Heavy-ion physics

Korinna Zapp[a]
[a]Lund University, Lund, Sweden

Collisions of heavy ions at collider energies provide us with a unique opportunity to study strongly interacting matter at extreme temperatures and densities in the laboratory. Under these conditions quarks and gluons become deconfined to form a new state of matter, the quark-gluon plasma. Heavy ion physics has seen three major discoveries in the last 30 years. The first is that the QGP is the least dissipative material known and behaves like an almost perfect liquid. The second is that jets which are the manifestations of highly energetic quarks and gluons are strongly suppressed and modified compared to proton–proton collisions. This so-called jet quenching can be understood as the partial equilibration of a far-from-equilibrium system in a thermal QGP. The third main discovery is that particles with low transverse momentum produced in small collision systems like high multiplicity proton–proton and proton–ion collisions show many features that were believed to be signs for QGP formation. On the other hand, no jet quenching has been observed so far in small collision systems. These lectures are meant to give an overview over all relevant aspects of heavy ion physics at a phenomenological level.

1 Introduction

A large number of particles is produced in collisions of heavy atomic nuclei at present day colliders. In Pb+Pb collisions at $\sqrt{s_{\mathrm{NN}}} = 2.76\,\mathrm{TeV}$ in CERN's Large Hadron Collider LHC, for instance, up to 1600 primary charged particles are produced per unit rapidity in the central rapidity region[1]. These particles originate from a system of very high energy density, and following an argument by Bjorken [1] the initial

This chapter should be cited as: Heavy-ion physics, Korinna Zapp, DOI: 10.23730/CYRSP-2025-009.215,
in: Proceedings of the 2023 European School of High-Energy Physics,
CERN Yellow Reports: School Proceedings, CERN-2025-009, DOI: 10.23730/CYRSP-2025-009, p.215.

[1]For comparison, the number of primary charged particles per unit rapidity in typical proton-proton collisions at the same $\sqrt{s}$ is only around 4–5.

energy density shortly after the collision of the nuclei can be estimated from the radius of the nuclei R, an initial proper time $\tau_0 = \mathcal{O}(1\,\text{fm/c})$ and the measured transverse energy at mid-rapidity as

$$\epsilon_0 \simeq \frac{1}{\pi R^2 \tau_0} \left.\frac{\mathrm{d}\, E_\perp}{\mathrm{d}\,\eta}\right|_{\eta=0} \simeq 25\,\text{GeV/fm}^3 \tag{1}$$

for the scenario described above [2]. This corresponds to roughly 25 times the density inside a proton. In a strongly interacting system of such high density there has to be re-scattering in the final state, and scattering drives a system towards thermal equilibrium. *It is one of the central questions in heavy ion physics how and to what extent the system produced in heavy ion collisions thermalises.*

The colliders with heavy ion programs at high beam energies are currently the Relativistic Heavy Ion Collider RHIC at BNL on Long Island and the Large Hadron Collider LHC at CERN in Geneva. RHIC accelerates a variety of different ions to centre-of-mass energies of up to $\sqrt{s_{\text{NN}}} = 200\,\text{GeV}$, i.e. 200 GeV per nucleon pair. The largest data sets are for Au+Au collisions. The heavy ion program at the LHC has so far focused on Pb+Pb collisions at centre-of-mass energies up to $\sqrt{s_{\text{NN}}} = 5\,\text{TeV}$.

2 The quark-gluon plasma

QCD is an *asymptotically free* theory, which means that the coupling decreases as the energy scale increases and the corresponding length scale decreases[2]. It can therefore be expected that when nuclear matter is compressed and/or heated eventually the strong interaction will become so weak that quarks and gluons start propagating as nearly free particles instead of being confined in colour neutral bound states. This state of matter consisting of deconfined quarks and gluons is called the *quark-gluon plasma* (QGP). Another interesting feature of the QGP is that chiral symmetry is restored in the plasma phase.

In a gas of non-interacting particles in thermal equilibrium the pressure is given by the Stefan-Boltzmann law:

$$p(T) = \frac{\pi^2}{90}\left(N_{\text{B}} + \frac{7}{8}N_{\text{F}}\right)T^4\,, \tag{2}$$

where N_{B} and N_{F} are the number of bosons and fermions, respectively. A simple counting of degrees of freedoms leads to $N_{\text{B}}^{(\text{QGP})} = 2[\text{spin}]\cdot(N_c^2-1)[\text{colour}] = 16$ and $N_{\text{F}}^{(\text{QGP})} = 2[\text{anti-/particle}]\cdot 2[\text{spin}]\cdot N_c[\text{colour}]\cdot N_f[\text{flavour}] = 36$ for the QGP, while in a hadron gas with temperatures between the pion and the ρ mass one has $N_{\text{B}}^{(\text{hg})} = 3$ and $N_{\text{F}}^{(\text{hg})} = 0$. One thus expects a strong increase of the pressure at the transition from the hadronic phase to the QGP.

Thermodynamic properties of strongly interacting matter can be obtained from Lattice QCD, i.e. by solving QCD numerically on a discrete space-time lattice. However, this works reliably only for vanishing baryon chemical potential, i.e. for systems with vanishing net baryon number. For this case Lattice QCD indeed finds a rapid increase of the pressure and a cross-over into the plasma phase at a pseudo-critical temperature $T_{\text{c}} = (154 \pm 9)\,\text{MeV} \approx 1.7\cdot 10^{12}\,\text{K}$ [3], which corresponds to roughly 10^5 times the temperature in the centre of the sun. Furthermore, the Lattice results show that the so-called trace anomaly $(\epsilon - 3p)/T^4$, which vanishes for an ideal gas, is still non-zero at $T = 400\,\text{MeV}$ and beyond [3]. This suggests that the QGP is strongly coupled even at temperatures well above the

[2]This is caused by the "anti-screening" of the colour charge due to the gluon self-interaction.

pseudo-critical temperature.

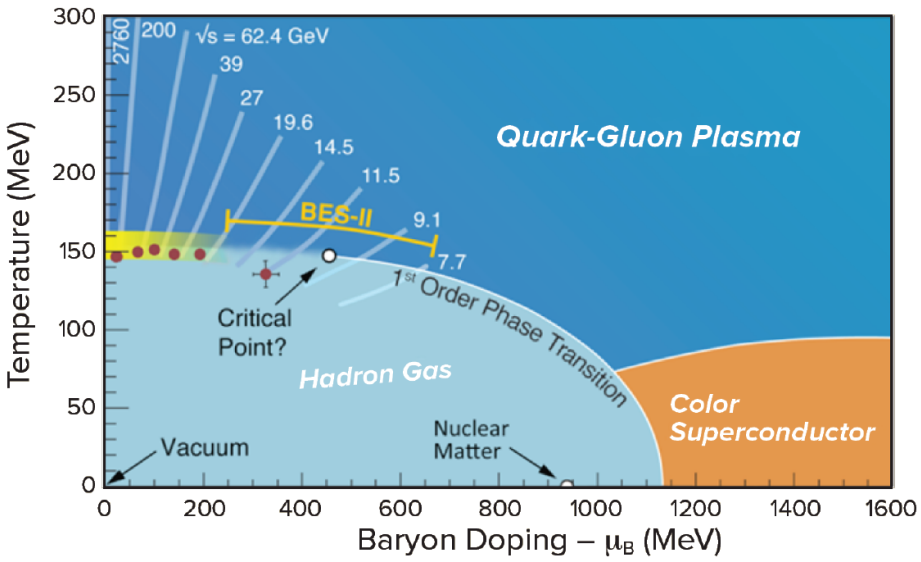

Fig. 1: The phase diagram of QCD with experimental points shown in red, the conjectured first-order phase transition line ending in the critical point and the regions probed by experiments at different beam energies indicated by the light blue lines. The Lattice QCD result for the cross-over temperature at small baryon chemical potential μ_B is shown as the yellow band. Figure from [4].

Figure 1 shows the phase diagram of QCD with baryon chemical potential μ_B on the horizontal and temperature on the vertical axis. Ordinary hadronic matter is found at low temperature T and not too high μ_B and the QGP at high temperature. At high baryon density a colour superconducting phase is expected, but very little is known about this region. The experimentally found points for the phase boundary are shown as the red points. They are extracted from the measured abundances of different hadron species using the *statistical hadronisation model*. The basic assumption of this model is that the system is in thermal equilibrium when it hadronises. The different hadron species are then produced with different probabilities dictated mainly by their mass. It is argued that the composition of the hadronic system does not change afterwards because the density is already too low for number changing scattering processes to occur in the hadronic phase. One can then calculate the expected hadron yields for a grand canonical ensemble and fit the data to extract T and μ_B at the phase boundary. An example for a statistical hadronisation fit is shown in the left plot of figure 2, while the right plot displays the obtained points in the phase diagram. The points at low μ_B are consistent with the Lattice QCD result for the pseudo-critical temperature. These points come from the highest beam energies. The net baryon density in the central rapidity region decreases with beam energy, because the nuclei become more and more transparent. At the LHC the chemical potential extracted from the statistical model is consistent with zero. Another point is remarkable, namely that the statistical hadronisation fits find that strangeness is not suppressed at top RHIC and LHC energies. This indicates that the temperature in the QGP is high enough that strange quarks can be produced thermally. Strangeness enhancement, or rather the disappearance of strangeness suppression, has long been regarded as a sign of QGP formation.

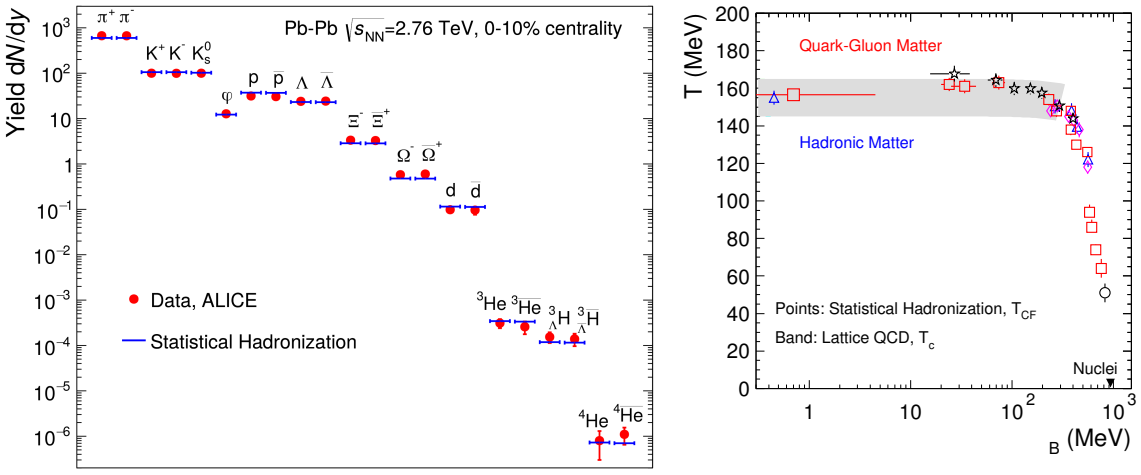

Fig. 2: Left: Hadron yields measured by the ALICE experiment in Pb+Pb collisions at $\sqrt{s_{\mathrm{NN}}} = 2.76\,\mathrm{TeV}$ and statistical hadronisation fit. **Right:** Statistical hadronisation points for different beam energies in the phase diagram. Figures from [9].

