



HAL
open science

Some aspects of the theory of heavy ion collisions

Francois Gelis

► **To cite this version:**

Francois Gelis. Some aspects of the theory of heavy ion collisions. Rept.Prog.Phys., 2021, 84 (5), pp.056301. 10.1088/1361-6633/abec2e . hal-03157106

HAL Id: hal-03157106

<https://hal.science/hal-03157106>

Submitted on 12 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Some Aspects of the Theory of Heavy Ion Collisions

FRANÇOIS GELIS

Institut de Physique Théorique
CEA/Saclay, Université Paris-Saclay
91191, Gif sur Yvette, France

August 31, 2022

Abstract

We review the theoretical aspects relevant in the description of high energy heavy ion collisions, with an emphasis on the learnings about the underlying QCD phenomena that have emerged from these collisions.

1 Introduction to heavy ion collisions

Elementary forces in Nature The interactions among the elementary constituents of matter are divided into four fundamental forces: gravitation, electromagnetism, weak nuclear forces and strong nuclear forces. All these interactions except gravity have a well tested a microscopic quantum description in terms of local gauge theories, in which the elementary matter fields are spin-1/2 fermions, interacting via the exchange of spin-1 bosons. In this framework, a special role is played by the Higgs spin-0 boson (the only fundamental scalar particle in the Standard Model), whose non-zero vacuum expectation value gives to all the other fields a mass proportional to their coupling to the Higgs. The discovery of the Higgs boson at the Large Hadron Collider in 2012 has so far confirmed all the Standard Model expectations. In this picture, gravity has remained a bit of an outlier: even though the classical field theory of gravitation (general relativity) has been verified experimentally with a high degree of precision (the latest of these verifications being the observation of gravitational waves emitted during the merger of massive compact objects - black holes or neutron stars), the quest for a theory of quantum gravity has been inconclusive until now (and possible experimental probes are far out of reach for the foreseeable future).

Strong nuclear force Quantum chromodynamics (QCD), the microscopic theory that governs strong nuclear interactions, was formulated in the early 1970's under the guidance of several experimental clues. In particular, deep inelastic scattering of electrons off proton targets indicated that the proton charge is concentrated into smaller constituents (unresolved in the scattering) of spin 1/2 (this follows from the measured structure functions), that interact weakly at high momentum transfer. These observations paved the way towards a non-Abelian gauge theory with the property of *asymptotic freedom* [1, 2], i.e. a theory in

which the coupling strength decreases at short distance. Combined with some insights from hadron spectroscopy, this led to an $SU(3)$ gauge theory, with spin-1/2 matter fields (the quarks) in the fundamental representation. The fundamental property of QCD that resolved the tension between the fact that quarks must interact strongly enough to form bound hadronic states and the fact that they appear to be weakly interacting in deep inelastic scattering experiments is asymptotic freedom: namely, the property that the running of the coupling due to quantum corrections is such that the strong coupling constant becomes small at short distance and large on distance scales compared to the size of a hadron.

We now know that there are six families of quarks: up, down, strange, charm, bottom and top, ranging from nearly massless to about 175 GeV for the top quark [3]. The nucleons that compose the atomic nuclei of ordinary matter are built solely from the up and down quarks, and the heavier quarks appear only in more massive hadrons (at the exception of the top quark, whose lifetime is so short that it decays before bound states can be formed). QCD has received ample experimental support as the correct microscopic theory for describing strong nuclear interactions. However, because of asymptotic freedom, the most quantitative comparisons between theory and experiments are based on hard processes (i.e., processes involving at least one hard momentum particle in the final state). Although this is sufficient to ascertain the fact that strong nuclear interactions are indeed well described by QCD, these experiments leave unexplored another important aspect of strong interactions, that has to do with the rich properties of nuclear matter in extreme conditions of temperature or density.

Asymptotic freedom, confinement and deconfinement A crucial property of QCD, consequence of asymptotic freedom, is color confinement [4], namely the fact that isolated quarks or gluons cannot exist but instead combine into bound states –the hadrons– in which their color charge is “hidden”. Thus, trying to pull a quark out of a hadron (for instance in a high energy collision with another hadron) merely creates more hadrons. Conversely, when one packs many hadrons in a small volume, the average distance between their constituents decreases, and therefore they interact more and more weakly¹. Given this, we may expect that the forces that bind quarks inside hadrons eventually become weak enough to allow the quarks to become unconfined, i.e., free to wander in the entire volume of the system. This state of nuclear matter is called the *quark-gluon plasma* (QGP). Note that this transition is non-perturbative, since it happens at an energy scale where the coupling constant is still too large to apply reliably perturbation theory.

However, it is possible to formulate QCD non-perturbatively by discretizing Euclidean space-time on a lattice. This setup provides a way of computing certain observables without resorting to an expansion in powers of the coupling. Some of the quantities that one may calculate in lattice QCD are related to the confinement/deconfinement transition: e.g., the expectation value of the trace of a Wilson loop (that one may relate to the potential between a pair of infinitely heavy quark and antiquark), or the entropy density (that measures the number of active degrees of freedom in the system). Moreover, in lattice QCD, one may vary several parameters of the theory, like the number of quark families and their masses, in order to investigate the role they play in the observed phenomena. Some of these results are summarized in the plot of Figure 1, taken from [5]. For instance, pure-gluon QCD (equivalent to QCD with infinitely massive quarks) has a first order deconfinement transition at a temperature of the order of 270 MeV. A first order transition may also exist in the opposite

¹Let us clear out a possible misconception related to this: a small coupling constant in pairwise parton interactions does not necessarily imply that the system as a whole is weakly interacting. The latter is true only if the mean free path is large compared to the De Broglie wavelength of the constituents, and the inverse mean free path is the product of a cross-section by a density. In other words, in such a dense system, there may be strong collective effects despite a weak coupling.

limit, with massless quarks. In this limit, the classical QCD Lagrangian also has a chiral symmetry, spontaneously broken at low temperature (this transition is also a first order transition). The physical spectrum of light quarks in Nature lies in between these two




Figure 1: Nature of the transition at null chemical potential as a function of quark masses. From [5].

extreme situations, and there is now a consensus that this physical point corresponds to a mere crossover transition, i.e., a perfectly smooth (but rather rapid) transition from hadrons to deconfined quarks, that takes place at a temperature around 155 MeV. Even in the absence of a genuine transition in the thermodynamical sense, this transition exhibits signs of deconfinement accompanied by a restoration of chiral symmetry (up to explicit residual violations due to the non-zero quark masses).

Lattice QCD and baryon chemical potential In statistical equilibrium, the parameters that control the possible states of nuclear matter are the temperature, the chemical potentials associated to conserved quantities (the most important of which is the baryon chemical potential μ_B), and possibly some external fields (such as magnetic fields). However, a non-zero chemical potential is a very serious obstacle for lattice QCD. Indeed, the chemical potential turns the exponential of the action integrated over the quark fields into a non-positive measure, that cannot be sampled by Monte-Carlo. At very small μ_B/T , various lattice techniques [6, 7, 8, 9, 10] can be used to track the transition to non-zero values of μ_B , but these methods loose their accuracy when the chemical potential becomes of the order of the temperature. Note that a perturbative approach becomes possible at very large μ_B and/or large T , since these parameters control the relevant scale for the running coupling. At high T and low μ_B , these analytic calculations are consistent with lattice computations. At high μ_B and very low temperature, they indicate the presence of several color super-conducting phases [11, 12, 13, 14] (in these phases, the ground state of the system exhibits a non-zero quark-quark condensate, very similar to Cooper pairing in BCS super-conductivity).

Outside the regions accessible to lattice simulations or to perturbation theory, our knowledge of the phase diagram of nuclear matter is mostly speculative. It is for instance expected that the system is strongly interacting near the transition line, implying small transport co-

efficients, consistent with the measurements of final state correlations among the produced particles (see later the section on hydrodynamics).

Heavy Ion Collisions In the history of the early Universe, the confinement transition was crossed when the Universe was about one microsecond old, but as far as we know this did not leave any visible imprint accessible to present astronomical observations. In the early 1980's emerged the idea to collide heavy nuclei in order to produce in the laboratory nuclear matter at high temperature and density, possibly sufficient to reach and go beyond the critical line. Subsequently, several experiments have had all or part of their scientific program devoted to the study of heavy ion collisions:

- **Bevatron** (Billions of eV Synchrotron) :
From 1954 to 1993 at Lawrence Berkeley National Laboratory, U.S
- **AGS** (Alternating Gradient Synchrotron) :
Since 1960 at Brookhaven National Laboratory, U.S
Now used as injector for RHIC
- **SPS** (Super Proton Synchrotron) :
Since 1976 at CERN
Now the injector for the LHC
- **SIS-18** (Schwer-Ionen-Synchrotron) :
Since 2001 at GSI
- **RHIC** (Relativistic Heavy Ion Collider) :
Since 2000 at Brookhaven National Laboratory, U.S
- **LHC** (Large Hadron Collider) :
Since 2009 at CERN

The first experimental hints of a deconfinement transition were observed at the CERN SPS [15, 16, 17, 18, 19, 20, 21, 22], that collided heavy ions at a center of mass energy of 17 GeV, and in the subsequent experimental programs at higher energies (the RHIC at Brookhaven National Laboratory, and the Large Hadron Collider at CERN) the focus has shifted from assessing the production of a quark-gluon plasma towards measuring quantitatively some of its properties [23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36].

Experimental handles In heavy ion collisions, a few experimental handles are available to vary the conditions in which the quark gluon plasma may be formed. One of them is the atomic number of the nuclei used in the collisions, whose main effect is to change the volume of the interaction zone (but as we shall see in the next section, this has also an incidence on the so-called saturation momentum). When performing collisions with a given species of ions, another variable that has a direct effect on the volume is the impact parameter of each collision. Although the impact parameter is not directly measurable, some observable quantities (such as the total multiplicity in the final state, or the transverse energy) are strongly correlated with the impact parameter. Finally, the collision energy can in principle be varied (but of course, this is in practice highly constrained by the accelerator design), which affects the initial energy density (i.e., temperature) and the net baryon density of the matter produced in a collision. This translates in different reaches in the phase diagram for various heavy ion experiments, as sketched in Figure 2.

