and Satz launched this line of investigation, suggesting it as a quantitative means of characterizing the formation and properties of deconfined matter, and today we know of no other measurement that has been advocated as a more direct experimental signature for the deconfinement transition. And, there is hardly any other measurement whose phenomenological analysis has turned out to be more involved. In this section, we shall describe both the appeal of studying quarkonia in the hot matter produced in heavy ion collisions and the practical difficulties. The theoretical basis for the argument of Matsui and Satz has evolved considerably within the last two decades [728]. Moreover, the debate over how to interpret these measurements is by now informed by data on J/ψ suppression in nucleus–nucleus collisions at the CERN SPS [43, 72], at RHIC [19] and at the LHC [1]. There is also a good possibility that qualitatively novel information will become accessible in future high statistics runs at RHIC and LHC.

A sketch of the basic idea of Matsui and Satz is shown in Fig. 2.15. In very general terms, one expects that the attractive interaction between the heavy quark and antiquark in a putative bound state is sensitive to the medium in which the heavy particles are embedded, and that this attraction weakens with increasing temperature. If the distance between the heavy quark and antiquark is much smaller than $1/T$, there will not be much quark–gluon plasma between them. Equivalently, typical momentum scales in the medium are of order the temperature T , and so the

medium cannot resolve the separation between the quark–antiquark pair if they are much closer together than $1/T$. However, if the distance is larger, then the bound state is resolved, and the color charges of the heavy quarks are screened by the medium, see Fig. 2.15. For the temperatures T that one expects to attain in a heavy ion collision in which quark–gluon plasma with a temperature T is created, only those quarkonia with radii that are smaller than some length scale of order $1/T$ can form. These basic arguments support the idea that quarkonium production rates are an indicator of whether quark–gluon plasma is produced, and at what temperature.

In Section 3.3, we review lattice calculations of the heavy quark static free energy $F_{Q\bar{Q}}(r)$. This static potential is typically defined via how the correlation function of a pair of Polyakov loops, namely test quarks at fixed spatial positions whose worldlines wrap around the periodic Euclidean time direction, falls off as the separation between the test quarks is increased. This static potential is renormalized such that it matches the zero temperature result at small distances. Calculations of $F_{Q\bar{Q}}(r)$ were the earliest lattice results which substantiated the core idea that a quarkonium bound state, placed in hot QCD matter, dissociates (“melts”) above a critical temperature. As we now discuss, phenomenological models of quarkonium in matter are based upon interpreting $F_{Q\bar{Q}}(r)$ as the potential in a Schrödinger equation whose eigenvalues and eigenfunctions describe the heavy Q – \bar{Q} bound states. There is no rigorous basis for this line of reasoning, and if pushed too far it faces various conceptual challenges as we shall discuss in Section 3.3. However, these models remain valuable as a source of semi-quantitative intuition.

At zero temperature, lattice results for $F_{Q\bar{Q}}(r)$ in QCD without dynamical quarks are well approximated by the ansatz $F_{Q\bar{Q}}(r) = \sigma r - \frac{\alpha}{r}$, where the linear term that dominates at long distance is characterized by the string tension $\sigma \simeq 0.2 \text{ GeV}^2$ and the perturbative Coulomb term α/r is dominant at short distances. In QCD with dynamical quarks, beyond some radius r_c the potential flattens because as the distance between the external Q and \bar{Q} is increased, it becomes energetically favorable to break the color flux tube connecting them by producing a light quark–antiquark pair from the vacuum which, in a sense, screens the potential. With increasing temperature, the distance r_c decreases, that is, the colors of Q and \bar{Q} are screened from each other at increasingly shorter distances. This is seen clearly in the Fig. 3.5 in Section 3.3. These lattice results are well parametrized by a screened potential of the form [728, 521]

$$F_{Q\bar{Q}}(r) = -\frac{\alpha}{r} + \sigma r \left(\frac{1 - e^{-\mu r}}{\mu r} \right), \quad (2.45)$$