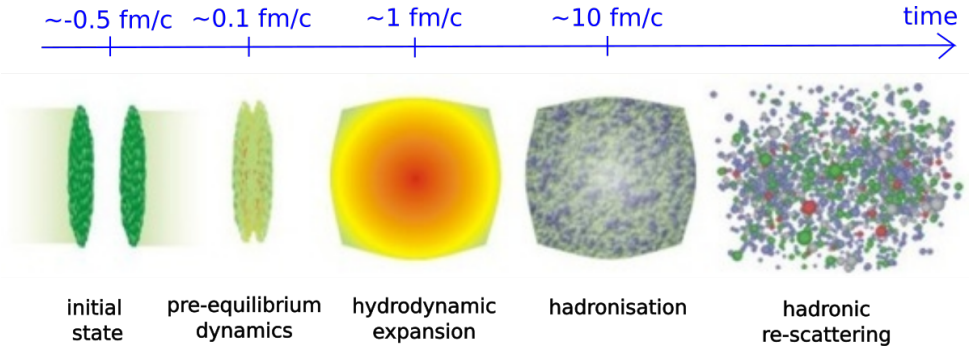

Fig. 3: Different stages of a heavy ion collision at top RHIC and LHC energies, see text for details. Figure courtesy to S. Bass.

Figure 3 summarises the current understanding of heavy ion collisions (see [5–8] for reviews). The evolution of the system produced in a heavy ion collision at high energies encompasses different stages. Starting from the initial state before the actual collision, the system shortly after the collision, at around $0.1\,\mathrm{fm/c}$, enters a phase of pre-equilibrium dynamics that rapidly equilibrates it to a degree that at around $1\,\mathrm{fm/c}$ viscous hydrodynamics becomes applicable. It follows an extended phase of hydrodynamic expansion during which the system cools until it reaches the pseudo-critical temperature and hadronises. Re-scattering continues for a while in the hadronic phase until the density becomes too low and the system reaches the kinetic freeze-out. After that the particles free-stream to the detector and resonances decay.

In the following sections I will briefly go through the different stages of a heavy ion collision. After that, I will discuss hard probes, small collision systems and ultra-peripheral collisions.

3 The initial state

When two nuclei collide the amount of overlap can be quantified in different ways. The *impact parameter* b is the transverse distance between the centres of the nuclei. A head-on collision thus has zero impact parameter. The impact parameter is not an observable quantity and inferring it from data inevitably relies on models. The experiments therefore use *centrality*, which is defined as fraction of the geometric cross section. If the nuclei were hard spheres ("billard balls") the geometric cross section would be given by $\sigma_{\text{geo}} = \pi(R_A + R_B)^2$, where R_A and R_B are the radii of the colliding nuclei. The cross section for a collision with impact parameter up to b is then πb^2 and the centrality is the ratio, i.e. $b^2/(R_A + R_B)^2$. The most central collisions ($b = 0$) thus have, somewhat counter-intuitively, zero centrality.

Experimentally, centrality is determined from the event multiplicity and/or the number of forward-going neutrons measured in the zero degree calorimeters (spectator neutrons) [10]. The measured centrality is interpreted in terms of a *Glauber model* [11] to extract the number of nucleons participating the in the collision, N_{part}, and the number of binary nucleon–nucleon collisions, N_{coll}.

Simple versions of the (optical) Glauber model [12] assume that nucleons travel on straight lines, that the nuclei are large enough that edge effects can be neglected, and that there are no correlations among the nucleons inside a nucleus. The nuclei are represented by a smooth density, usually the Woods-Saxon potential

$$n_A(r) = \frac{n_0}{1 + \exp\left(\frac{r-R}{d}\right)} \,. \tag{3}$$

Integrating over the beam direction one obtains the nuclear thickness function

$$T_A(s) = \int_{-\infty}^{\infty} \mathrm{d}z \, n_A(\sqrt{s^2 + z^2}) \,. \tag{4}$$

The overlap between two nuclei colliding with impact parameter b is then found by integrating the product of the two thickness functions

$$T_{AB}(b) = \int \mathrm{d}^2 s T_A(\mathbf{s}) T_B(\mathbf{s} - \mathbf{b}) \,. \tag{5}$$

The number of binary nucleon–nucleon collisions is then obtained by multiplying the thickness function with the inelastic nucleon–nucleon cross section: $N_{\text{coll}} = T_{AB}(b)\, \sigma_{\text{inel}}^{\text{NN}}$. The geometric cross section is then the cross section for having $N_{\text{coll}} \geq 1$. It is given by

$$\sigma_{\text{geo}} = \int \mathrm{d}^2 b \left[1 - e^{-T_{AB}(b)\, \sigma_{\text{inel}}^{\text{NN}}}\right] \,. \tag{6}$$

The number of participating nucleons can be calculated from the probability for a nucleon to pass through nucleus A at impact parameter b without interaction $\mathcal{P}_0^{(A)}(b) = [1 - \sigma_{\text{inel}}^{\text{NN}} T_A(b)/A]^A$ and is given by

$$N_{\text{part}}(b) = \int \mathrm{d}^2 s \, T_A(\mathbf{s}) \left\{1 - \mathcal{P}_0^{(B)}(\mathbf{s} - \mathbf{b})\right\} + \int \mathrm{d}^2 s \, T_B(\mathbf{s}) \left\{1 - \mathcal{P}_0^{(A)}(\mathbf{s} + \mathbf{b})\right\} \,. \tag{7}$$

As a general rule soft (i.e. low $p_\perp$) particle production scales with the number of participants N_{part} while hard (high $p_\perp$) processes are proportional to the number of binary collisions N_{coll}.

Monte Carlo (MC) versions of the Glauber model are a simple way of dealing with the event-by-event fluctuations in the nuclear density. They distribute nucleons inside the nuclei according to the nuclear potential. When the transverse distance between two nucleons from different nuclei is smaller than $\sqrt{\sigma_{\text{inel}}^{\text{NN}}/\pi}$ they are counted has having an interaction. In this way the determination of N_{part} and N_{coll} is straightforward.

Soft particle production probes the gluon density in the nucleons at small x. The scale evolution of the parton distribution functions leads to a rapid rise of the gluon density at small x when Q^2 increases. The DGLAP equations are linear and generate an ever increasing gluon density, but on physical grounds one expects that eventually the density gets so high that gluon recombination becomes important and competes with the splitting processes producing soft gluons. The generic expectation is that this will slow down the evolution and eventually lead to *gluon saturation*. At this point the gluon density cannot increase further. The typical transverse momentum of saturated gluons defines the saturation scale Q_{s}, which can also be regarded as the (transverse) size of saturated gluons. At top RHIC and LHC energies the saturation scale is found to be of the order of a few GeV. This has lead to the development of the *Colour Glass Condensate* (CGC) [13, 14] framework. This picture is valid at high energies when the valence quarks can be approximated as "frozen" by time dilation. They act as sources for saturated gluons with typical momenta $\mathcal{O}(Q_{\text{s}})$. The gluons have occupation number $1/\alpha_{\text{s}}$, i.e. they form an over-occupied state. When the saturation scale is high enough $\alpha_{\text{s}}(Q_{\text{s}}) \ll 1$, which means that the gluons fields are strong but weakly coupled. They can then be described using classical field theory. The gluon fields obey an evolution equation, the so-called JIMWLK equation. The interaction between two nuclei then leads to the strong colour fields decaying to partons.

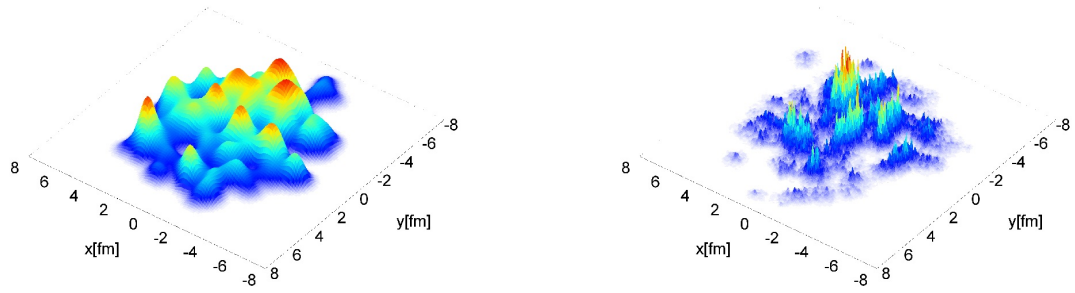

Fig. 4: Energy density in the transverse plane for a Glauber model (left) and the IP-Glasma model encoding gluon saturation in the CGC framework (right), figure from [15].

A characteristic of the CGC is that the typical length scale for fluctuations of the energy density in the initial state is $1/Q_{\text{s}}$. A naive but phenomenological successful alternative approach is to model the energy density by assuming that each participant found in a MC Glauber simulation adds a Gaussian profile of energy density in the transverse plane. The parameters are then found by fitting to data. In case of the calculation shown in figure 4 the width of the Gaussian is $0.4\,\text{fm}$, which is much larger than the $1/Q_{\text{s}}$ of models based on gluon saturation (cf. RHS of figure 4).

4 Pre-equilibrium dynamics

At very early times after the collision of two nuclei the produced system is very far from thermal equilibrium. Despite the rapid longitudinal expansion the system is well described by viscous hydrodynamics at (proper) times of roughly 1 fm/c. The realisation that at that time the system is still very anisotropic has lead to the creation of the term *hydrodynamisation*. Hydrodynamisation denotes the evolution to a point where the system is well described by viscous hydrodynamics (but is not in local thermal equilibrium).