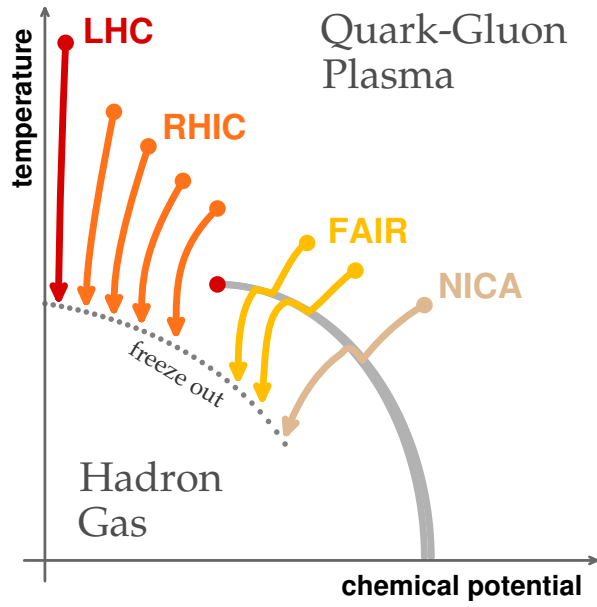


Figure 2: Sketch of the phase-diagram of strongly interacting nuclear matter, and approximate reach of various experimental facilities.




Figure 3: Main stages of the evolution of the matter produced in a heavy ion collision.

Main stages of a heavy ion collision From a theoretical point of view, an ultrarelativistic collision between two nuclei can be conveniently divided in several stages, sketched in the figure 3. Such a picture stems from the approximate boost invariance (i.e., independence on the spatial rapidity variable $\eta_s \equiv \frac{1}{2} \ln((t+z)/(t-z))$) of such collisions, and the successive stage are ordered by the proper time τ . The collision itself is very short due to the Lorentz contraction of the nuclei in the laboratory frame. Just after the collision, the matter produced is predominantly made of non-equilibrated gluons (they are not even on-shell at the very beginning, and the system is better treated in terms of fields rather than particles). This matter is strongly interacting due to a large gluon occupation number, and evolves towards equilibration (both kinetic and chemical, since quark-antiquark pairs are produced in the process). In the subsequent stages, the bulk evolution of the system is remarkably well described by nearly ideal (i.e., with very small values of the viscous transport coefficients) relativistic hydrodynamics. The expansion of the system causes the system to cool down, and at some point the temperature reaches the confinement temperature. In the framework of hydrodynamics, as long as the system remains close to equilibrium, the crossing of the confinement transition is rather transparent since it is encoded in the equation of state. Soon after, the system becomes dilute, the mean free path increases, and a description of its expansion in terms of kinetic theory rather than hydrodynamics becomes preferable. In such a description, the values of the various cross-sections control when each type of reaction stops: the inelastic processes stop first (chemical freezeout), soon followed by a kinetic freezeout after which the momenta of the particles remain unchanged (afterwards, all particles therefore fly on straight lines at constant velocity until they hit a detector).

2 Initial state, Color Glass Condensate

Multiparton interactions and gluon saturation Let us start with the very first moments of a heavy ion collision. This is the realm of the highest momentum scales of the entire collision, and one may thus expect that this stage is amenable to a perturbative QCD treatment. The situation is however more complicated. The main difficulty is the fact that the typical transverse momentum of the produced particles in such a collision is rather low (around 1 GeV), implying that high energy heavy ion collisions probe the partonic content of the incoming nuclei at very small values of the longitudinal momentum fraction x (the fraction of the momentum of a hadron carried by one of its constituents), i.e., in a regime where the gluon density is large. The consequence of this is that processes initiated by more than one parton in each nucleus become possible, which invalidates the usual factorization schemes (based on single parton densities), as illustrated in Figure 4. One may derive a




Figure 4: Left: scattering in the dilute regime. Right: multi-gluon scattering in the saturation regime.

simple criterion for the onset of these multi-parton phenomena (known as *gluon saturation*) by combining the estimated gluon recombination cross-section and the number of gluons per unit of transverse area [37, 38, 39]. Gluon saturation happens when the product of these

two quantities is greater than one, which can also be framed as an upper bound for the momentum transfer Q in deep inelastic scattering (the inverse of this scale plays the role of a spatial resolution in such a scattering), $Q \lesssim Q_s$, where Q_s is the so-called saturation momentum. Q_s depends both on the atomic number of the nuclei, and on the collision energy via the momentum fraction x ,

$$Q_s^2 \sim A^{1/3} x^{-\lambda}, \quad (1)$$

where the exponent λ has been determined phenomenologically from deep inelastic data and estimated to be $\lambda \approx 0.25$. The growth at small x follows from the growth of the gluon density, while the factor $A^{1/3}$ is a measure of the thickness of a nucleus in the direction of the collision axis. The variations of the saturation momentum as a function of A and x are shown in Figure 5. Recalling that the typical value of x scales as p_\perp/\sqrt{s} where p_\perp is




Figure 5: Saturation domain. From [40].

the transverse momentum of produced particles and \sqrt{s} the energy in the center of mass of a nucleon-nucleon collision, it appears that the bulk of particle production in heavy ion collisions at the energy of the LHC is potentially affected by gluon saturation, which calls for a theoretical framework going beyond the usual collinear factorization. Indeed, the standard parton distribution are single-parton densities and do not contain the necessary information about the multi-parton initial states that become important in the saturation regime.

Color Glass Condensate Extending the framework of collinear factorization by defining multi-parton densities in the same way as the usual parton distributions is not practical. Instead, one exploits the fact that gluon saturation is also a regime of large gluon occupation number, which allows to treat the gluon field as classical in a first approximation [41, 42]. In such a description, the relevant information about an incoming nucleus is the color current it carries, that acts as a source for the color field. Moreover, this current is produced by the partons that are comparatively fast (in the observer's frame), which implies that it is nearly time independent thanks to time dilation. The degrees of freedom in such a description are thus color currents (one for each projectile) coupled to gluon fields, with an effective action

$$\mathcal{S}_{\text{eff}} \equiv -\frac{1}{4} F_a^{\mu\nu} F_{\mu\nu a} + J_a^\mu A_{\mu a}. \quad (2)$$




Figure 6: Illustration of the McLerran-Venugopalan model.

This setup, known as the *Color Glass Condensate* [43, 44, 45, 46, 47, 48] (CGC) is illustrated in Figure 6. For a fast moving projectile in the $+z$ direction, the color current has the form

$$J_a^\mu(x) = \delta^{\mu+} \rho_a(x^-, \mathbf{x}_\perp), \quad (3)$$

where $\rho_a(x^-, \mathbf{x}_\perp)$ is a function that describes the spatial distribution of the color charges in the object under consideration. In this expression, we have introduced the light-cone coordinates, $x^\pm \equiv (t \pm z)/\sqrt{2}$ (the notation $\delta^{\mu+}$ means that the only non-zero component of the current is $J^+ \propto J^- + J^3$). These coordinates are convenient when discussing the kinematics of an object moving at the speed of light, since x^+ acts as the time variable for this object (the fact that the above current does not depend on x^+ simply reflects the fact that this object is time-independent) and x^- as a longitudinal coordinate as measured in the rest frame of the object. The support of the x^- dependence of the current is very narrow and centered around $x^- = 0$, due to Lorentz contraction. Note also that the current must be covariantly conserved, $[D_\mu, J^\mu] = 0$. Since the covariant derivative contains the color field, the color current may be affected by its own radiated field (the light-cone gauge $A^- = 0$ mitigates this difficulty, since the current J^+ can be altered only by the field A^-). The \mathbf{x}_\perp dependence of $\rho_a(x^-, \mathbf{x}_\perp)$ reflects the positions in the transverse plane of the color sources at the instant of the collision (the duration of the collision, controlled by the thickness of the Lorentz contracted nuclei, is much shorter than the typical timescales of the internal motions of the constituents of a hadron – thus J^μ needs only to provide a snapshot of the hadron content). But of course, the configuration of these charges is not known event-by-event, and the best we may hope to know is a statistical distribution of these densities, encoded in a functional $W[\rho]$. This functional density is not something that we can calculate perturbatively in QCD, since it depends on aspects such as confinement, the nuclear wavefunction, etc... For a large nucleus, the McLerran-Venugopalan model [41, 42], in which $W[\rho]$ is a Gaussian,

$$W[\rho] = \exp\left(-\int d^2\mathbf{x}_\perp \frac{\rho_a(x^-, \mathbf{x}_\perp)\rho_a(x^-, \mathbf{x}_\perp)}{2\mu^2(x^-, \mathbf{x}_\perp)}\right), \quad (4)$$

is often employed due to its simplicity (in some cases, it even allows analytical calculations). In this distribution, the mean value of the charge distribution at a point (x^-, \mathbf{x}_\perp) is zero,

and $\mu^2(x^-, \mathbf{x}_\perp)$ is a measure of its fluctuations. At lowest order in the CGC effective theory, this parameter is a placeholder for the value of the saturation momentum,

$$Q_s^2(\mathbf{x}_\perp) \propto g^2 \mu^2(\mathbf{x}_\perp) \ln \left(\frac{\mu^2(\mathbf{x}_\perp)}{\Lambda_{QCD}} \right), \quad \text{with } \mu^2(\mathbf{x}_\perp) \equiv \int dx^- \mu^2(x^-, \mathbf{x}_\perp). \quad (5)$$

(This correspondence may be established by calculating the DIS cross-section in the McLerran-Venugopalan model [49].) Although a possible heuristic justification for this Gaussian model is the central limit theorem, thanks to the fact that a large nucleus has many constituents per unit of transverse area, one should keep in mind that this distribution of ρ_a 's is not derived from first principles in QCD (since doing so would require to control QCD in a non-perturbative regime). Another reason why the Gaussian in eq. (4) does not have a very fundamental standing is that this shape is not preserved when one includes one-loop corrections: indeed, these corrections contain large logarithms of energy that turn $W[\rho]$ into a non-Gaussian energy-dependent distribution.

Power counting in the saturated regime Let us now describe how a typical CGC calculation is organized. The color glass condensate may be viewed as a Yang-Mills theory coupled to an external source [50, 51], which diagrammatically means that all graphs contain only gluon propagators. Their endpoints can be attached to the sources, to gluon vertices, or to the observable of interest. In the saturation regime, the power counting for these graphs is a bit peculiar due to the large gluon occupation number. The 3-gluon and 4-gluon vertices are respectively of order g and g^2 , while the external source can be as large as g^{-1} (this order of magnitude is reached when the occupation number reaches its maximal value, of order g^{-2}). Therefore, the order of magnitude of a generic connected graph \mathcal{G} is

$$\mathcal{G} \sim g^{-n_E} g^{2n_L} (gJ)^{n_J}, \quad (6)$$

where n_E is the number of external gluons, n_L the number of loops and n_J the number of sources J in the graph. We see from this formula that when $J \sim g^{-1}$, the order of the graph does not depend on the number of sources, implying that there is an infinity of graphs contributing at each order (for instance, the *leading order* is made of all the trees). The saturation regime is therefore a strongly interacting non-perturbative situation, despite the fact that the coupling constant may be small at high energy (the typical scale at which the running strong coupling constant should be evaluated is governed by the saturation momentum).