where $\mu \equiv \mu(T)$ can be interpreted at high temperatures as a temperature-dependent Debye screening mass. For suitably chosen $\mu(T)$, this ansatz reproduces the flattening of the potential found in lattice calculations at the finite large distance

value $F_{Q\bar{Q}}(\infty) = \sigma/\mu(T)$. Taking this $Q\bar{Q}$ free energy $F_{Q\bar{Q}}(r, T)$ as the potential in a Schrödinger equation, one may try to determine which bound states in this potential remain, as the potential is weakened as the temperature increases. Such potential model studies have led to predictions of the dissociation temperatures T_d of the charmonium family, which range from $T_d(J/\psi) \simeq 2.1 T_c$ for J/ψ to $T_d(\psi') \simeq 1.1 T_c$ for the more loosely bound and therefore larger $2s$ state. The deeply bound, small, $1s$ state of the bottomonium family is estimated to have a dissociation temperature $T_d(\Upsilon(1S)) > 4 T_c$, while dissociation temperatures for the corresponding $2s$ and $3s$ states were estimated to lie at $1.6 T_c$ and $1.2 T_c$ respectively [728, 521]. Because the leap from the static quark–antiquark potential to a Schrödinger equation is not rigorously justified, the uncertainties in quantitative results obtained from these potential models are difficult to estimate. (For more details on why this is so, see Section 3.3.) However, these models with their inputs from lattice QCD calculations do provide qualitative support for the central idea of Matsui and Satz that quarkonia melt in hot QCD matter, and they provide support for the qualitative expectation that this melting proceeds *sequentially*, with smaller bound states dissociating at a higher temperature.

Figure 2.16 shows data of the Upsilon resonances $1s$, $2s$ and $3s$ in the dimuon invariant mass distribution measured by the CMS collaboration in p+p and Pb+Pb collisions at the LHC. While all three resonance states are clearly visible in p+p collisions, the higher excited $2s$ and $3s$ states are strongly suppressed if not absent in Pb+Pb collisions. These data are the most direct experimental support to date for the sequential quarkonium suppression pattern which is a generic prediction of all models of quarkonium suppression, and according to which only the tightly bound $1s$ state with $T_d(\Upsilon(1S)) > 4 T_c$ is expected to survive in the hot and dense QCD matter produced at the LHC. A more detailed discussion of these data requires the understanding of so-called feed-down corrections, which are contributions to the $1s$ yield from the decay of higher excited states. In proton–proton collisions, a significant fraction ($\sim 40\%$) of the observed $\Upsilon(1s)$ mesons arises from the production of the excited $2s$ and $3s$ states which subsequently decay to $\Upsilon(1s)$. Therefore, if the higher excited states melt in the hot matter, one expects that their “feed-down” to the $\Upsilon(1s)$ state is absent and the measured yield of $\Upsilon(1s)$ is reduced accordingly. Indeed, despite significant experimental uncertainties that still exist for these very first data, there is already evidence for a reduction in the number of $\Upsilon(1s)$ mesons produced in Pb+Pb collisions that can be explained in just this way. The analogous argument has also been instrumental in interpreting earlier data on the suppression of the yield of J/ψ mesons in nucleus–nucleus collisions at the CERN SPS [43, 72] and at RHIC in terms of the dissociation of higher excited states like the ψ' and χ_c whose decays contribute to the J/ψ yield in proton–proton collisions. In particular, assuming that directly produced $1s$ charmonium states survive and




Figure 2.16 The invariant mass distribution of dimuons in Pb+Pb (above) and p+p (below) collisions measured by the CMS collaboration. In comparison to the benchmark measurement in p+p, the higher Υ resonances are strongly suppressed. Figures taken from Ref. [269].

higher excited states melt completely at CERN SPS and RHIC energies provides a natural interpretation for the fact that the suppression of the J/ψ yield and its centrality dependence in nucleus–nucleus collisions at the CERN SPS and RHIC are comparable. However, since these earlier studies did not have experimental access

to the excited quarkonium states, their support for a picture of sequential melting was less direct.

The analysis of data on charmonium has to take into account a significant number of important confounding effects. Here, we cannot discuss the phenomenology of these effects in detail, but we provide a list of the most important ones.