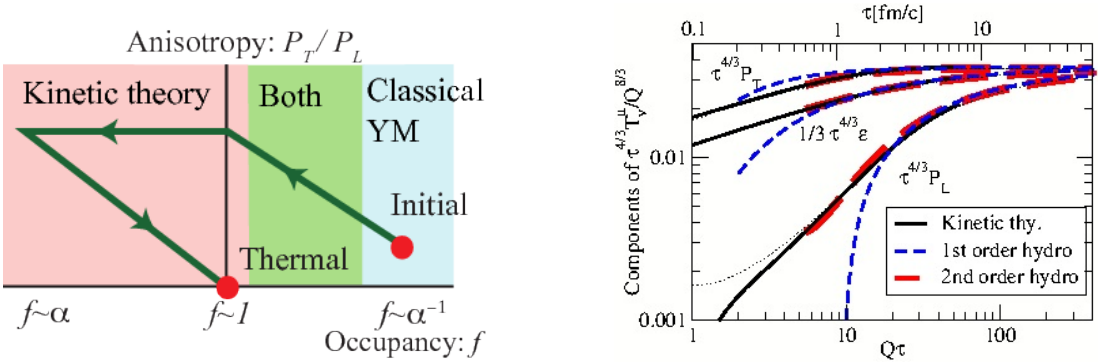

Fig. 5: Left: Cartoon of the evolution from a CGC like initial condition to thermal equilibrium in the presence of strong longitudinal expansion in a weak coupling scenario, figure from [16]. **Right:** Time evolution of rescaled components of the energy-momentum tensor in the effective kinetic theory of QCD at high temperature and hydrodynamics, figure from [17].

The rapid hydrodynamisation can be understood in terms of weak and strong coupling dynamics. In the former scenario the evolution of the system is sketched in figure 5 (left): Starting from a CGC-like initial condition with occupancy $\mathcal{O}(1/\alpha_s)$ the system initially gets more anisotropic and dilute because the interactions are not strong enough to counteract the rapid expansion. Once the systems becomes, under occupied the anisotropy does not increase further. During the last phase the system approaches thermal equilibrium through radiative break-up of the now under occupied modes.

This evolution can be calculated quantitatively using an effective *kinetic theory* of QCD at high temperatures [18]. It formulates a set of Boltzmann equations

$$-(\partial_t + \mathbf{v}\cdot\nabla_x) f(\mathbf{x},\mathbf{p},t) = \mathcal{C}_{\text{“}1\leftrightarrow 2\text{”}}[f] + \mathcal{C}_{2\leftrightarrow 2}[f] \tag{8}$$

for the phase space densities $f(\mathbf{x},\mathbf{p},t)$ of (anti-)quarks and gluons. The dynamics is encoded in the collision kernels on the right hand side, where $\mathcal{C}_{2\leftrightarrow 2}[f]$ describes elastic scattering and $\mathcal{C}_{\text{“}1\leftrightarrow 2\text{”}}[f]$ nearly collinear splitting and merging processes in the presence of multiple coherent scattering. Solving the Boltzmann equations numerically shows that the energy-momentum tensor coincides with the expectation from second order viscous hydrodynamics starting from times $\lesssim 1$ fm/c (figure 5 right).

Strong coupling calculations rely in the *AdS/CFT correspondence* relating strongly coupled conformal field theories to weakly coupled type IIB string theories in a five-dimensional AdS space. Most of the results have been obtained for $\mathcal{N} = 4$ Super-Yang-Mills theory, which shares many similarities with QCD at high temperature. Heavy ion collisions are in this scenario usually modeled as collision of two shock waves [19]. Thermalisation of the produced system is related to the creation of a black hole

in the fifth dimension of the string theory. Also in this case the system quickly reaches hydrodynamic behaviour on timescales $\mathcal{O}(1/T)$.

5 Hydrodynamics

One of the most compelling pieces of evidence for hydrodynamic behaviour is that the anisotropic flow coefficients are well described by hydrodynamics. What is meant by this is the following: the system created in mid-central collisions has an elliptical shape in the transverse plane. In a hydrodynamic scenario the expansion of the system is driven by pressure gradients. In the case of an elliptically shaped system the pressure gradient is larger along the short axis than along the long axis. As a result, matter is pushed out preferentially in the direction of the short axis, which leads to an isotropic momentum distribution. This is quantified by a Fourier decomposition of the momentum distribution in azimuthal angle ϕ

$$\frac{\mathrm{d}\,N}{\mathrm{d}\,\phi} = \frac{N}{2\pi}\left[1 + 2\sum_n v_n \cos(n(\phi - \Psi_n))\right] , \tag{9}$$

where the event plane angle Ψ gives the orientation of the collision system in azimuth. The coefficients v_n are the so-called *flow coefficients*. The elliptical shape of the overlap region induces a large elliptical momentum anisotropy, i.e. a large v_2. Figure 6 shows a measurement of v_2 compared to a hydrodynamic calculation. The data are very well described by the calculation up to transverse momenta of roughly 2 GeV, beyond which other particle production mechanisms become dominant. The different hadron species show a mass ordering that is characteristic for *collective flow*, where all particles flow with a common velocity.

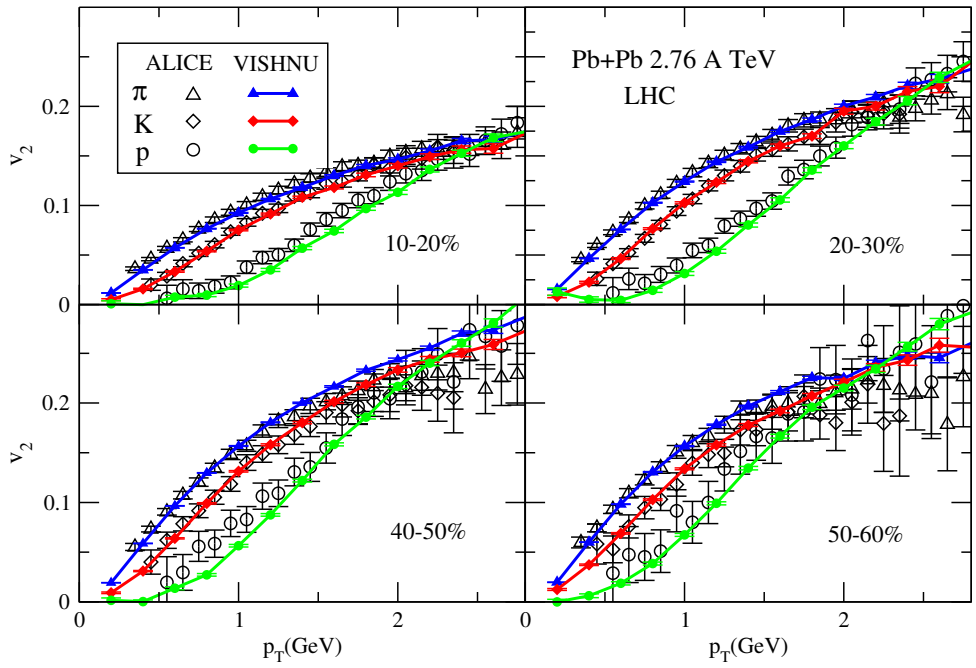

Fig. 6: Measurement of the elliptic flow coefficient v_2 for different hadron species in different collision centralities compared to a hydrodynamic calculation, figure from [20].

Hydrodynamics [21, 22] can be viewed as a low energy effective theory describing long distance, late time behaviour of averaged macroscopic features of a system. It is applicable to a very generic set of theories and assumes that the system is in local thermal equilibrium, i.e. a temperature can be defined locally at each point but may vary from point to point. Microscopic dynamics of the theory enter through the equation of state and a set of coefficients. The equation of state relates the energy density ϵ and the pressure p. Hydrodynamics is valid for distances that are long compared to the mean free path, times that are long compared to the inverse scattering rate and systems with sufficiently smooth variations.

A fluid in thermal equilibrium is described by the energy-momentum tensor

$$T^{\mu\nu} = \epsilon u^\mu u^\nu + p(g^{\mu\nu} + u^\mu u^\nu)\,, \tag{10}$$

where $g^{\mu\nu}$ is the metric and u^μ the fluid velocity. We here allow for small deviations from global thermal equilibrium such that energy density and fluid velocity vary with position and time, $\epsilon = \epsilon(x)$ and $u^\mu = u^\mu(x)$. Energy-momentum conservation $\partial_\mu T^{\mu\nu} = 0$ leads to

$$u^\mu \partial_\mu \epsilon + (\epsilon + p)\partial_\mu u^\mu = 0 \tag{11}$$

$$(\epsilon + p)u^\mu \partial_\mu u^\nu + (g^{\nu\mu} + u^\nu u^\mu)\partial_\mu p = 0\,. \tag{12}$$

Together with the equation of state these form a closed system that can be solved to obtain the energy-momentum tensor.

In order to allow for perturbations with larger gradients a more general form of the energy-momentum tensor is needed

$$T^{\mu\nu} = \epsilon u^\mu u^\nu + p\Delta^{\mu\nu} + \Pi^{\mu\nu}\,, \tag{13}$$

where $\Delta^{\mu\nu} = g^{\mu\nu} + u^\mu u^\nu$ and $\Pi^{\mu\nu}$ is the viscous stress tensor. It can be decomposed into a traceless part and the remainder

$$\Pi^{\mu\nu} = \pi^{\mu\nu} + \pi_{\text{bulk}}\Delta^{\mu\nu}\,, \tag{14}$$

where the shear stress tensor $\pi^{\mu\nu}$ and the bulk viscous pressure π_{bulk} parametrise deviations from ideal fluid dynamics. Viscous fluid dynamics can be organised as a gradient expansion

$$\pi_{\text{bulk}} = -\zeta\partial_\mu u^\mu + \ldots \tag{15}$$

$$\pi^{\mu\nu} = -2\eta\left(\frac{1}{2}\Delta^{\mu\alpha}\Delta^{\nu\beta} + \frac{1}{2}\Delta^{\mu\beta}\Delta^{\nu\alpha} + \frac{1}{3}\Delta^{\mu\nu}\Delta^{\alpha\beta}\right)\partial_\alpha u_\beta + \ldots \tag{16}$$

where at first order the bulk viscosity $\zeta = \zeta(\epsilon)$ and the shear viscosity $\eta = \eta(\epsilon)$ appear. The second order comes with many more parameters (relaxation times, more transport coefficients, ...). The resulting increasingly complicated evolution equations have to be solved numerically except for the simplest cases.