Leading order At leading order, the infinite series of tree diagrams that one needs to sum can always (for inclusive observables) be expressed in terms of classical solutions of the Yang-Mills equations,

$$[D_\mu, F^{\mu\nu}] = J^\nu, \quad (7)$$

with retarded boundary conditions (the retarded nature of the boundary condition follows from the fact that inclusive measurements do not put any restriction on the final state). Since in the CGC the incoming projectiles are completely encoded in the source J^μ , the classical initial condition is simply to have a null field (or more generally a pure gauge) in the remote past, before the collision has happened. By causality, space-time is naturally divided into four domains shown in Figure 7. Domain 0 corresponds to space-time points not yet reached by any of the nuclei. Domain 1 is causally connected with the right moving nucleus, but not with the left moving one. For domain 2, it is the opposite. And finally the domain 3 contains the outcome of the collision. Given a null initial condition in domain 0, the




Figure 7: Decomposition of space-time in four domains, when solving the classical Yang-Mills equations.

gauge field can be obtained analytically in the domains 1,2, and also at the lower boundary of domain 3 (i.e., at a proper time $\tau = 0^+$). At later times in domain 3, no analytical solution is known, but it is rather straightforward to solve the classical Yang-Mills equations numerically [52, 53, 54, 55, 56, 57].

Next to Leading Order CGC calculations can in principle be pushed to next-to-leading order, i.e., one-loop. The main difficulty in doing this is that, like with the leading order, there is an infinite set of diagrams contributing at NLO. These are all the one-loop graphs embedded in the external classical gauge field obtained at LO. For any inclusive observable, there exists an exact relationship between the LO (tree level) and NLO (one loop) results, that schematically reads [58, 59]

$$\mathcal{O}_{\text{NLO}} = \frac{1}{2} \int \frac{d^3 \mathbf{k}}{(2\pi)^3 2E_{\mathbf{k}}} \int_{x^0=y^0=-\infty} d^3 \mathbf{x} d^3 \mathbf{y} \left[e^{+ik \cdot x} \frac{\delta}{\delta \mathcal{A}_{\text{ini}}(x)} \right] \left[e^{-ik \cdot y} \frac{\delta}{\delta \mathcal{A}_{\text{ini}}(y)} \right] \mathcal{O}_{\text{LO}}, \quad (8)$$

where \mathcal{A}_{ini} is the initial condition for the classical field in the LO calculation. This relation-




Figure 8: Left: typical one-loop graph contributing to an inclusive observable. Right: interpretation of Eq. (8). In this graphical sketch, the horizontal direction represents space and the vertical direction represents time. The red dots represent the external source J .

ship indicates that the NLO may be obtained from the LO simply by replacing at two points x, y the classical initial condition by plane waves $e^{+ik \cdot x} e^{-ik \cdot y}$, integrated over all on-shell momenta k (but the time evolution continues to be governed by the classical equation of motion). This is consistent with the fact that the first quantum correction (order \hbar^1) only affects the initial state, and the first quantum effect that alters the evolution arises at the

order \hbar^2 . In ordinary quantum mechanics, this is best seen in the phase-space formulation (also known as Moyal-Groenewold formulation), in which each quantum operator \mathbf{A} is replaced by its Wigner transform,

$$A(X, p) \equiv \int ds e^{i\frac{p \cdot s}{\hbar}} \langle X + \frac{s}{2} | \mathbf{A} | X - \frac{s}{2} \rangle, \quad (9)$$

defined on the classical phase-space variables (X, p) . In this mapping, the commutator of two such operators \mathbf{A}, \mathbf{B} becomes the *Moyal bracket* of their Wigner transforms:

$$\{\{A, B\}\}(X, p) \equiv \frac{2}{i\hbar} A(X, p) \sin\left(\frac{\hbar}{2}\left(\overleftarrow{\partial}_p \overrightarrow{\partial}_x - \overleftarrow{\partial}_x \overrightarrow{\partial}_p\right)\right) B(X, p). \quad (10)$$

If we denote W the Wigner transform of the density operator ρ of a system, the von Neumann equation $i\hbar\partial_t\rho = [H, \rho]$ becomes

$$\partial_t W = \{\{\mathcal{H}, W\}\} = \underbrace{\{\mathcal{H}, W\}}_{\text{Poisson bracket}} + \mathcal{O}(\hbar^2), \quad (11)$$

where \mathcal{H} is the Wigner transform of H (i.e., \mathcal{H} is the classical Hamiltonian).

Large logarithms and JIMWLK evolution When calculating such one-loop corrections, an aspect of the CGC degrees of freedom that we have ignored until now becomes important, namely that one must introduce a cutoff to separate the color sources from the gauge fields. For instance, such a cutoff should limit the range of integration over the momentum k in Eq. (8). This separation is based on the longitudinal momentum (or equivalently, the rapidity $y \equiv \frac{1}{2} \ln((k_0 + k_z)/(k_0 - k_z))$), and is mandatory when evaluating loops to avoid double countings. In practice, the longitudinal component of the loop momentum must remain below the cutoff, because all the higher momentum modes are already included in the color current. This leads to all one-loop correction to be sensitive (proportional to the logarithm of the cutoff on longitudinal momentum) to the cutoff [60, 61, 43, 44, 45, 58, 59, 62]. However, since this cutoff is an ad-hoc parameter of the CGC effective theory rather than a physical parameter, it should not enter in physical observables. This paradox is resolved by the fact that the cutoff dependence is universal, in the sense that it depends on the nature of the two colliding projectiles, but not on the inclusive observable one is measuring. Therefore, it is possible to absorb the cutoff dependence into a redefinition of the distributions $W[\rho]$ that define the color source content of the projectiles, turning them into cutoff dependent objects. For this to be feasible, one should perform an average of the ρ -dependent observable, weighted by the distributions of ρ 's of each projectile,

$$\langle \mathcal{O} \rangle \equiv \int [D\rho_1 D\rho_2] W_1[\rho_1] W_2[\rho_2] \mathcal{O}(\rho_1, \rho_2). \quad (12)$$

The cutoff dependence of $W[\rho]$ is controlled by the so-called JIMWLK equation, schematically of the form

$$\frac{\partial W[\rho]}{\partial \log \Lambda} = \frac{\delta}{\delta \rho_a} \chi_{ab} \frac{\delta}{\delta \rho_a} W[\rho], \quad (13)$$

where χ_{ab} depends on the LO classical field created by the source ρ (the possibility to transfer the cutoff dependence from the observable to the distribution $W[\rho]$ is made possible by the fact that the operator in the right hand side of the JIMWLK equation is self-adjoint, via integration by parts). Thus, by evolving the distributions $W[\rho]$ of each projectile to values of the longitudinal momentum in the immediate vicinity of the scales relevant for the observable of interest, one resums all the leading logarithms, i.e., the powers $(g^2 \log \Lambda)^n$

where each logarithm of the cutoff is accompanied by a factor g^2 . This is very similar in spirit to collinear factorization, the unphysical cutoff Λ playing the role of a factorization scale that should disappear from observables.

Since it is a functional equation, the JIMWLK equation is difficult to solve, even numerically. The only known approach so far uses the fact that the JIMWLK equation acts like a diffusion equation in the functional space of the ρ 's (in a treatment more rigorous than this general discussion, one would use Wilson lines built from the ρ 's rather than the ρ 's themselves), and therefore can be rewritten as a Langevin equation [63]. Then, after discretization of the transverse plane, this stochastic equation is amenable to a numerical treatment in order to obtain an ensemble of ρ 's evolved to the relevant value of the cutoff [64, 65, 66]. Let us also mention recent improvements: a modification of the Langevin equation has been proposed to include the effects of a running coupling [67], and the full NLO corrections to the JIMWLK equation have also been evaluated [68, 69] (but not yet implemented in a numerical code).

Balitsky-Kovchegov equation To avoid this computationally heavy approach, it is also possible to truncate the JIMWLK equation. The first thing to note is that the functional form of the JIMWLK equation is equivalent to an infinite sequence of equations for the correlation functions of Wilson lines constructed from the ρ 's

$$U(\mathbf{x}_\perp) \equiv T \exp ig \int dx^- \frac{1}{\nabla_\perp^2} \rho(x^-, \mathbf{x}_\perp). \quad (14)$$

These equations are nested: the equation that drives the cutoff dependence of the 2-point correlation function depends on a 4-point function, etc... A possible approximation (that may be justified in the limit of a large number of colors) consists in factorizing this 4-point function as a product of two 2-point functions, which has the effect of closing the evolution equation of the latter. The resulting equation, known as the Balitsky-Kovchegov equation [70, 71], reads

$$\frac{\partial T_{\mathbf{x}\mathbf{y}}}{\partial \log \Lambda} = \frac{\alpha_s N_c}{2\pi^2} \int d^2 z_\perp \frac{(\mathbf{x}_\perp - \mathbf{y}_\perp)^2}{(\mathbf{x}_\perp - \mathbf{z}_\perp)^2 (\mathbf{y}_\perp - \mathbf{z}_\perp)^2} \left\{ T_{\mathbf{x}\mathbf{z}} + T_{\mathbf{z}\mathbf{y}} - T_{\mathbf{x}\mathbf{y}} - T_{\mathbf{x}\mathbf{z}} T_{\mathbf{z}\mathbf{y}} \right\}, \quad (15)$$

where

$$T_{\mathbf{x}\mathbf{y}} \equiv 1 - N^{-1} \text{tr} \langle U(\mathbf{x}_\perp) U^\dagger(\mathbf{y}_\perp) \rangle \quad (16)$$

(with Wilson lines taken in the fundamental representation of $su(N)$). $T_{\mathbf{x}\mathbf{y}}$ is also proportional to the scattering amplitude of a quark-antiquark dipole (at the transverse coordinates \mathbf{x}_\perp and \mathbf{y}_\perp , respectively) off a high energy nucleus. In this equation, the first three terms, linear, correspond to the BFKL equation, and the last term, non-linear, is a correction due to gluon saturation, that becomes sizeable when the scattering amplitude approaches the unitarity limit $T = 1$. Since it is an equation for an ordinary function, the BK equation is much easier to solve numerically. Note also that the previous equation has now been improved by running coupling corrections [72, 73], by the full next-to-leading log corrections [74], and by a resummation of collinear logarithms [75, 76]. These improvements have allowed a successful phenomenology of small- x phenomena in deep-inelastic scattering and forward proton-proton or proton-nucleus collisions based on the BK evolution equation [77, 78, 79, 80].