(1) *Cold nuclear matter effects*

The interaction of the heavy quark–antiquark pair with ordinary confined hadronic matter can be a significant source of quarkonium dissociation [202]. The operational procedure for separating such hadronic phase effects (often referred to as “cold nuclear matter effects”) is to measure them separately in proton–nucleus collisions [525], and to establish then to what extent the number of J/ψ mesons produced in nucleus–nucleus collisions drops below the yield extrapolated from proton–nucleus collisions [26].

(2) *Collective dynamics of the heavy ion collision: “explosive expansion”*

Lattice calculations are done for heavy quark bound states that are at rest in a hot, static, medium. In heavy ion collisions, however, even if the droplet of hot matter equilibrates rapidly, its temperature drops quickly during the subsequent explosive expansion. The observed quarkonium suppression must therefore result from a suitable time average over a dynamical medium. This is challenging in many ways. One issue that arises is the question of how long a bound state must be immersed in a sufficiently hot heat bath in order to melt. Or, phrased better, how long must the temperature be above the dissociation temperature T_d in order to prevent an heavy quark and antiquark produced at the initial moment of the collision from binding to each other and forming a quarkonium meson?

(3) *Collective dynamics of the heavy ion collision: “hot wind”*

Another issue that faces any data analysis is that quarkonium mesons may be produced moving with significant transverse momentum through the hot medium. In their own reference frame, the putative quarkonium meson sees a hot wind. Phenomenologically, the question arises whether this leads to a stronger suppression since the bound state sees some kind of blue-shifted heat bath (an idea which we will refine in Section 8.7), or whether the bound state is less suppressed since it can escape the heat bath more quickly.

(4) *Formation of quarkonium bound states*

Neither quarkonia nor equilibrated quark–gluon plasma are produced at time zero in a heavy ion collision. Quarkonia have to form, for instance by a colored $c\bar{c}$ pair radiating a gluon to turn into a color-singlet quarkonium state. This formation process is not fully understood in elementary interactions or in heavy ion collisions. However, since the formation process takes time, it is *a priori*

unclear whether any observed quarkonium suppression is due to the effects of the hot QCD matter on a formed quarkonium bound state or on the precursor of such a bound state, which may have different attenuation properties in the hot medium. And, it is unclear whether the suppression is due to processes occurring after the liquid-like strongly coupled plasma reaches approximate local thermal equilibrium or earlier, before equilibration.

(5) *Recombination as a novel mechanism of quarkonium formation*

QCD is flavor neutral and thus charm is produced in $c\bar{c}$ pairs in primary interactions. If the average number of pairs produced per heavy ion collision is $\lesssim 1$, then all charmonium mesons produced in heavy ion collisions must be made from a c and a \bar{c} produced in the same primary interaction. At RHIC and even more so at the LHC, however, more than one $c\bar{c}$ pair is produced per collision, raising the possibility of a new charmonium production mechanism in which a c and a \bar{c} from different primary $c\bar{c}$ pairs meet and combine as the quark-gluon plasma falls apart into hadrons to form a charmonium meson [780]. If this novel quarkonium production mechanism were to become significant, it could reduce the quarkonium suppression or even turn it into quarkonium enhancement. Since there are more c and \bar{c} produced in heavy ion collisions at LHC energies than at RHIC energies, and more at low transverse momentum p_T than at high p_T , and more in central collisions than in peripheral collisions, one seeks signatures of this recombination mechanism in particular in the low- p_T -dominated total J/ψ yields in sufficiently central heavy ion collisions at the LHC [1].

As discussed in Section 8.7 and Chapter 9, calculations based on the AdS/CFT correspondence can provide information relevant for phenomenological modeling, in particular by calculating heavy quark potentials within a moving heat bath and by determining meson dispersion relations. The above discussion illustrates the context in which such information is useful, but it also emphasizes that such information is not sufficient. An understanding of quarkonium production in heavy ion collisions relies on phenomenological modelling as the bridge between experimental observations and the theoretical analysis of the underlying properties of hot QCD matter.