The shear viscosity is related to momentum transport in the fluid. When the shear viscosity is large compared to the entropy density (large η/s) momentum perturbations are transported over large distances by *quasi-particles*, while at small η/s there are no well-defined quasi-particles. The shear viscosity thus also determines to what extend perturbations in the initial conditions are washed out. This can be seen in figure 7, which shows how an initial energy density with fluctuations looks after evolving

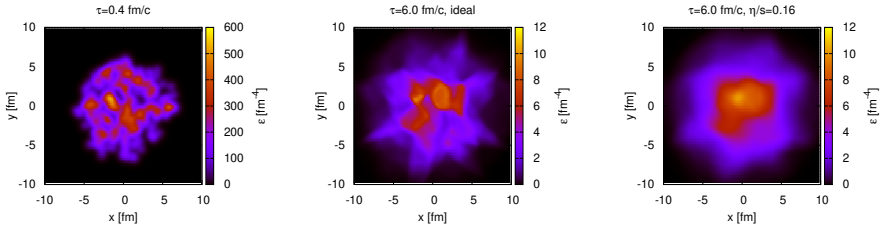

Fig. 7: Initial energy density distribution at proper time $\tau = 0.4\,\text{fm/c}$ in the transverse plane (left) and after evolving to time $\tau = 0.6\,\text{fm/c}$ with ideal hydrodynamics (middle) and with viscous hydrodynamics with a small shear viscosity $\eta/s = 0.16$, figure from [23].

it for a short time with and without a small shear viscosity. With shear viscosity the fluctuations have been smeared out to a larger extent.

The efficiency with which elliptic flow is generated from a given spatial anisotropy also depends on the shear viscosity — the smaller the shear viscosity the larger v_2 [24]. The values of η/s that have been inferred from measurements of the flow coefficients are very close to $1/(4\pi)$, which is a conjectured lower bound obtained from AdS/CFT calculations [25]. The QGP is thus the least dissipative material currently known.

Given the shape of the overlap region of two colliding nuclei the elliptic flow coefficient v_2 is dominant, but higher coefficients are also present. While a smooth spatial distribution with a mirror symmetry does not generate odd harmonics ($n = 3, 5, \ldots$) event-by-event fluctuations in the initial state give rise to these.

The measurements and the hydrodynamic modeling have become so precise that they are sensitive to small deformations of the initial state, e.g. due to deformations of the colliding nuclei. This is for instance the case in data from Ru+Ru and Zr+Zr collisions at RHIC [27].

6 Hadronisation

Hydrodynamics describes the evolution of a fluid, but in the detector particles are measured. The usual way of translating the fluid's energy-momentum tensor into a particle spectrum is the *Cooper-Frye prescription* [28]. The starting point is the observation that, neglecting viscous corrections, the occupation number in each fluid element is given by a Fermi-Dirac or Bose-Einstein distribution

$$\frac{\mathrm{d}N_i}{\mathrm{d}^3x\,\mathrm{d}^3p} = f_i(p^\mu; T(x), u^\mu(x)) \propto \left[e^{\frac{u_\mu(x)p^\mu}{T(x)}} \pm 1 \right]^{-1} . \tag{17}$$

To find the particle spectrum one considers the number of particles passing through a surface Σ, which can be written in terms of the particle number current density $j_i^\mu(x)$ for particle species i

$$N_i = \int\limits_\Sigma \mathrm{d}^3\sigma_\mu\, j_i^\mu(x) = \int\limits_\Sigma \mathrm{d}^3\sigma_\mu \left[\int \frac{\mathrm{d}^3p}{(2\pi)^3} \frac{p^\mu}{E} f_i(x,p) \right] , \tag{18}$$

where $\mathrm{d}\sigma_\mu$ is the area element of the surface Σ. The particle spectrum is thus found to be given by

$$E\frac{\mathrm{d}N_i}{\mathrm{d}^3p} = \frac{1}{(2\pi)^3}\int\limits_\Sigma p_\mu \mathrm{d}\sigma^\mu\, f_i(x,p)\,. \tag{19}$$

The momentum spectrum of particles obtained from two different surfaces is the same only if the particles are free-streaming between the two surfaces. The obvious choice for Σ is thus the freeze-out surface, i.e. the surface of last scattering. In practice it is often assumed that the temperature is constant on the freeze-out surface.

When the partonic distribution function is known the distribution of hadrons can be obtained from the *quark coalescence* picture [29, 30]. The idea of quark coalescence is simply put that quarks and anti-quarks combine to form hadrons. The number of mesons, for instance, is then again given by an integral over a space-like surface Σ

$$N_\mathrm{M} = g_\mathrm{M}\int_\Sigma (p_{1\mu}\mathrm{d}\sigma_1^\mu)\,(p_{2\mu}\mathrm{d}\sigma_2^\mu)\,\frac{\mathrm{d}^3p_1}{(2\pi)^3E_1}\,\frac{\mathrm{d}^3p_2}{(2\pi)^3E_2} f_q(x_1,p_1) f_{\bar{q}}(x_2,p_2) f_\mathrm{M}(x_1,x_2,p_1,p_2)\,, \tag{20}$$

where g_M is a statistical factor and f_M is the probability for a quark and anti-quark to form a meson, e.g.

$$f_\mathrm{M}(x_1,x_2,p_1,p_2) \propto \exp\left(\frac{(x_1-x_2)^2}{2\Delta_x}\right)\cdot\exp\left(\frac{(p_1-p_2)^2}{2\Delta_p}\right)\,. \tag{21}$$

For baryons corresponding expressions can be written down.

The coalescence picture predicts that v_2 of hadrons scales with the number of constituent quarks, a feature that is clearly visible in RHIC data (and to a lesser extent in LHC data).

7 Hadronic re-scattering

After the chemical freeze-out where quarks and gluons are converted into hadrons re-scattering continues for a while in the hadronic phase. There are two ways of dealing with this. The first is to run hydrodynamics down to the kinetic freeze-out, which requires additional input to adequately describe the hadronic phase. The second option is to explicitly simulate re-scattering in the hadronic phase with transport codes based on the Boltzmann equation. This requires knowledge of a large number of hadronic resonances and cross sections. Examples for hadronic transport codes are UrQMD [31] and SMASH [32].

8 Hard probes

Particles with very high transverse momentum are produced in hard partonic scattering processes characterised by a large momentum transfer Q. Due to the uncertainty principle these processes occur on a timescale $\mathcal{O}(1/Q)$ and are therefore the first processes to happen in a collision of two nuclei. The produced high-$p_\perp$ particles then travel through the QGP on their way to the detector. The time and length scale of hard scattering processes are too short to feel the nuclear environment. This is basically the same argument as underlying factorisation theorems that allow to write the cross section for a hard process as a convolution of a partonic cross section $\hat{\sigma}_{ij}$ for producing the hard particle in a scattering of partons i

and j with the parton distribution functions $f_i(x, Q^2)$

$$\sigma(P_1, P_2) = \sum_{i,j} \int_0^1 \mathrm{d}x_1 \, \mathrm{d}x_2 \, f_i(x_1, Q^2) f_j(x_2, Q^2) \hat{\sigma}_{ij}(x_1 P_1, x_2 P_2, \alpha_\mathrm{s}, Q^2) \,. \tag{22}$$

Here the partonic cross section $\hat{\sigma}_{ij}$ encodes the short distance physics and is insensitive to the nature of the colliding hadrons. It has an expansion in powers of the strong coupling α_s and can be calculated order by order in perturbation theory. The parton distribution function (PDF) $f_i(x, Q^2)$ is the density of partons of type i carrying a fraction x of the hadron's longitudinal momentum inside the beam hadron when probed at a scale Q^2. The PDFs of protons and neutron bound inside nuclei are somewhat different from the free proton PDF and there are *nuclear PDF* fits available.

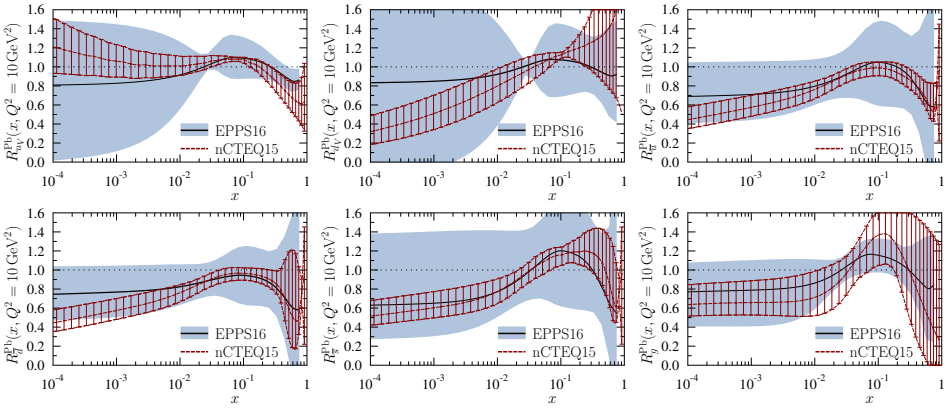

Fig. 8: Nuclear modifications of the PDFs for Pb nuclei at a scale of $Q^2 = 10\,\mathrm{GeV}^2$ obtained from the EPPS16 and nCTEQ15 nuclear PDF sets, figure from [33].

The approach taken by the EPPS collaboration [33] is to define the bound proton PDF as

$$f_i^{\mathrm{p/A}}(x, Q^2) = R_i^A(x, Q^2) f_i^\mathrm{p}(x, Q^2) \,, \tag{23}$$

where f_i^p is a free proton PDF and R_i^A encodes the nuclear modification. The free proton PDF used for the EPPS16 nuclear PDF set is CT14NLO. R_i^A is parametrised and fitted to data in a procedure that is similar to the way in which free proton PDFs are constructed. The reason for not directly fitting the entire nuclear PDF (instead of the nuclear modification of the free proton PDF) is that the amount of data available for nuclear PDF fits is much smaller than for free proton PDFs. The neutron PDF is obtained from isospin symmetry. Figure 8 shows a comparison of the nuclear modification of the PDF from the EPPS16 [33] and nCTEQ15 [34] nuclear PDF sets. Other nuclear PDF sets, e.g. DSSZ [35], are also available but not shown. Generally, the uncertainties are sizable due to sparse data from nuclear collisions. The nuclear modifications are typically of moderate size and become smaller at higher scales. From a theoretical standpoint the production of high-$p_\perp$ particles in hard scattering processes is thus under control.