3 Pre-equilibrium evolution just after the collision

Energy-momentum tensor at Leading Order As we have seen in the previous section, large nuclei at high energy may be described by using the CGC framework, in which the

large momentum degrees of freedom are treated as random color currents coupled to the color field. At leading order in the strong coupling constant, all expectation values are given by tree diagrams, whose sum is the classical solution of Yang-Mills equations with a null retarded boundary condition. Having in mind a description of the subsequent stages of the collision in terms of relativistic hydrodynamics, it is therefore natural to calculate the components of the energy momentum tensor. In the CGC framework, the dominant contribution comes from the gluons (the valence quarks have a very small contribution at small x , and the sea quark distribution is suppressed by a power of α_s compared to that of the gluons). In a classical field, they are given by the following formulas

$$T_{\text{LO}}^{00} = \frac{1}{2} \underbrace{[\mathbf{E}^2 + \mathbf{B}^2]}_{\text{class. fields}}, \quad T_{\text{LO}}^{0i} = [\mathbf{E} \times \mathbf{B}]^i,$$

$$T_{\text{LO}}^{ij} = \frac{\delta^{ij}}{2} [\mathbf{E}^2 + \mathbf{B}^2] - [\mathbf{E}^i \mathbf{E}^j + \mathbf{B}^i \mathbf{B}^j], \quad (17)$$

in terms of the chromo-electric and chromo-magnetic fields (note that this tensor is traceless in classical Yang-Mills theory – a non-zero trace would arise from loop corrections via the β -function, and from explicit quark masses when quarks are taken into account). Let us first mention the very special configuration of the \mathbf{E} and \mathbf{B} fields just after the collision: at $\tau = 0^+$, these two fields are both parallel to the collision axis [81], which leads to the following form of $T_{\text{LO}}^{\mu\nu}(\tau = 0^+)$,

$$T_{\text{LO}}^{0i} = 0, \quad T_{\text{LO}}^{11} = T_{\text{LO}}^{22} = T_{\text{LO}}^{00}, \quad T_{\text{LO}}^{33} = -T_{\text{LO}}^{00}. \quad (18)$$

In other words, the matter is produced at rest, with a negative longitudinal pressure (i.e., the system resists longitudinal expansion). Such a negative pressure means that this system should not be viewed as a collection of on-shell particles, but rather as fields. At later times




Figure 9: Time evolution at Leading Order of the ratios T^{11}/T^{00} and T^{33}/T^{00} .

(see Figure 9), the longitudinal pressure progressively builds up to become positive at a time around $Q_s \tau \sim 1$, which is indeed the time at which the color fields may be interpreted as nearly on-shell gluons [82]. However, it is also clear from this plot that the ratio of longitudinal to transverse pressure remains very small. In fact, at leading order, this ratio decreases as τ^{-2} , which is characteristic of a free streaming system (i.e., a system whose self-interactions are too weak to compete with the longitudinal expansion).

Next to Leading Order, Instabilities Such a behavior of the longitudinal pressure is not consistent with hydrodynamical evolution, where the ratio P_L/P_T would instead

increase to eventually approach unity. For an underlying QCD description to allow a smooth matching to a subsequent hydrodynamical expansion, there should be a range of times where the two descriptions lead to similar behaviors. It turns out that higher order corrections in the CGC description are potentially more important than the power counting suggests. Indeed, the power counting correctly states that loop corrections (i.e. corrections beyond the classical field approximation) are suppressed by additional powers of the coupling constant, but it implicitly assumes that the coefficients in this power expansion remain of order one at all times. It is this assumption that turns out to be incorrect, because the classical solutions of Yang-Mills equations are subject to instabilities that make them exponentially sensitive to their initial conditions [83, 84, 85, 86], combined with the fact that one-loop corrections can be expressed in terms of small perturbations to the initial condition of the LO classical field.




Figure 10: Graphs that have the leading time behavior in the presence of instabilities.

Resummation, Classical Statistical Approximation By a more careful power counting that keeps track of these terms that have an exponential growth with time [87], it is possible to determine the set of graphs that provide the leading contributions at large time. These graphs are shown in the Figure 10 (and in the left part of Figure 8, we show for comparison a typical next-to-leading order contribution). It turns out that the sum of this infinite set of higher-loop order graphs can be obtained by letting the initial value of the gluon field fluctuate around its classical value, with a Gaussian distribution which is completely determined by the one-loop result. Schematically, this resummed result reads:

$$\mathcal{O}_{\text{resummed}} = \int [Da] \exp \left[-\frac{1}{2\hbar} \int_{\mathbf{x}, \mathbf{y}} a(\mathbf{x}) \mathbf{\Gamma}_2^{-1}(\mathbf{x}, \mathbf{y}) a(\mathbf{y}) \right] \mathcal{O}_{\text{LO}}[\mathcal{A}_{\text{in}} + a]. \quad (19)$$

Such a resummation is of course only a part of the full answer (with an appropriate choice of the variance $\mathbf{\Gamma}_2$, one may obtain the complete LO and NLO results, but only a subset of the higher orders), known in the literature as the *Classical Statistical Approximation* (CSA). We have written the explicit dependence of the Gaussian distribution with respect to \hbar to emphasize the quantum nature of these fluctuations of the field. In the special case of heavy ion collisions, the variance of the Gaussian fluctuations, $\mathbf{\Gamma}_2$, can be determined analytically at the time $\tau = 0^+$, i.e., just after the collision [88]. Note that in the vacuum, this variance reads (this is just a sketch, that ignores the complications due to gauge fields)

$$\langle a(x)a(y) \rangle = \mathbf{\Gamma}_2(x, y) = \int \frac{d^3\mathbf{k}}{(2\pi)^3 E_{\mathbf{k}}} \frac{1}{2} e^{i\mathbf{k}\cdot(x-y)}. \quad (20)$$

The factor 1/2 in the integrand can be interpreted as the zero point occupation of the vacuum (at one-loop, one may show that the only quantum effect is the fact that the ground

state is not empty but subject to zero point fluctuations – the resummation considered here is an approximation in which this is extended to higher-loop orders).

Once the variance Γ_2 is known, Eq. (19) can be evaluated numerically on a lattice in a straightforward way, since it simply amounts to reproducing the leading order classical CGC computation with a fluctuating initial condition. The Gaussian integral in Eq. (19) can be estimated by Monte-Carlo sampling. These simulations, with the variance given in Eq. (20), lead to an increase of the ratio of longitudinal to transverse pressure [89]. However, the interpretation of this result is obscured by the fact that this setup has no continuum limit when the lattice spacing goes to zero. This can be understood easily from the fact that Eq. (20) corresponds to a flat momentum distribution of gluons, that extends to arbitrarily large momenta (only cut-off by the inverse lattice spacing).

CSA with a compact fluctuation spectrum An alternative to the variance given in Eq. (20) that does not have this ultraviolet problem would be to replace the factor $1/2$ by another gluon distribution that has a fast enough fall-off at large momentum,

$$\langle a(x)a(y) \rangle_{\text{alt}} = \int \frac{d^3\mathbf{k}}{(2\pi)^3 E_{\mathbf{k}}} f_0(k) e^{ik \cdot (x-y)}. \quad (21)$$

Although such a distribution cannot be derived from first principles, unlike Eq. (20), a handwaving argument in favor of it is that after a time of order Q_s^{-1} the gluons produced in a collision are nearly on-shell with a compact distribution that extends up to momenta $k \sim Q_s$ (with an occupation number of order g^{-2} within this support). With such a spectrum of initial field fluctuations, the behavior of the ratio of pressures P_L/P_T is at odds with what was obtained with Eq. (20), everything else being equal: with Eq. (21), one has $P_L/P_T \sim \tau^{-2/3}$, showing no sign of isotropization [90, 91]. In this scenario, it is argued that isotropization is delayed until the gluon occupation number becomes of order one, which would happen eventually at a time $Q_s \tau \sim \alpha_s^{-3/2}$.

Going beyond the classical statistical approximation in a field theoretical framework is possible with the two-particle irreducible formalism [92, 93, 94]. This formalism amounts to a self-consistent determination of the propagator (which in a many-body context also contains the information about the particle distribution) by resumming a self-energy –itself a function of the propagator– on the propagator. The 2PI framework can be renormalized [95, 96], thereby avoiding the issues with the CSA and zero point vacuum fluctuations, and can thus be used to track the real-time evolution of a system starting from any quantum state. The main drawback of this approach is that it is very demanding in terms of computational resources, especially in the case of an expanding system like the one formed in a heavy ion collision. At the time of this writing, there has only been one “proof of concept” implementation for an expanding system [97], in which the questions related to isotropization were not investigated.

Kinetic theory and Boltzmann equation A less costly alternative is *kinetic theory*, that one may obtain from the 2PI approach provided one makes two additional approximations:

- Quasi-particle approximation: this amounts to assuming that the propagator describes on-shell infinitely long lived particles. With this assumption (that can only be valid in a system where the mean free path is much larger than the De Broglie wavelength of the particles), the only unknown in the propagator is the particle distribution $f(x, \mathbf{p})$.
- Gradient approximation: this amounts to assuming that the spatial variations of the system due to its off-equilibriumness occur only on time and distance scales much larger than the De Broglie wavelength of the particles.

With these two approximations, the Kadanoff-Baym equation of motion of the 2PI formalism reduces to a much simpler Boltzmann equation, schematically of the form

$$\begin{aligned}
(\partial_t + \mathbf{v}_{\mathbf{p}_1} \cdot \nabla_{\mathbf{x}}) f_1 &= C_{\mathbf{p}_1}[f] & (22) \\
&= \int_{\mathbf{p}_{2,3,4}} |M(12 \rightarrow 34)|^2 \delta(p_1 + p_2 - p_3 - p_4) \\
&\quad \times \{f_3 f_4 (1 + f_1)(1 + f_2) - f_1 f_2 (1 + f_3)(1 + f_4)\}. & (23)
\end{aligned}$$

The first line is the generic Boltzmann equation obtained when using these two approximations, with a *collision term* local in x that can a priori contain arbitrary orders in the distribution f . In the second line, we have written specifically the Boltzmann equation that includes only $2 \rightarrow 2$ elastic reactions (but this truncation is an extra approximation that goes beyond the quasi-particle and gradient approximations).

Testing semi-classical approximations via kinetic theory Besides being a tool to study the evolution of the particle distribution, the Boltzmann equation can also provide some insights about the limitations of the classical statistical approximation, because identical classical approximations may be applied to the right hand side of the Boltzmann equation. To that effect, the correspondence is the following [98, 99]

1. CSA with Eq. (21) \iff keep only the terms cubic in f ,
2. CSA with Eq. (20) \iff replace $f \rightarrow f + \frac{1}{2}$ in the previous approximation (note that in this approximation, the collision integral has the correct cubic and quadratic terms, and also some spurious linear terms not present in the exact collision term).