The production of the heavy charm and beauty quarks requires a high momentum scale due to

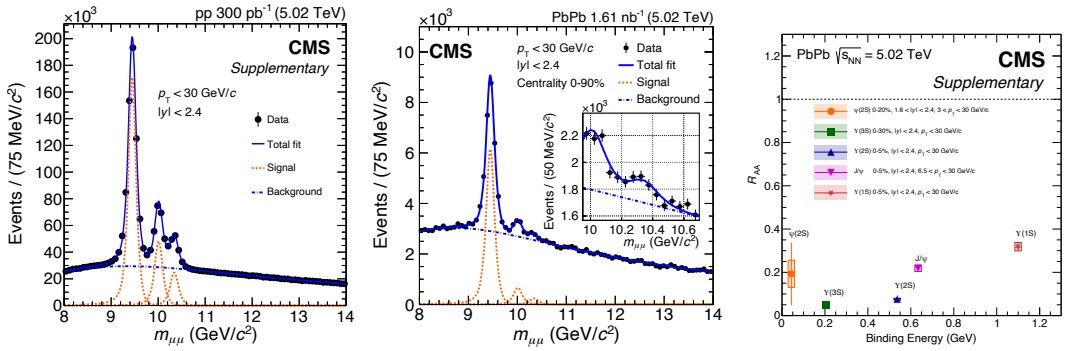

Fig. 9: Left and middle: Invariant mass distribution of muon pairs showing the Υ(1S), Υ(2S) and Υ(3S) peaks in p+p (left) and Pb+Pb (middle) collisions as measured by CMS. **Right:** Quarkonium yields per binary collision in Pb+Pb collision divided by the respective p+p yields showing the suppression of quarkonia states in nuclear collisions, figures from [37].

Table 1: Masses, binding energies and radii of the lightest charmonium and bottomonium states.

state	J/Ψ	Ψ'	Υ	Υ'	Υ''
mass [GeV]	3.10	3.68	9.46	10.02	10.36
ΔE [GeV]	0.64	0.05	1.10	0.54	0.20
r [fm]	0.25	0.45	0.14	0.28	0.39

the large masses of the produced particles. An effect that has long been regarded as a signature for QGP formation is the suppression of quarkonium states due to *colour screening* [36]. The idea is that when the quarkonium states are placed in a deconfined medium consisting of quasi-free colour charges the charges of the heavy quarks are screened. When the screening length is smaller than the size of the quarkonium state the quarkonium is expected to dissolve. Different quarkonium states have different radii and should thus disappear at different QGP temperatures. Observing this so-called sequential suppression should thus in principle allow to measure the QGP temperature. Table 1 lists the binding energies and radii of the most easily accessible quarkonium states. As seen in Figure 9 there is indeed an indication of a binding energy dependent suppression of quarkonium states in heavy ion collisions.

There is, however, a competing effect and that is the *regeneration* of quarkonia by statistical hadronisation at the phase boundary [38]. At the LHC the charm cross section is so large that typically more than one $c\bar{c}$ pair is produced per Pb+Pb collision. When the number of charm quarks in the system is large enough there is a sizable probability that a c and a $\bar{c}$ quark will end up close in phase space and form a charmonium state during hadronisation. This effect can even lead to an enhancement of charmonium states.

Hard processes are expected to scale with the number of binary collisions, therefore a common way of quantifying modifications due to a nuclear environment is to take the ratio to the corresponding

quantity in p+p collisions scaled by the number of binary collisions

$$R_{AA}(\{\cdot\}) = \frac{\left.\frac{\mathrm{d}\sigma_X}{\mathrm{d}\{\cdot\}}\right|_{AA}}{\langle N_{\mathrm{coll}}\rangle \left.\frac{\mathrm{d}\sigma_X}{\mathrm{d}\{\cdot\}}\right|_{pp}} \,. \tag{24}$$

Quantities of this type are called *nuclear modification factors*. In the case of electroweak processes the different isospin composition in p+p and heavy ion collisions also has to be taken into account. Deviations from unity in nuclear modification factors are signs that a nucleus–nucleus collision is more than the incoherent superposition of nucleon–nucleon collisions and that there are additional effects from the nuclear environment.

The electroweak bosons W, Z and γ are produced in hard scattering processes but don't participate in the strong interaction and thus escape from the QGP without interaction. Measurements of nuclear modification factors for these bosons are indeed consistent with the expectation from nuclear PDF and isospin effects [39, 40]. This is an important confirmation of scaling with N_{coll} and a cross-check that nuclear PDF effects are under control.

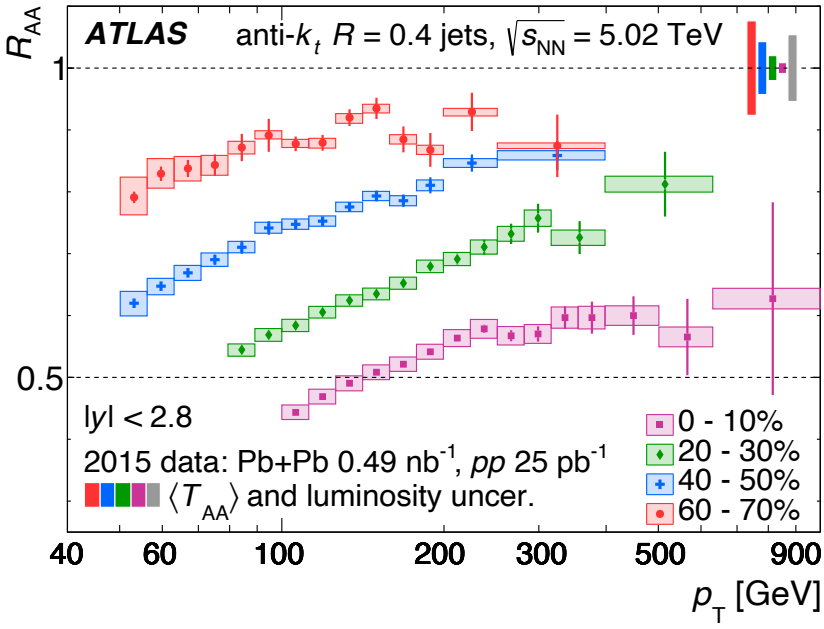

Fig. 10: Nuclear modification factor for jets as a function of the jet $p_\perp$ in Pb+Pb collisions of different centrality as measured by ATLAS, figure from [41].

The situation is very different for jets, which are the result of the fragmentation of highly energetic quarks and gluons. Figure 10 shows that jets are suppressed by almost a factor of two out to very high $p_\perp$ in central Pb+Pb collisions at the LHC. This phenomenon (together with modifications of the internal structure of jets) is referred to as *jet quenching*. The energetic quarks and gluons interact in the QGP and thereby loose energy. Jet quenching can thus be seen as the partial thermalisation of a far-from-equilibrium system and can inform us about equilibration in QCD.

Another observation is that the transverse momentum imbalance in di-jets increases, because one of the jets looses more energy than the other [42]. Interestingly, the azimuthal angle between the two jets remains unchanged. Looking at the energy distribution inside jets it is found that the jet profile shows a suppression at intermediate distances from the jet axis and an enhancement at the periphery of the jet, i.e. at large distances from the jet axis [43]. The jet fragmentation functions show a consistent enhancement of soft fragments and a suppression at intermediate momenta [44]. One can thus conclude that hard structures inside jets survive largely unmodified while there is a clear enhancement of soft activity at large angles relative to the jet axis. This indicates that hard partonic systems loose energy predominantly through radiation of soft gluons while rare hard or semi-hard emissions don't play a major role [45].

There are two scenarios for how hard partons interact in the QGP, namely the strong coupling and the weak coupling scenario. In the former the hard parton does not resolve quasi-particles in the QGP and one has to employ AdS/CFT techniques to calculate the energy loss of hard partons at strong coupling, which is caused by a kind of drag force. The advantage of this approach is that an exact solution can be obtained with no uncertainties in the parton-medium interaction. A fundamental problem is that jets are a weak coupling phenomenon and don't have a natural counterpart at strong coupling. A $q\bar{q}$ pair in a QGP is dual to a classical string falling into a black hole in the fifth dimension of AdS space. This construction has been used as a proxy for a "holographic jet" [46, 47].

In the weak coupling scenario hard partons scatter off quasi-particles in the QGP. It should be noted that the relevant scale for distinguishing between the two scenarios is the momentum transfer between the hard parton and the QGP (not the momentum of the hard parton). Perturbative techniques can be used to calculate the energy loss of a hard parton in QCD at week coupling. Two types of processes can occur, namely elastic scattering and bremsstrahlung. Thermalisation through elastic scattering is slow and the energy loss is dominated by QCD bremsstrahlung. In bremsstrahlung due to multiple scattering a quantum mechanical interference occurs because scatterings within the formation time of a gluon emission act coherently. This gives rise to the QCD analogue of the Landau-Pomerantchuk-Migdal (LPM) effect.

The concept of *formation time*, i.e. the time it takes to radiate a gluon, plays a central role in the discussion of energy loss of hard partons. The emission of a gluon at finite emission angle is kinematically only possible when the emitting parton is off-shell with a virtual mass $m_{\text{virt}} = p^2$. The time it takes the virtual parton with energy E to "decay" by emitting a gluon with energy ω and relative transverse momentum $k_\perp$ is the gluon formation time and can be estimated from the uncertainty principle to be $\mathcal{O}(1/m_{\text{virt}})$ in the virtual parton's rest frame. In a general frame the formation time acquires an additional boost factor E/m_{virt} and becomes

$$t_{\text{form}} = \frac{1}{m_{\text{virt}}}\frac{E}{m_{\text{virt}}} = \frac{E}{2p_\mu k^\mu} \simeq \frac{E}{\omega E\theta^2} \simeq \frac{\omega}{k_\perp^2}\,, \tag{25}$$

where p_μ and k_μ are the parton and gluon momentum, respectively, and θ is the emission angle.

To illustrate the LPM effect [48], let's consider a highly energetic quark with energy E traveling through a coloured background medium and emitting an almost collinear gluon with energy $\omega \ll E$. Due to multiple scattering in the background the gluon undergoes Brownian motion and acquires a transverse

momentum relative to the quark $\langle k_\perp^2 \rangle = \hat{q}L$, where L is the path length traveled and $\hat{q}$ is the transport parameter. Then the gluon formation time becomes

$$t_\mathrm{f} \simeq \frac{\omega}{k_\perp^2} \simeq \frac{\omega}{\hat{q}t_\mathrm{f}} \quad \Rightarrow \quad t_\mathrm{f} = \sqrt{\frac{\omega}{\hat{q}}}\,. \tag{26}$$

The number of scatterings during the formation time is given by $N_\mathrm{coh} = t_\mathrm{f}/\lambda$, where λ is the mean free path. The incoherent gluon spectrum is parametrically $\mathrm{d}^2I^\mathrm{incoh}/\mathrm{d}\omega\mathrm{d}y \propto \alpha_\mathrm{s}/(\omega\lambda)$, where y is the longitudinal spatial coordinate. To estimate the gluon energy spectrum including coherence effects one has to divide by the number of scatterings acting coherently, i.e. effectively as one

$$\frac{\mathrm{d}^2I^\mathrm{coh}}{\mathrm{d}\omega\mathrm{d}y} \simeq \frac{1}{N_\mathrm{coh}}\frac{\mathrm{d}^2I^\mathrm{incoh}}{\mathrm{d}\omega\mathrm{d}y} \propto \frac{\alpha_\mathrm{s}}{\omega\lambda}\lambda\sqrt{\frac{\hat{q}}{\omega}} = \alpha_\mathrm{s}\sqrt{\hat{q}}\,\omega^{-3/2}\,. \tag{27}$$

To obtain the total energy loss this spectrum has to be weighted with gluon energy and integrated over the gluon energy and the path traveled in the medium

$$\Delta E = \int_0^L \mathrm{d}y \int_0^{\omega_\mathrm{c}} \mathrm{d}\omega\,\omega \frac{\mathrm{d}^2I}{\mathrm{d}\omega\mathrm{d}y} \propto \alpha_\mathrm{s}\hat{q}L^2\,, \tag{28}$$

where $\omega_\mathrm{c} = \hat{q}L^2$ is the highest gluon energy that can be radiated coherently, i.e. the gluon energy for which the formation time equals the path length L in the medium and the entire medium acts coherently. The L^2 dependence of the energy loss is characteristic of the LPM effect.

A high-$p_\perp$ parton produced in a hard scattering process undergoes a scale evolution in QCD during which it emits gluons[3]. Inserting numbers for typical jets at the LHC into equation (25) one finds that the first few emission happen very quickly, usually before the QGP forms. However, the time needed for the total scale evolution through successive gluon emissions is of the order of a few fm/c and thus comparable to the transverse size of the QGP. Interactions in the QGP and gluon emission related to the QCD scale evolution thus happen at the same time.

When a quark emits a gluon the resulting quark and gluon together still form a colour triplet. Whether the quark–gluon system is resolved as a quark carrying a triplet charge and a gluon carrying an octet charge, or is seen as a quark with a triplet charge depends on the resolution scale. For interactions in the QGP this means that a scattering with a transverse momentum transfer larger than the inverse transverse separation of the quark and the gluon will resolve two partons. A soft scattering with transverse momentum transfer smaller than the inverse transverse separation, on the other hand, will only 'see' a quark [49]. This is a type of *colour coherence* and can naturally explain why hard small angle structures inside jets are not affected by jet quenching. The hope that by measuring up to which opening angle structures stay coherent to get a handle on the medium resolution power has so far not been fulfilled.

Many jet quenching measurements and particularly those targeted at the internal structure of quenched jets suffer from the so-called *selection bias*. The jet $p_\perp$ spectrum is very steeply falling and

[3]This is basically the same as the DGLAP evolution of the PDFs and is simulated by parton showers in Monte Carlo event generators or taken care of by analytical resummation.

the energy loss has large fluctuations. As a consequence, in a sample of jets selected based on the final $p_\perp$ those that lost only very little energy will always dominate. The nuclear modification factor is thus mostly sensitive to the no-quenching probability. In jet shape and jet sub-structure observables a related effect becomes visible. As mentioned earlier, the first few splittings of a hard parton happen very early and can thus be expected to be unmodified by the QGP. But they are decisive in determining the general shape of the jet. Already in p+p collisions jets with the same $p_\perp$ are not all the same: some have a soft fragmentation pattern where the energy is shared among many particles already at parton level, while some have a hard fragmentation pattern with only few energetic partons. The soft fragmenting jets are more susceptible to medium modifications and typically loose significantly more energy than the hard fragmenting ones. In the presence of energy loss the soft fragmenting jets are thus more likely to fall below the minimum $p_\perp$ required and thereby disappear from the sample. A sample of jets in heavy ion collisions is thus biased towards a harder fragmentation pattern with a harder and narrower core than in p+p collisions. This is indeed observed in data. This bias is hard to control on a quantitative level and complicates the interpretation of the internal structure of quenched jets.

The QGP modifies hard partons, and hard partons modify the QGP. The energy and momentum lost by hard partons is transferred into the QGP and manifests itself in the form of additional soft particle production at large angles relative to the jet axis [50]. The measurements go out to $\Delta r = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2} = 1$, but most theoretical calculations are available only for smaller $\Delta r < 0.3$. In that regime the enhancement of soft particles is in some calculations clearly caused by the so-called *medium response* [51, 52], in others it is due to other effects [53] or it is a mixture [54]. The situation is thus so far inconclusive.

9 Small collision systems

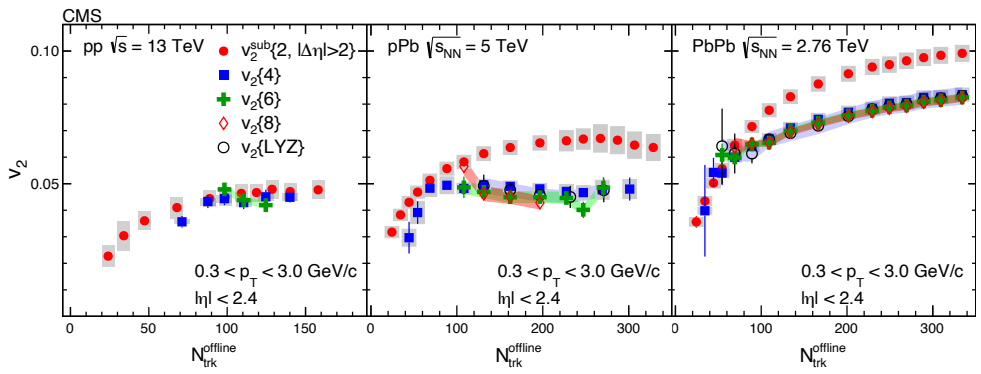

Fig. 11: Elliptic flow coefficient v_2 measured using different methods in p+p, p+Pb and Pb+Pb collisions, figure from [55].

The traditional heavy ion physicists' view was that many correlations observed among soft particles produced in heavy ion collisions are signs of collective flow in the final state and thus indicative of QGP formation. In p+p collisions, on the other hand, the density of produced particles was thought to be too low for sizable final state re-scattering. The expectation was therefore that the particles produced in p+p and other small collision systems like p+Pb free-stream to the detector. It therefore came as a

surprise when the LHC experiments discovered that many of the characteristics of heavy ion collisions are also found in small collision systems. An example is v_2, which is shown in figure 11 for p+p, p+Pb and Pb+Pb collisions. Although the magnitude of v_2 increases with system size the overall behaviour is similar in all three systems. Another example is strangeness enhancement, which shows a smooth increase with multiplicity from p+p via p+Pb to Pb+Pb collisions [56]. Interestingly, no jet quenching has so far been observed in small collision systems.

Different explanations for the observed correlations in small collision systems have been put forward. The first is that it is due to response to the initial geometry. The mechanism that creates the azimuthal anisotropy could then be hydrodynamic flow as in nucleus–nucleus collisions [57], or it could be the so-called *escape mechanism* [58, 61] which requires only $\mathcal{O}(1)$ scattering per particle. Alternatively, correlations present in the initial state could get imprinted on the final state without final state interactions. The CGC framework, for instance, predicts such initial state correlations [59]. A third explanation is inspired by p+p physics and proposes string interactions, in this case string shoving, as an explanation for the observed correlations [60]. Presently, no preference for one of these scenarios can be identified.

Despite the observed similarities there are also fundamental differences between small and large collision systems. One is that in small systems multiplicity is generated to a larger extent by jets than in large systems. As a consequence there is no strong correlation between centrality and multiplicity as in heavy ion collisions [62]. Small systems are also affected by auto-correlations of different kinds that can complicate making meaningful measurements. One example for this is that the charged and neutral kaon multiplicities have a different dependence on the charged particle multiplicity. The reason is simply that by requiring a large charged particle multiplicity one biases the events towards having more charged than neutral kaons. Similarly, a rapidity shift of a central di-jet system depending on the amount of transverse energy in the forward region [63] could be explained by energy conservation [64]. These complications are part of the reason why no consistent understanding of small collision systems has emerged so far.

10 Ultra-peripheral collisions

Ultra-relativistic ions are sources of very strong electromagnetic fields. Electric fields can go up to $10^{16} - 10^{18}$ V/m and magnetic fields reach $10^{14} - 10^{16}$ T. These are the strongest electromagnetic fields in the universe, but they are extremely short lived. In ultra-peripheral collisions where the impact parameter is too large for strong interactions to occur a new class of processes with incoming photons becomes accessible. These processes can be roughly divided into photon–photon scattering where two photons, one from each nucleus, interact and photo-nuclear reactions, where a photon from one nucleus interacts with the other nucleus. An example for the first class is the first evidence for light-by-light scattering reported by ATLAS (figure 12). The scattering $\gamma + \gamma \rightarrow \gamma + \gamma$ proceeds in the Standard Model via a lepton loop as shown in figure 12. In certain BSM extensions, however, it can proceed via an axion-like particle in the s-channel and the measurements can be used to set limits on the axion-like particles [66].