With the approximation 2 of the collision integral, one can for instance reproduce quantitatively the ultraviolet sensitivity (i.e., the lack of continuum limit) of the CSA when zero point fluctuations are included [100]. With the approximation 1, the Boltzmann equation (for an expanding system) leads to the behavior $P_L/P_T \sim \tau^{-2/3}$. It is also possible to understand by kinetic arguments why this approximation misses important physics for a longitudinally expanding system [101]. In such a system, isotropization results from a competition between the expansion of the system, that drives the particles towards a more anisotropic distribution, and collisions that tend to redistribute the directions of the velocities (for collisions that are sufficiently isotropic, which is the case for a dense system due to a strong Debye screening). However, as shown in Figure 11, when two particles (labeled 1 and 2) from a very anisotropic distribution scatter out-of-plane, momentum conservation implies that the two final state particles (labeled 3 and 4) end in the empty region of phase-space. Such a scattering is forbidden when one keeps only the cubic terms in the collision integral, because these terms correspond to stimulated emission, that can only happen when one of the produced particles goes in an already populated region. The large-angle scattering process of Figure 11 is allowed only by the terms quadratic in f , that are not present in the approximation of Eq. (21). The difference between a semi-classical approximation such as Eq. (21) and one that includes the zero-point fluctuations can then be seen on the time evolution of the ratio P_L/P_T , as shown in Figure 12 in the case of a scalar theory (see [101] for details): in this figure, we see that the classical approximation leads to a decrease of this ratio as $\tau^{-2/3}$, while with the full collision term it first decreases (while the expansion outpaces the scatterings) and then increases (when the expansion rate has become low enough) to approach unity. It is only for very small couplings (i.e., unrealistically large values of the ratio of the shear viscosity to entropy, η/s – see the next section for a discussion of the range of values of this ratio expected in heavy ion collisions) that this classical approximation agrees with the evolution driven by the full collision term long enough to reach the asymptotic scaling regime.




Figure 11: Typical $2 \rightarrow 2$ scattering in a highly anisotropic system.

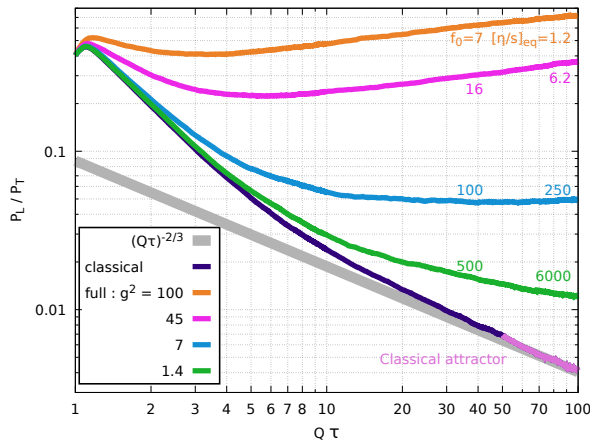


Figure 12: Comparison of the classical and full evolution of P_L/P_T in kinetic theory for a longitudinally expanding system. The dark purple curve is the kinetic theory analogue of the classical statistical approximation based on eq. (21).

Kinetic approach in Yang-Mills theory A similar computation has also been performed in the more realistic setting of Yang-Mills theory [102, 103], with similar results as shown in Figure 13. There also, one sees the classical approximation depart from the full evolution fairly quickly. For instance, the curve $\lambda = 0.5$, i.e. $\alpha_s = 0.02$ for $N_c = 2$, deviates from the classical approximation around $Q_s \tau \approx 2$. Moreover, this happens much earlier than the presumed range of validity $Q_s \tau \lesssim \alpha_s^{-3/2} \approx 350$ predicted within the classical approximation itself (in fact, we see from this plot that the point $Q_s \tau = 350$ on the classical evolution is orders of magnitude off the correct trajectory). In addition, this computation has shown a very good agreement with second order hydrodynamics already at times where isotropization is still far from being achieved, thereby providing a justification for the applicability of hydrodynamics as early as $\tau \approx 0.6$ fm/c (compare the red and solid black curves in Figure 14).

Fixed points of kinetic evolution Kinetic theory has also been used to obtain results about the fate of a system that undergoes longitudinal expansion [104, 105, 106, 107], with only generic assumptions about the strength of the collisions. The starting point of this




Figure 13: Isotropization in Yang-Mills theory in a kinetic description. In this figure $\lambda \equiv g^2 N_c$ is the 't Hooft coupling. The solid black curve is the kinetic theory analogue of the classical statistical approximation based on eq. (21). From [103].




Figure 14: Comparison of kinetic theory with second order hydrodynamics. From [103].

study is the Boltzmann equation in the *relaxation time approximation*,

$$\left(\partial_\tau - \frac{p_z}{\tau} \right) f(\tau, \mathbf{p}) = -\frac{f - f_{\text{eq}}}{\tau_R}. \quad (24)$$

In this equation, τ_R is a relaxation time that controls how fast the particle distribution relaxes to its local equilibrium value. This parameter may be chosen in various ways:

- $\tau_R = \infty$: for a collisionless system,
- $\tau_R \sim \epsilon^{-1/4}$: for a “conformal” system, i.e., a system where the collision rate scales as the inverse temperature,
- $\tau_R = \text{const}$: for a fixed collision rate (although this is not very realistic with expansion).

Then, one may define the following moments,

$$L_n \equiv \int_{\mathbf{p}} \mathbf{p}^2 P_{2n}(p_z/p) f(\tau, \mathbf{p}), \quad g_n \equiv \tau \partial_\tau \ln L_n. \quad (25)$$




Figure 15: Kinetic evolution of the exponents g_n for many distinct initial conditions. The dotted line indicates the interacting fixed point. From [105].

(In this definition, $P_{2n}(z)$ is the order- $2n$ Legendre polynomial. Note that $L_0 = \epsilon = P_L + 2P_T$, $L_1 = P_L - P_T$.) The quantity g_n gives the exponent in the power law behavior of the corresponding moment. The Boltzmann equation can then be transformed into an infinite set of linear equations that govern the evolution of the moments L_n, L_{n+1} and L_{n-1} , in which all but one coefficients depend only on the left hand side of the Boltzmann equation. Then, it was observed that these coupled equations admit two (and only two) fixed points:

- A *free-streaming fixed point*, obtained for $\tau_R = \infty$, where all the L_n behave as τ^{-1} (i.e., all the g_n go to -1), with fixed ratios among them. In particular, one has $L_1/L_0 \rightarrow -\frac{1}{2}$, and P_L/P_T decreases like τ^{-2} .
- An *interacting fixed point*, obtained for $\tau_R \sim \epsilon^{-1/4}$, where $g_0 \rightarrow -\frac{4}{3}$, $g_1 \rightarrow -2$. At this fixed point, the system has a locally isotropic particle distribution, and an energy density that decreases like $\tau^{-4/3}$.

Note that in a scaleless system where the collision rate scales as $\epsilon^{-1/4}$ (this is the conformal case mentioned above), the ratio τ/τ_R always goes to infinity as $\tau \rightarrow \infty$, and the system therefore always converges to the interacting fixed point. The figure 15 shows the evolution of the first three g_n 's for an ensemble of initial conditions, as a function of τ/τ_R . After a short transient evolution, all these evolutions coalesce on a universal “attractor” (indicated by the thick black curve), on which the exponents are the free streaming ones or the interacting ones depending on the value of the ratio τ/τ_R .

4 Hydrodynamical evolution; Late stages

The next, and in fact main, stage of the bulk evolution of the system formed in the collision of two heavy ions is a rather long period of relativistic hydrodynamical expansion. Hydrodynamics [108, 109, 110, 111, 112] is a coarse grained description whose starting point is the local conservation laws, for energy and momentum and for any other conserved charge (such as baryon number),

$$\partial_\mu T^{\mu\nu} = 0, \quad \partial_\mu J_B^\mu = 0, \quad \dots \quad (26)$$

However, these equations are too general to constrain uniquely the evolution of the system. For this, it is necessary to express the energy-momentum tensor and the other conserved currents in terms of a small number of quantities.

Perfect fluid The simplest case is that of a perfect, i.e., non dissipative, fluid. In this case, the energy-momentum tensor depends only on the local energy density ϵ , pressure p and fluid 4-velocity vector u^μ ,

$$\begin{aligned} T^{\mu\nu}_{\text{perfect}} &= (\epsilon + p)u^\mu u^\nu - p g^{\mu\nu} \\ &= \epsilon u^\mu u^\nu + p \Delta^{\mu\nu} \quad \text{with} \quad \Delta^{\mu\nu} \equiv u^\mu u^\nu - g^{\mu\nu}, \end{aligned} \quad (27)$$

where $g^{\mu\nu}$ is the Minkowski metric tensor with $(+, - - -)$ signature. The tensor $\Delta^{\mu\nu}$ is a projector on the local rest frame of the fluid, $\Delta^{\mu\nu} u_\nu = 0$, in terms of which one may define time and spatial derivatives, $D \equiv u^\mu \partial_\mu$ and $\nabla^\mu \equiv \Delta^{\mu\nu} \partial_\nu$, in the fluid rest frame. The equations of ideal hydrodynamics therefore read

$$D\epsilon = -(\epsilon + p)\nabla_\mu u^\mu, \quad Du^\mu = -(\epsilon + p)^{-1} \nabla^\mu p. \quad (28)$$

These equations are the relativistic analogue of Euler fluid equations. The first equation indicates that the local variation of the energy density is proportional to the variation of the volume of fluid cells, since $dV/V = dt \nabla_\mu u^\mu$. The second equation relates the acceleration of the fluid to the gradient of its pressure. Note that the first equation implies that entropy is conserved (indeed this equation is equivalent to $d(\epsilon V) + p dV = T dS = 0$).