Both CMS and ATLAS have measured the anomalous magnetic moment a_τ of the τ in $\gamma + \gamma \rightarrow \tau + \bar{\tau}$ events [67, 68]. Due to the large τ mass a_τ is 280 times more sensitive to BSM physics than

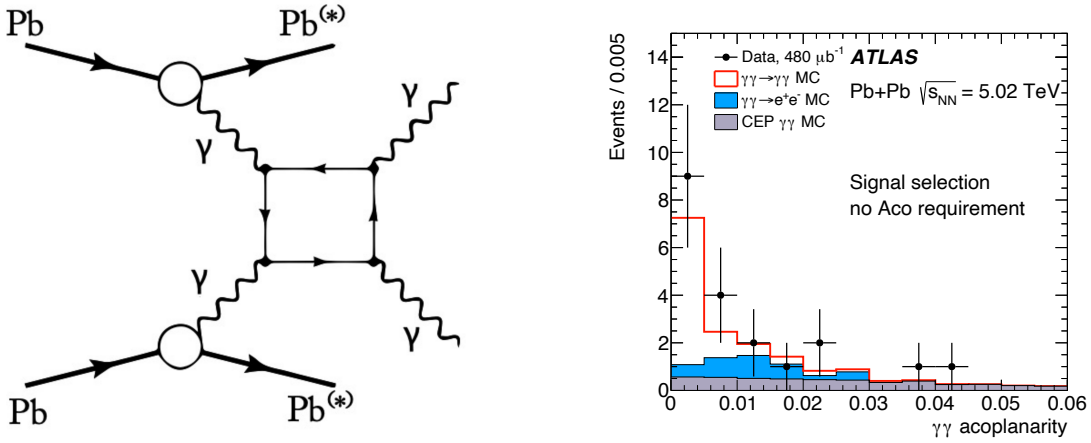

Fig. 12: First evidence for light-by-light scattering in ultra-peripheral Pb+Pb collisions by ATLAS, figure from [65].

the anomalous magnetic moment of the muon. With better statistics expected in the near future the measurements will reach similar precision as the currently best determinations of a_τ.

A similar process, namely the Breit-Wheeler process $\gamma + \gamma \to e^+ + e^-$ can be used to measure the photon spin via the angular asymmetry of the produced leptons [69]. These measurements can therefore be used to place limits on dark photons [70].

An example for an interesting photo-nuclear reaction is diffractive J/Ψ production. By measuring the coherent (where the photon interacts with the whole nucleus) and incoherent (where the photon resolves the nuclear structure) production modes one is sensitive to the nuclear and sub-nucleonic structure, respectively [71, 72].

These are just a few examples for the kind of physics that ultra-peripheral collisions give access to. With larger data sets becoming available this area is certainly going to expand.

11 Conclusions

Heavy ion collisions offer a unique opportunity to study strongly interacting matter under extreme conditions. The temperature and density reached in these collisions are so high that QCD becomes weakly coupled enough that quarks and gluons become deconfined and propagate as quasi-free particles. This new phase of strongly interacting matter is called the quark-gluon plasma. Most of the research in heavy ion physics at high energies revolves around the questions how the QGP comes into existence and what its properties are.

Heavy ion collisions proceed through a series of distinct phases. Important characteristics of the collision are defined in the initial state prior to the actual collision. These include the nuclear geometry, which is usually modeled with Glauber models. Soft particles production probes the gluon density at small x where gluon saturation is expected to set in. This has lead to the development of the Colour Glass Condensate as a framework for modeling heavy ion collisions in terms of saturated gluon fields. Directly after the collision of the nuclei the produced system is far from thermal equilibrium and rapidly expanding in the longitudinal direction. The early times dynamics can be described at weak coupling

using kinetic theory and at strong coupling using the AdS/CFT correspondence. Both scenarios lead to equilibration of the produced matter to the extent that viscous hydrodynamics becomes applicable on the phenomenologically required (proper) time scale $\tau \lesssim 1\,\mathrm{fm/c}$. What follows is an extended phase of hydrodynamic expansion during which the system develops collective flow that is reflected in the momentum distributions of final state hadrons. The QGP has a very low shear viscosity to entropy density ratio $\eta/s \gtrsim 1/4\pi$ closed to the conjectured lower bound. The QGP is thus the least dissipative system presently known: it behaves like an almost perfect liquid. The small shear viscosity is indicative of strong residual interactions among the partons in the QGP. Finally the system has cooled so much that it transitions into the hadronic phase at the pseudo-critical temperature $T_\mathrm{c} \simeq 155\,\mathrm{MeV}$. At this point the strong coupling is so large again that there is little first principles understanding of the hadronisation process and one has to rely on phenomenological models. Hydrodynamic calculation rely in the Cooper-Frye prescription to convert the energy-momentum tensor of the fluid into hadron spectra. One idea that turned out to be useful in the context of heavy ion collisions is quark coalescence that models how quarks and anti-quarks combine to from hadrons. Scattering continues for a while in the hadronic phase and can be modeled with hydrodynamics or transport theory.

The production of jets, which are the manifestations of highly energetic partons, is found to be suppressed in heavy ion collisions — a phenomenon that is called jet quenching. The partons fragmenting into jets are produced very early in the collision in hard scattering processes that are well understood theoretically. They then propagate through the QGP and interact strongly with it. This leads to energy loss of the hard partons in what can be viewed as the partial thermalisation of the hard partons. The lost energy is transferred to the QGP and manifests itself in the form of additional soft particles situated at large angles from the jet axis. It is an open question whether the hard partons resolve quasi-particles in the QGP, i.e. whether their interaction with the QGP is best characterised by weak or strong coupling dynamics. The interpretation of jet quenching measurements is hindered by the selection bias, which biases samples of jets in heavy ion collisions towards the least modified jets and a hard fragmentation pattern.

Quarkonium states are also suppressed by colour screening in the QGP, but a quantitative understanding of the data is complicated by a number of confounding factors, most notably the regeneration of charmonium states by statistical hadronisation.

Soft particles produced in small collision systems such as high multiplicity p+p and p+A collisions show surprisingly many features such as correlations and strangeness enhancement that were thought to be signs for QGP formation. On the other hand, no jet quenching has so far been observed in small collision systems. It is not clear whether the effects come from correlations in the initial state that get imprinted on the final state or from final state interactions among particles or strings.

Ultra-peripheral collisions come with extremely strong electromagnetic fields and allow to study a large variety of processes with incoming photons. One example is the first evidence for light-by-light scattering.

Heavy ion physics is a field with a rich and diverse phenomenology that requires conceptually new approaches both theoretically and experimentally. The theoretical tools currently used include classical field theory, thermal field theory, AdS/CFT techniques, kinetic theory, relativistic hydrodynamics, effec-

tive field theories, lattice QCD and phenomenological models. This list is not complete and is likely to get longer as progress is made in understanding the many aspects of heavy ion collisions.

References

[1] J. D. Bjorken, Phys. Rev. D **27** (1983), 140-151 doi:10.1103/PhysRevD.27.140

[2] J. Adam *et al.* [ALICE], Phys. Rev. C **94** (2016) no.3, 034903 doi:10.1103/PhysRevC.94.034903 [arXiv:1603.04775 [nucl-ex]].

[3] A. Bazavov *et al.* [HotQCD], Phys. Rev. D **90** (2014), 094503 doi:10.1103/PhysRevD.90.094503 [arXiv:1407.6387 [hep-lat]].

[4] A. Aprahamian, A. Robert, H. Caines, G. Cates, J. A. Cizewski, V. Cirigliano, D. J. Dean, A. Deshpande, R. Ent and F. Fahey, *et al.*

[5] W. Busza, K. Rajagopal and W. van der Schee, Ann. Rev. Nucl. Part. Sci. **68** (2018), 339-376 doi:10.1146/annurev-nucl-101917-020852 [arXiv:1802.04801 [hep-ph]].

[6] U. W. Heinz, [arXiv:hep-ph/0407360 [hep-ph]].

[7] N. Armesto and E. Scomparin, Eur. Phys. J. Plus **131** (2016) no.3, 52 doi:10.1140/epjp/i2016-16052-4 [arXiv:1511.02151 [nucl-ex]].

[8] P. Foka and M. A. Janik, Rev. Phys. **1** (2016), 154-171 doi:10.1016/j.revip.2016.11.002 [arXiv:1702.07233 [hep-ex]].

[9] A. Andronic, P. Braun-Munzinger, K. Redlich and J. Stachel, Nature **561** (2018) no.7723, 321-330 doi:10.1038/s41586-018-0491-6 [arXiv:1710.09425 [nucl-th]].

[10] B. Abelev *et al.* [ALICE], Phys. Rev. C **88** (2013) no.4, 044909 doi:10.1103/PhysRevC.88.044909 [arXiv:1301.4361 [nucl-ex]].

[11] M. L. Miller, K. Reygers, S. J. Sanders and P. Steinberg, Ann. Rev. Nucl. Part. Sci. **57** (2007), 205-243 doi:10.1146/annurev.nucl.57.090506.123020 [arXiv:nucl-ex/0701025 [nucl-ex]].

[12] K. J. Eskola, K. Kajantie and J. Lindfors, Nucl. Phys. B **323** (1989), 37-52 doi:10.1016/0550-3213(89)90586-5

[13] E. Iancu, A. Leonidov and L. McLerran, [arXiv:hep-ph/0202270 [hep-ph]].

[14] F. Gelis, E. Iancu, J. Jalilian-Marian and R. Venugopalan, Ann. Rev. Nucl. Part. Sci. **60** (2010), 463-489 doi:10.1146/annurev.nucl.010909.083629 [arXiv:1002.0333 [hep-ph]].

[15] B. Schenke, P. Tribedy and R. Venugopalan, Phys. Rev. Lett. **108** (2012), 252301 doi:10.1103/PhysRevLett.108.252301 [arXiv:1202.6646 [nucl-th]].

[16] A. Kurkela, Nucl. Phys. A **956** (2016), 136-143 doi:10.1016/j.nuclphysa.2016.01.069 [arXiv:1601.03283 [hep-ph]].

[17] A. Kurkela and Y. Zhu, Phys. Rev. Lett. **115** (2015) no.18, 182301 doi:10.1103/PhysRevLett.115.182301 [arXiv:1506.06647 [hep-ph]].

[18] P. B. Arnold, G. D. Moore and L. G. Yaffe, JHEP **01** (2003), 030 doi:10.1088/1126-6708/2003/01/030 [arXiv:hep-ph/0209353 [hep-ph]].

[19] P. M. Chesler and W. van der Schee, Int. J. Mod. Phys. E **24** (2015) no.10, 1530011 doi:10.1142/S0218301315300118 [arXiv:1501.04952 [nucl-th]].

[20] H. Song, S. Bass and U. W. Heinz, Phys. Rev. C **89** (2014) no.3, 034919 doi:10.1103/PhysRevC.89.034919 [arXiv:1311.0157 [nucl-th]].

[21] U. W. Heinz, Landolt-Bornstein **23** (2010), 240 doi:10.1007/978-3-642-01539-7_9 [arXiv:0901.4355 [nucl-th]].

[22] P. Romatschke, Int. J. Mod. Phys. E **19** (2010), 1-53 doi:10.1142/S0218301310014613 [arXiv:0902.3663 [hep-ph]].

[23] B. Schenke, S. Jeon and C. Gale, Phys. Rev. Lett. **106** (2011), 042301 doi:10.1103/PhysRevLett.106.042301 [arXiv:1009.3244 [hep-ph]].