Boost invariant ideal flow In heavy ion collisions at ultrarelativistic energies, the longitudinal momenta of the produced particles are typically much larger than their transverse momenta. This leads to a strong correlation between the longitudinal momentum of a particle and its longitudinal position in coordinate space. More precisely, one has

$$y \equiv \frac{1}{2} \ln \left(\frac{p^0 + p^3}{p^0 - p^3} \right) \approx \eta_s \equiv \frac{1}{2} \ln \left(\frac{x^0 + x^3}{x^0 - x^3} \right). \quad (29)$$

If the energy density of the fluid at some initial proper time τ_0 is independent of the spatial rapidity η_s (this is approximately the case in the Color Glass Condensate framework, since the rapidity dependence comes from the JIMWLK evolution of the distributions of color sources, which is significant only on scales $\delta\eta_s \sim \alpha_s^{-1}$), the subsequent hydrodynamical evolution of the fluid is boost invariant. In this case, the evolution of a perfect fluid is governed by a single equation,

$$\frac{d\epsilon}{d\tau} = -\frac{\epsilon + p}{\tau}. \quad (30)$$

(If the fluid is not locally isotropic, the pressure in the right hand side should be replaced by the longitudinal pressure.) In a conformal theory (i.e., a theory with only massless particles and no running coupling), the energy-momentum tensor is traceless and one has $p = \epsilon/3$ in equilibrium. Therefore, this leads to

$$\begin{aligned}\epsilon, p &\sim \tau^{-4/3}, \\ T &\sim \epsilon^{1/4} \sim \tau^{-1/3}, \\ s &\sim T^3 \sim \tau, \\ sV &\sim s\tau \sim \text{const},\end{aligned}\tag{31}$$

where s is the entropy density. Recall that the assumption of boost invariance, and therefore these scaling laws, are only true as long as the longitudinal expansion dominates over the transverse one (this is hidden in the assumption that $V \sim \tau$), and are therefore expected to change when the proper time becomes comparable to the diameter of the colliding nuclei. Note also an important fact, equally valid for solving Eq. (30) or the general hydrodynamical equations: in order to close the system of equations and obtain a solution, it is necessary to use an *equation of state* that relates for instance the pressure to the energy density.

Viscous corrections In order to go beyond the simple description in terms of a perfect fluid, one should first alter Eq. (27) by writing

$$T^{\mu\nu} = T_{\text{perfect}}^{\mu\nu} + \pi^{\mu\nu} + \Pi \Delta^{\mu\nu},\tag{32}$$

where we have split the deviation from the perfect fluid into a traceless tensor $\pi^{\mu\nu}$ and a term $\Pi\Delta^{\mu\nu}$ that has a non-zero trace. The equations of motion are $\partial_\mu T^{\mu\nu} = 0$, combined with an equation of state and constituent equations that express $\pi^{\mu\nu}$ and Π in terms of gradients. In a system which is not too far from local equilibrium, these expressions may be expanded in powers of the gradients², and at lowest order one may write

$$\begin{aligned}\pi^{\mu\nu} &= -\eta \sigma^{\mu\nu} \quad \text{with} \quad \sigma^{\mu\nu} \equiv \nabla^\mu u^\nu + \nabla^\nu u^\mu - \frac{2}{3} \Delta^{\mu\nu} (\nabla_\rho u^\rho), \\ \Pi &= -\zeta (\nabla_\rho u^\rho).\end{aligned}\tag{33}$$

The coefficients η and ζ (respectively, the shear and bulk viscosities) describe how the stress tensor responds to a small gradient of the fluid velocity. The resulting hydrodynamical equations are the relativistic analogue of the Navier-Stokes equations.

When applied to a boost invariant system, the resulting hydrodynamical equations lead to

$$\frac{d\epsilon}{d\tau} = -\frac{\epsilon + p - \frac{4}{3}\frac{\eta}{\tau}}{\tau},\tag{34}$$

resulting in a slower decrease of the energy density compared to the case of a perfect fluid. This equation indicates that the first order gradient expansion it was obtained from is legitimate as long as $\eta \ll \tau(\epsilon + p)$. Using the thermodynamic relation $\epsilon + p = sT$, this condition can be turned into

$$\frac{\eta}{s} \ll \tau T,\tag{35}$$

where the left hand side of the inequality is a local property of the fluid while the right hand side is a property of the flow itself. From kinetic theory, the ratio η/s may be estimated

²This expansion should be regarded as rather formal, as it may not lead to a convergent series [113].

to be of the order of λT where λ is the mean free path. Thus, the inequality also reads $\lambda \ll \tau$, implying that the system cannot be described by hydrodynamics at times that are smaller than the time between two successive scatterings of a particle. Conversely, the hydrodynamical description improves as τ increases. In a scale invariant system, we may estimate the relative magnitude of the first viscous correction in the right hand side of Eq. (34) as follows,

$$\frac{\eta}{\tau(\epsilon + p)} \sim \frac{T^3}{\tau T^4} \sim \tau^{-2/3}, \quad (36)$$

where in the last step we use the behavior of T from ideal hydrodynamics.

Causality, Second order hydrodynamics The first order gradient expansion in Eq. (33) leads to some pathologies in a relativistic context, because the correction to the stress tensor follows instantaneously any modification to the velocity field. This causality violating behavior eventually leads to numerical instabilities when solving the corresponding hydrodynamical equations. A possible practical strategy to fix this problem is to modify Eqs. (33) into relaxation equations, in order to introduce a delay between changes of the gradients and the resulting variation of the stress tensor. There is no unique way of doing this, but on timescales longer than the ad-hoc relaxation time, all these models lead to identical physical predictions.

At a more fundamental level, the modifications introduced by turning Eqs. (33) into relaxation equations can be motivated from the study of second order terms in the gradient expansion [114, 115]. For instance, the second gradient order in $\pi^{\mu\nu}$ contains a term of the form $\eta\tau_\pi D\sigma^{\mu\nu}$, where τ_π has the dimension of a time. If we take this second order expansion as is, it displays similar causality issues as the first order one. These problems may be avoided by replacing $\sigma^{\mu\nu}$ in the time derivatives that appear at second order by the first order relationship between $\sigma^{\mu\nu}$ and $\pi^{\mu\nu}$, thereby producing a term in $-\tau_\pi D\pi^{\mu\nu}$. By doing this, the second order constituent relation becomes a differential equation in time, with a relaxation time τ_π . The benefit of this point of view compared to the more phenomenological one described before is that it allows to relate the relaxation time to the underlying microscopic theory in an unambiguous fashion.

Equation of state A key ingredient in order to turn the equations of hydrodynamics into a closed set of equations is an equation of state that relates for instance the entropy density to the temperature, or equivalently the pressure and the energy density. Perturbation theory, improved by the resummation of hard thermal loops[117], allows to obtain robust results for temperatures only a couple of times above the critical temperature [118, 119, 120, 121, 122, 123, 124, 125], but becomes less and less reliable as T_c is approached from above. A non-perturbative first-principles alternative is lattice QCD [126]. At null baryon chemical potential, the advances in computing hardware and algorithms are by now allowing to perform unquenched (i.e., with virtual quark loops) simulations with realistic quark masses. At low temperatures compared to the deconfinement transition, these computations may be continued by an equation of state based on a gas of hadron resonances.

This approach works well only at zero baryon chemical potential, where it has shown that the deconfinement transition is a continuous crossover rather than a discontinuity for realistic quark masses. At $\mu_B > 0$, the determinant resulting from integrating out the fermion fields is complex valued, leading to a *sign problem* that precludes any direct approach based on a Monte-Carlo sampling. When μ_B/T is small enough, various workarounds are possible: reweighting, analytic continuation from calculations at imaginary μ_B (for which there is no sign problem), that allow to make an incursion into the territory of positive chemical




Figure 16: Top: equation of state obtained from lattice QCD at zero baryon chemical potential. Bottom: speed of sound as a function of temperature. Boxes: extraction from LHC data (see text). From [116].

potentials [127, 128]. It is expected that the crossover at small μ_B becomes a first order phase transition at larger μ_B , the beginning of the transition line being a second order critical point (the red dot in Figure 2). However, the quest in lattice simulations for such a second order critical point has remained rather inconclusive until now.

There has been a recent attempt to extract directly some information about the equation of state from LHC heavy ion data [116]. In this work, the authors used a hydrodynamical simulation to estimate the effective temperature T_{eff} and effective volume V_{eff} of a hypothetical homogeneous fluid that would evolve into a system with the same energy and entropy as the QGP at the time where the particles decouple (see the discussion of *freeze-out* later in this section). They observed that this effective temperature is related to the mean transverse momentum $\langle p_{\perp} \rangle$ of the final state particles by $T_{\text{eff}} \approx 3.07 \langle p_{\perp} \rangle$ (with a proportionality coefficient roughly independent of the equation of state and transport coefficients). Given ALICE data in the bin of centrality [0, 5]%, one gets $T_{\text{eff}} = 222 \pm 9$ MeV. The total entropy is inferred from the number of produced charged particles, $S \approx 6.7 N_{\text{ch}}$, while the effective volume V_{eff} also comes from the hydrodynamical simulation, giving an entropy density $s = 20 \pm 5 \text{ fm}^{-3}$, and $s/T_{\text{eff}}^3 = 14 \pm 3.5$, in agreement with the lattice equation of state. In particular, this is much higher than the value $\sim 3\text{-}4$ expected in the confined phase, suggesting a large number of degrees of freedom consistent with deconfinement. By repeating

this analysis at two collision energies (2.76 and 5.02 TeV), one may also estimate the speed of sound via

$$c_s^2(T_{\text{eff}}) = \left. \frac{sdT}{Tds} \right|_{T_{\text{eff}}} = \frac{d \ln \langle p_{\perp} \rangle}{d \ln(dN_{\text{ch}}/d\eta)} = 0.24 \pm 0.04,$$

again in agreement with lattice computations.

Transport coefficients Among the transport coefficients that enter in hydrodynamics, the one that has received most interest is the shear viscosity η . As we mentioned earlier, the ratio η/s is the ratio of the mean free path to the quantum wavelength of the particles. This allows to make some simple estimates in several limits. Firstly, in the perturbative limit (weak coupling, and low enough particle density), this leads to [129, 130, 131, 132]

$$\frac{\eta}{s} \sim \frac{1}{\alpha_s^2 \ln(\alpha_s^{-1})} \gg 1. \quad (37)$$

(Although we do not write the prefactor here, it can be determined in the weak coupling limit.) Another limit is that of a strongly coupled plasma. Although this limit is not accessible in QCD, this calculation is possible in a supersymmetric cousin of QCD thanks to the AdS/CFT correspondence, leading the following result [133]

$$\frac{\eta}{s} = \frac{1}{4\pi}. \quad (38)$$

Note that such a constant value, independent of the coupling, is consistent with the fact that quantum mechanics prevents this ratio from becoming arbitrarily small since the quantum wavelength is a lower value for the mean free path (but such an argument does not give the value of the constant ratio one would reach).

Out of equilibrium, there is another interesting situation, where the coupling constant is weak but the gluon occupation number is large, possibly as large as α_s^{-1} . In this case, the scattering rate should contain a factor $f(1+f)$ where f is the occupation number of the scattering centers (when the occupation number is small, only a factor f is necessary, leading to the usual formula $\lambda^{-1} \sim n\sigma$). Although the scattering cross-section is proportional to α_s^2 (up to logarithms), the factor $f(1+f)$ leads to a mean free path that does not contain any power of the coupling constant. This situation of weak coupling but high density, which is relevant in the very early stages of a heavy ion collision, illustrates the fact that a strong coupling is not the only possibility for having a small ratio η/s : more generally, the system should be *strongly interacting* for this to be true.