[24] P. Romatschke and U. Romatschke, Phys. Rev. Lett. **99** (2007), 172301 doi:10.1103/PhysRevLett.99.172301 [arXiv:0706.1522 [nucl-th]].

[25] P. Kovtun, D. T. Son and A. O. Starinets, Phys. Rev. Lett. **94** (2005), 111601 doi:10.1103/PhysRevLett.94.111601 [arXiv:hep-th/0405231 [hep-th]].

[26] K. Aamodt *et al.* [ALICE], Phys. Lett. B **708** (2012), 249-264 doi:10.1016/j.physletb.2012.01.060 [arXiv:1109.2501 [nucl-ex]].

[27] G. Nijs and W. van der Schee, SciPost Phys. **15** (2023) no.2, 041 doi:10.21468/SciPostPhys.15.2.041 [arXiv:2112.13771 [nucl-th]].

[28] F. Cooper and G. Frye, Phys. Rev. D **10** (1974), 186 doi:10.1103/PhysRevD.10.186

[29] V. Greco, C. M. Ko and P. Levai, Phys. Rev. C **68** (2003), 034904 doi:10.1103/PhysRevC.68.034904 [arXiv:nucl-th/0305024 [nucl-th]].

[30] C. B. Dover, U. W. Heinz, E. Schnedermann and J. Zimanyi, Phys. Rev. C **44** (1991), 1636-1654 doi:10.1103/PhysRevC.44.1636

[31] M. Bleicher, E. Zabrodin, C. Spieles, S. A. Bass, C. Ernst, S. Soff, L. Bravina, M. Belkacem, H. Weber and H. Stoecker, *et al.* J. Phys. G **25** (1999), 1859-1896 doi:10.1088/0954-3899/25/9/308 [arXiv:hep-ph/9909407 [hep-ph]].

[32] J. Weil *et al.* [SMASH], Phys. Rev. C **94** (2016) no.5, 054905 doi:10.1103/PhysRevC.94.054905 [arXiv:1606.06642 [nucl-th]].

[33] K. J. Eskola, P. Paakkinen, H. Paukkunen and C. A. Salgado, Eur. Phys. J. C **77** (2017) no.3, 163 doi:10.1140/epjc/s10052-017-4725-9 [arXiv:1612.05741 [hep-ph]].

[34] I. Schienbein, J. Y. Yu, K. Kovarik, C. Keppel, J. G. Morfin, F. Olness and J. F. Owens, Phys. Rev. D **80** (2009), 094004 doi:10.1103/PhysRevD.80.094004 [arXiv:0907.2357 [hep-ph]].

[35] D. de Florian, R. Sassot, P. Zurita and M. Stratmann, Phys. Rev. D **85** (2012), 074028 doi:10.1103/PhysRevD.85.074028 [arXiv:1112.6324 [hep-ph]].

[36] T. Matsui and H. Satz, Phys. Lett. B **178** (1986), 416-422 doi:10.1016/0370-2693(86)91404-8

[37] A. Tumasyan *et al.* [CMS], [arXiv:2303.17026 [hep-ex]].

[38] A. Andronic, P. Braun-Munzinger, K. Redlich and J. Stachel, Phys. Lett. B **652** (2007), 259-261 doi:10.1016/j.physletb.2007.07.036 [arXiv:nucl-th/0701079 [nucl-th]].

[39] S. Chatrchyan *et al.* [CMS], JHEP **03** (2015), 022 doi:10.1007/JHEP03(2015)022 [arXiv:1410.4825 [nucl-ex]].

[40] G. Aad *et al.* [ATLAS], Eur. Phys. J. C **79** (2019) no.11, 935 doi:10.1140/epjc/s10052-019-7439-3 [arXiv:1907.10414 [nucl-ex]].

[41] M. Aaboud *et al.* [ATLAS], Phys. Lett. B **790** (2019), 108-128 doi:10.1016/j.physletb.2018.10.076 [arXiv:1805.05635 [nucl-ex]].

[42] S. Chatrchyan *et al.* [CMS], Phys. Lett. B **712** (2012), 176-197 doi:10.1016/j.physletb.2012.04.058 [arXiv:1202.5022 [nucl-ex]].

[43] S. Chatrchyan *et al.* [CMS], Phys. Lett. B **730** (2014), 243-263 doi:10.1016/j.physletb.2014.01.042 [arXiv:1310.0878 [nucl-ex]].

[44] G. Aad *et al.* [ATLAS], Phys. Lett. B **739** (2014), 320-342 doi:10.1016/j.physletb.2014.10.065 [arXiv:1406.2979 [hep-ex]].

[45] J. Casalderrey-Solana, J. G. Milhano and U. A. Wiedemann, J. Phys. G **38** (2011), 035006 doi:10.1088/0954-3899/38/3/035006 [arXiv:1012.0745 [hep-ph]].

[46] P. M. Chesler, K. Jensen, A. Karch and L. G. Yaffe, Phys. Rev. D **79** (2009), 125015 doi:10.1103/PhysRevD.79.125015 [arXiv:0810.1985 [hep-th]].

[47] P. M. Chesler and K. Rajagopal, JHEP **05** (2016), 098 doi:10.1007/JHEP05(2016)098 [arXiv:1511.07567 [hep-th]].

[48] R. Baier, D. Schiff and B. G. Zakharov, Ann. Rev. Nucl. Part. Sci. **50** (2000), 37-69 doi:10.1146/annurev.nucl.50.1.37 [arXiv:hep-ph/0002198 [hep-ph]].

[49] J. Casalderrey-Solana, Y. Mehtar-Tani, C. A. Salgado and K. Tywoniuk, Phys. Lett. B **725** (2013), 357-360 doi:10.1016/j.physletb.2013.07.046 [arXiv:1210.7765 [hep-ph]].

[50] A. M. Sirunyan *et al.* [CMS], JHEP **05** (2018), 006 doi:10.1007/JHEP05(2018)006 [arXiv:1803.00042 [nucl-ex]].

[51] R. Kunnawalkam Elayavalli and K. C. Zapp, JHEP **07** (2017), 141 doi:10.1007/JHEP07(2017)141 [arXiv:1707.01539 [hep-ph]].

[52] C. Park, S. Jeon and C. Gale, Nucl. Phys. A **982** (2019), 643-646 doi:10.1016/j.nuclphysa.2018.10.057 [arXiv:1807.06550 [nucl-th]].

[53] Y. T. Chien and I. Vitev, JHEP **05** (2016), 023 doi:10.1007/JHEP05(2016)023 [arXiv:1509.07257 [hep-ph]].

[54] Y. Tachibana, N. B. Chang and G. Y. Qin, Phys. Rev. C **95** (2017) no.4, 044909 doi:10.1103/PhysRevC.95.044909 [arXiv:1701.07951 [nucl-th]].

[55] V. Khachatryan *et al.* [CMS], Phys. Lett. B **765** (2017), 193-220 doi:10.1016/j.physletb.2016.12.009 [arXiv:1606.06198 [nucl-ex]].

[56] J. Adam *et al.* [ALICE], Nature Phys. **13** (2017), 535-539 doi:10.1038/nphys4111 [arXiv:1606.07424 [nucl-ex]].

[57] R. D. Weller and P. Romatschke, Phys. Lett. B **774** (2017), 351-356 doi:10.1016/j.physletb.2017.09.077 [arXiv:1701.07145 [nucl-th]].

[58] A. Kurkela, U. A. Wiedemann and B. Wu, Phys. Lett. B **783** (2018), 274-279 doi:10.1016/j.physletb.2018.06.064 [arXiv:1803.02072 [hep-ph]].

[59] B. Schenke, S. Schlichting and R. Venugopalan, Phys. Lett. B **747** (2015), 76-82 doi:10.1016/j.physletb.2015.05.051 [arXiv:1502.01331 [hep-ph]].

[60] C. Bierlich, G. Gustafson and L. Lönnblad, Phys. Lett. B **779** (2018), 58-63 doi:10.1016/j.physletb.2018.01.069 [arXiv:1710.09725 [hep-ph]].

[61] L. He, T. Edmonds, Z. W. Lin, F. Liu, D. Molnar and F. Wang, Phys. Lett. B **753** (2016), 506-510 doi:10.1016/j.physletb.2015.12.051 [arXiv:1502.05572 [nucl-th]].

[62] J. Adam *et al.* [ALICE], Phys. Rev. C **91** (2015) no.6, 064905 doi:10.1103/PhysRevC.91.064905 [arXiv:1412.6828 [nucl-ex]].

[63] S. Chatrchyan *et al.* [CMS], Eur. Phys. J. C **74** (2014) no.7, 2951 doi:10.1140/epjc/s10052-014-2951-y [arXiv:1401.4433 [nucl-ex]].

[64] N. Armesto, D. C. Gülhan and J. G. Milhano, Phys. Lett. B **747** (2015), 441-445 doi:10.1016/j.physletb.2015.06.032 [arXiv:1502.02986 [hep-ph]].

[65] M. Aaboud *et al.* [ATLAS], Nature Phys. **13** (2017) no.9, 852-858 doi:10.1038/nphys4208 [arXiv:1702.01625 [hep-ex]].

[66] G. Aad *et al.* [ATLAS], JHEP **03** (2021), 243 [erratum: JHEP **11** (2021), 050] doi:10.1007/JHEP11(2021)050 [arXiv:2008.05355 [hep-ex]].

[67] A. Tumasyan *et al.* [CMS], Phys. Rev. Lett. **131** (2023), 151803 doi:10.1103/PhysRevLett.131.151803 [arXiv:2206.05192 [nucl-ex]].

[68] G. Aad *et al.* [ATLAS], Phys. Rev. Lett. **131** (2023) no.15, 151802 doi:10.1103/PhysRevLett.131.151802 [arXiv:2204.13478 [hep-ex]].

[69] J. Adam *et al.* [STAR], Phys. Rev. Lett. **127** (2021) no.5, 052302 doi:10.1103/PhysRevLett.127.052302 [arXiv:1910.12400 [nucl-ex]].

[70] I. Xu, N. Lewis, X. Wang, J. D. Brandenburg and L. Ruan, [arXiv:2211.02132 [hep-ex]].

[71] S. Acharya *et al.* [ALICE], Phys. Lett. B **817** (2021), 136280 doi:10.1016/j.physletb.2021.136280 [arXiv:2101.04623 [nucl-ex]].

[72] S. Acharya *et al.* [ALICE], [arXiv:2305.06169 [nucl-ex]].