Outside of the above limiting cases, one may also consider lattice QCD. Unlike the equation of state, the transport coefficients remain very difficult to even estimate in lattice QCD. Thanks to Green-Kubo's formulas, transport coefficients may be expressed in terms of the Fourier transform of a retarded current-current (the current should be the one that couples to the quantity whose transport one is interested in, e.g., a charge current for an electrical conductivity) correlation function at zero momentum,

$$\sigma \propto \lim_{\omega \rightarrow 0} \frac{\rho(\omega, \mathbf{k} = 0)}{\omega} \quad (39)$$

This formula expresses the transport coefficient in terms of the slope at zero energy of the corresponding spectral function $\rho(\omega, \mathbf{k})$. However, on the lattice, one has only a direct access to the imaginary time version of these correlation functions. This imaginary time correlator admits a spectral representation involving the relevant spectral function,

$$G(\tau, \mathbf{k}) = \int d\omega K_{\tau}(\tau, \omega) \rho(\omega, \mathbf{k}), \quad (40)$$

where $K_T(\tau, \omega)$ is a known temperature-dependent kernel. In a lattice approach, the left hand side of this equation would be obtained (with statistical errors, and a finite number of values of τ and \mathbf{k}) from numerical simulations, and one would then try to invert this integral relationship to obtain the spectral function $\rho(\omega, \mathbf{k})$. The difficulty with this is that it is a very ill-posed problem when the left hand side is only imperfectly known, implying that a direct inversion is unfeasible. Attempts at constraining the spectral function in this way have been made by using Bayes theorem in order to find the most likely spectral function compatible with Eq. (40) and with a set of prior assumptions about its shape (a minimal assumption would be that it is positive) [134, 135, 136, 137, 138].

More recently, a hybrid approach was proposed [139], that combines a skeleton loop expansion for the correlator $\langle[\pi_{ij}(x), \pi_{ij}(0)]\rangle$ that enters in Green-Kubo's formula, truncated at two loops, and a non-perturbative input for the couplings and the propagator that enters in this expansion obtained from the Euclidean functional renormalization group [140]. The




Figure 17: Theoretical evaluations of the viscosity to entropy ratio as a function of temperature. From [139].

results obtained with this approach are shown in Fig. 17, for a pure Yang-Mills theory (the result shown for QCD with quarks is an estimate based on the pure YM result, not an ab initio calculation). Interestingly, this calculation suggests that η/s has a rather pronounced temperature dependence, with a minimum around $1.25 T_c$ (the location of this minimum in units of the critical temperature is almost identical in Yang-Mills theory and in QCD). Moreover, the value of η/s at this minimum is only slightly above the value $1/(4\pi)$ obtained in the strong coupling limit in the AdS/CFT approach for $N = 4$ super-Yang-Mills theory.

Freeze-out Since at the end of the day experiments observe particles, it is necessary to convert the objects (the fluid energy density and its velocity) evolved by the hydrodynamical equations into distributions of particles. In fact, the need for a description in terms of particles arises much earlier than the time at which particles are detected because the density decreases as the system expands, and therefore the mean free path increases, leading to a situation where the conditions of applicability of hydrodynamics are not met anymore.

One should distinguish between a *chemical freeze-out*, at which the inelastic collisions become rare (after that, the chemical composition of the system is frozen), and a *kinetic freeze-out* where the elastic collisions also stop and the momentum distributions freeze. Experimentally, the temperature and chemical potential at the chemical freeze-out are well constrained from the ratios of abundances of various species of particles, as shown in the figure 18 [141, 142, 143, 144] (see also [145, 146] for a somewhat related approach to hadron

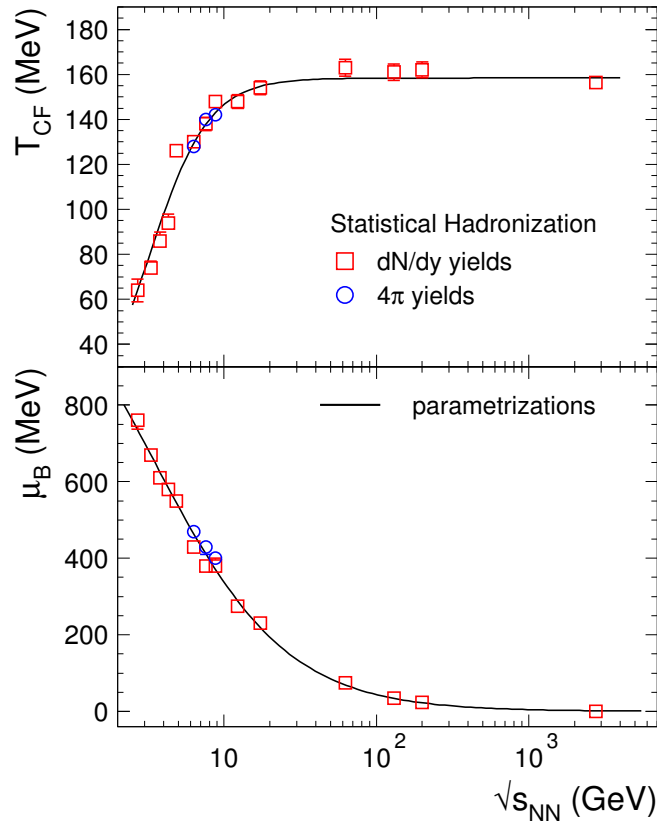



Figure 18: Top: measured hadron yields compared to the statistical hadronization model output. Bottom: temperature and chemical potential at chemical freeze-out, extracted from collisions at various energies. From [144].

formation). In the simplest version of this model, one assumes that all particle distributions are the equilibrium ones at a common temperature and chemical potential. A volume is also necessary to obtain absolute yields, but this parameter drops out in ratios of the yields of various particle species.

The conversion from a fluid to free particles should be done at the kinetic freeze-out, by choosing a locally space-like “surface” Σ , and by using the *Cooper-Frye formula*,

$$E_{\mathbf{p}} \frac{dN}{d^3\mathbf{p}} = \frac{1}{(2\pi)^3} \int_{\Sigma} d^3S_{\mu} P^{\mu} f(P \cdot u), \quad (41)$$

where f is the local distribution function. For a fluid in local thermal equilibrium, f is the Bose-Einstein or the Fermi-Dirac distribution, evaluated at the local fluid temperature. However, when the fluid is viscous, there should also be deviations from local thermal distributions,

$$f(p) = f_{\text{eq}}(p) + \delta f(p). \quad (42)$$

The form of the deviation $\delta f(p)$ is related to the transport coefficients (such as the shear viscosity) and therefore depends on the microscopic interactions in the fluid.

In this approach, the freeze-out conditions (temperature and chemical potential) are a priori free parameters that may be adjusted to best fit the spectra of produced particles. A more sophisticated alternative would be to convert the fluid into particles at an earlier time, and continue the evolution with kinetic equations [147, 148, 149]. In this fancier approach, the freeze-out temperature is controlled by the values of the various cross-sections used in the kinetic description (this would even allow different species to decouple at different times), and is no longer an ad hoc input of the model.

Flow anisotropies Experimentally, several predictions of hydrodynamical models may be compared with data. One of them is the transverse momentum spectra of the produced particles, that are sensitive to the temperature at which the freeze-out occurs (and to whether the particle distributions are the equilibrium ones or not).

Another main class of observables directly related to the hydrodynamical expansion of the quark gluon plasma consists in measuring angular correlations among the detected particles [150, 151, 152, 153, 154, 155, 156, 157, 158, 159]. The quantities used to express these correlations are the so-called v_n , defined as the Fourier coefficients of the azimuthal distribution of particles. These flow coefficients may be measured as a function of transverse momentum, of the centrality of the collisions, of the species of particles. The hydrodynamical expansion provides a one-to-one mapping between the spatial anisotropy of the initial distribution of energy density (and possibly its initial flow) and the final momentum anisotropy, that may be understood as an effect of pressure gradients (the fluid is accelerated in the direction of the pressure gradient). Moreover, the conversion of spatial inhomogeneities into momentum space anisotropies depends quite sensitively on the transport coefficients, mostly the shear viscosity: if the ratio η/s is too large, this conversion is very ineffective and hydrodynamics cannot explain the rather large values of the measured v_n s.

Recently, the JETSCAPE collaboration has used a Bayesian approach in order to extract the most likely values of the shear and bulk viscosities as a function of temperature [161, 162], as shown in Figure 20. One can see that the constraints provided by the data restricts the range of values of the viscosities, down to fairly low values after normalization by the entropy density, compared to the prior distribution used in this analysis. Although these results are very promising, one should keep in mind that, besides the transport coefficients one wants to extract, the model has other unknowns, especially in the modelling of the initial condition and in the details of the implementation of the freeze-out.




Figure 19: Second Fourier coefficient of the azimuthal distribution for identified particles. From [160].

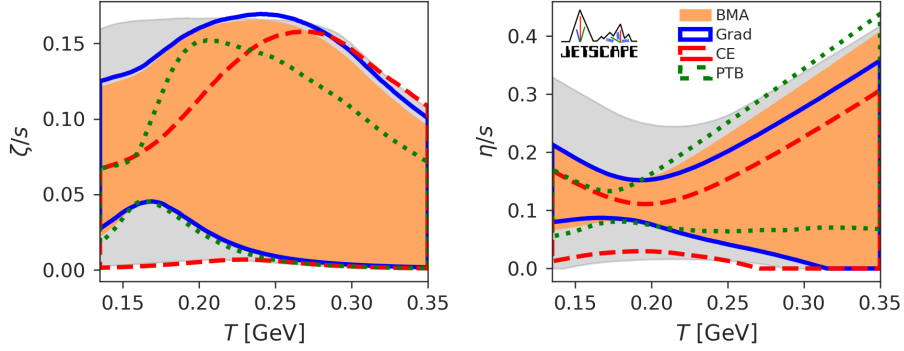


Figure 20: Bayesian extraction of the shear and bulk viscosities from Pb-Pb collisions at 2.76 TeV. Gray: 90% confidence interval for the prior. Red, Green, Blue, Orange: 90% confidence intervals for the posteriors in three models, and their average. From [161].

Hydrodynamics from kinetic theory In local thermal equilibrium, the hydrodynamical description has six unknowns (the local baryon density n_B , energy density ϵ and pressure P , plus the three independent components of the fluid velocity u^μ), and the conservation equations for the baryonic current and the energy-momentum tensor provide five equations. The system is closed by the equation of state that relates n_B, ϵ and P .

Away from this ideal situation, there are nine additional unknowns (II, five independent components of $\pi^{\mu\nu}$, and three components of the mismatch between the energy flow vector and the baryon number flow vector). In addition to the equation of state, we thus need 14 equations to close the system, that are made by the conservation equations (5 equations), and by 9 constitutive equations relating the stresses to the gradients. But these equations can also be obtained from the underlying microscopic dynamics, in the form of a Boltzmann equation $p^\mu \partial_\mu f = C_p[f]$. Indeed, by weighting the Boltzmann equation by 1, p^ν , $p^\nu p^\lambda$ and

integrating over \mathbf{p} , we get

$$\partial_\mu \int_{\mathbf{p}} p^\mu f = \int_{\mathbf{p}} C_{\mathbf{p}}[f] = 0, \quad (43)$$

$$\partial_\mu \int_{\mathbf{p}} p^\mu p^\nu f = \int_{\mathbf{p}} p^\nu C_{\mathbf{p}}[f] = 0, \quad (44)$$

$$\partial_\mu \int_{\mathbf{p}} p^\mu p^\nu p^\lambda f = \int_{\mathbf{p}} p^\nu p^\lambda C_{\mathbf{p}}[f]. \quad (45)$$

The right hand side of the first two equations is identically zero given the symmetry properties of the collision integral, for any distribution f . This set of 5 equations is in fact the five conservation equations. The last equation forms a set of 10 independent equations (given the symmetry of $p^\nu p^\lambda$). Note that by summing over $\nu = \lambda$, one recovers the conservation equation of particle number. Therefore, this set of equations contains only 9 novel equations, i.e., precisely the number needed to close our macroscopic description. This approach, supplemented by an expansion of f around the equilibrium distribution f_0 , is known as Grad's 14-moment method. It allows to relate the hydrodynamical description (in particular the transport coefficients [115, 163, 164, 165, 166, 167]) to the underlying microscopic interactions encoded in the collision integral.

Hydrodynamics far from equilibrium For a long time, it has been thought that the hydrodynamical description is a near-equilibrium effective description (this point of view stems in part from the fact that one may obtain hydrodynamics from a truncated gradient expansion). However, it was realized recently that this gradient expansion may be a non-convergent series with a null radius of convergence [168, 113]. This observation suggests to reconsider the conditions of applicability of the hydrodynamical description (indeed, if the radius of convergence is zero, it cannot serve as a parameter that defines what “close enough to equilibrium” means). Another, observational, reason for reassessing the applicability of hydrodynamics as an effective macroscopic description is that in heavy ion collisions it appears to perform much better than what one may naively expect by viewing it as the result of an expansion around equilibrium.

A step towards a better understanding of these questions has been the discovery (so far, in simple cases, where the dynamics and the flow are sufficiently symmetric) of attractors towards which hydrodynamical solutions evolve, even while gradients are still sizeable [169, 170, 171, 172]. In other words, these solutions quickly reach a universal behavior independent of the details of their initial conditions, long before the system is in a state of isotropic local equilibrium. On these attractors, the dissipative currents behave in a universal way in terms of the (possibly still large) gradients, which shed another light on the constitutive relations, indicating that their validity may not be conditioned by a gradient expansion.

Hydrodynamics in small systems Another pressing question, closely related to the previous point, is to determine what is the “minimal size” of a system that may be described by hydrodynamics. Indeed, as the system becomes smaller, the collectivity (that may be quantified by the ratio of the system size by the mean free path, a gross measure of the number of collisions per particle) is expected to decrease, leading to a situation less favorable for the applicability of hydrodynamics.

Experimentally, this question was put forward by the observation that certain correlation patterns (e.g., the “ridge” long range rapidity correlation observed in the two-particle spectrum) are seen in nucleus-nucleus collisions [173], proton-nucleus [30] and even proton-proton collisions [34] (provided one triggers on high multiplicity final states in the latter

cases). In nucleus-nucleus collisions, the ridge is interpreted as the result of the collective radial motion of the produced matter [174, 175, 176], i.e., as flow. Moreover, in this case, there is only a modest contribution to the correlation provided by the CGC initial condition.

In smaller systems (see Figure 21 for some recent data), the interpretation of the ridge correlation has been the subject of intense debates between two extreme positions: that the correlation can be entirely explained by collective flow [177, 178] (i.e., final state interactions), or entirely due to initial state correlations [179, 180, 181, 182]. By now, it seems that a consistent description of the flow observables in systems for which $dN_{\text{ch}}/d\eta \geq 10$ calls for a dominance of the final state interactions over the initial state correlations. At low multiplicity, the observed flow is probably the result of the combination of initial state correlations and final state interactions [183] (it has been observed in a realistic hydrodynamical model that the final state flow is less correlated with the initial geometry but more correlated with the initial momentum anisotropy at low multiplicity, and that this trend is reversed at high multiplicity).




Figure 21: Points: v_2 and v_3 flow coefficients in collisions of nuclei of various sizes, down to proton-proton collisions (ALICE collaboration). Bands: hydrodynamical model predictions. From [184]. See also [185, 186].

5 Tomography and hard probes

Until now, we have discussed mostly the bulk properties of the matter produced in heavy ion collisions, from its formation to its subsequent evolution. There, the quark-gluon-plasma or its precursor was the core subject of the study, with flow observables a tool for accessing (almost) directly some of their properties. Now, we turn to a class of observables in which this matter plays the role of a substrate that modifies them compared to a situation where the final state would be mostly empty. Here, the strategy is to measure these observables both in nucleus-nucleus collisions and in collisions of smaller projectiles –such as proton-proton collisions– where the formation of a quark-gluon plasma is not expected, and where the observables are supposedly well understood. The comparison between the two (with an appropriate rescaling to account for the different sizes of the two systems) provides informations on the properties of the matter produced in heavy ion collisions. A better control is provided by observables that are perturbative (meaning that they are characterized by a hard scale, allowing to use QCD in a regime of weak coupling), so that they can at least be computed with some degree of accuracy in proton-proton collisions. However, even

in this case, one should keep in mind that their modifications by the surrounding medium may involve much smaller momentum scales that render a perturbative treatment difficult or even impossible in nucleus-nucleus collisions.

Basics of QCD radiation The most important feature of QCD in this context is the fact that gluons are massless, and that three flavors of quarks (u,d, and s) have masses lower or comparable to the non-perturbative scale Λ_{QCD} . For this reason, gluon radiation is enhanced for soft gluons and for gluons emitted collinearly to their parent. For instance, the emission probability of a gluon off a parton goes like

$$dP \propto \alpha_s \frac{d^2 \mathbf{k}_\perp}{k_\perp^2} \frac{dz}{z}, \quad (46)$$

where \mathbf{k}_\perp is the transverse momentum of the gluon (relative to the direction of the momentum of its parent) and z is the fraction of its longitudinal momentum relative to that of its parent. The unwritten prefactor contains a color factor that depends on whether the emitter is a quark or a gluon. This probability diverges at small k_\perp and at small z . For radiation in vacuum (e.g., in the final state of a proton-proton collision), the consequences of these divergences are well understood:

- After one resums the soft radiation, the probabilities for partonic final states with a prescribed number of (massless) gluons is actually zero. This is of course largely irrelevant, since because of confinement the gluons are not the objects that are eventually detected.
- The observables that match the most closely the concept of perturbative parton are the so-called *jets*. Loosely speaking, a jet is a collimated beam of particles produced by the splitting of a common ancestor (quark or gluon). At a more operational level, defining jet cross-sections requires that one defines a procedure for deciding when two distinct partonic final states correspond to the same configuration of jets. For instance, two graphs that differ by a loop correction, or by an extra soft or collinear gluon, should contribute to the same jet final state. The jet definition is not unique (in particular, because there is no unique way of defining “soft” or “collinear”), but they all share a crucial property: they lead to finite cross-sections when all mass scales are sent to zero (observables that have this property are said to be *infrared and collinear safe*) – in other words, all the logarithms that one would get by integrating eq. (46) down to $k_\perp = 0$ or $z = 0$ cancel in these observables. The theoretical definition of what one means by a jet in turn defines the procedure (the so-called “jet algorithm”) for extracting jet cross-sections from experimental data: given a final state made of detected particles, the jet algorithm defines how they should be clustered into jets.
- Jet cross-sections are calculable without any non-perturbative input, but the price to pay for this is to give up on the idea of saying something about individual particles in the final state. It is also possible to consider cross-sections for producing a certain hadron with a given momentum, but their calculation requires to introduce *fragmentation functions*, that may be viewed as the inclusive probability that a certain quark or gluon turns into this hadron (plus any number of additional particles that we do not care about). This time, the collinear logarithms from eq. (46) do not cancel and they must be resummed, which introduces a scale dependence into the fragmentation function. This scale dependence is perturbative (governed by the DGLAP evolution equation [187, 188, 189], now known up to three-loop accuracy [190, 191, 192]), but the initial condition of this evolution is non-perturbative and must be extracted from

experimental input. The hadron production cross-section is obtained as the convolution of a partonic cross-section with the fragmentation function, evaluated at a certain scale (called the *factorization scale*). This scale is not a physical parameter, but rather a remnant of the truncation of the perturbative series at a finite order (the residual scale dependence decreases by going to higher loop order, and a fully non-perturbative calculation would have no such scale at all).

Production of hard probes in nucleus-nucleus collisions When extending these ideas to nucleus-nucleus collisions, a generic assumption is that the *production* of the hard object proceeds in the same perturbative way as in proton-proton collisions, the only change being a change of the parton distribution functions that describe the initial state. In the light of the earlier discussion of gluon saturation, it is clear that for this to be true the observable of interest should probe these distributions away from the non-linear saturation regime, in order to be dominated by processes that probe a single parton in each projectile. This is a reasonable assumption at high virtuality Q^2 , since the non-linear corrections are typically suppressed as powers of Q_s^2/Q^2 . Another thing to keep in mind is that the direct measurements of *nuclear parton distributions* are scarce. An approximate treatment consists in viewing a nucleus as an incoherent superposition of protons and neutrons. The parton distributions of protons are very well known. For the neutrons, one usually treats the neutron (udd) as the isospin partner of the proton (uud). This implies for instance that the u-quark distribution in a neutron is the same as the d-quark distribution in a proton, etc... But note that the measurement of the structure function F_2 in deep inelastic scattering off a proton does not allow to disentangle the u and d quark distributions. To separate them, one also needs DIS measurements with deuteron (plus the assumption that the binding of the deuteron is weak enough compared to the relevant virtuality scales so that its parton distributions are just the sum of the proton and neutron ones).

Gluon formation time The discussion of parton and jet energy loss requires to introduce the concept of *formation time* of a radiated gluon. Consider the emission of a gluon of momentum k off a colored particle (quark or gluon) of momentum $p + k$. According to the uncertainty principle, the virtuality of the line of momentum $p + k$ just before the gluon emission vertex defines the gluon “formation time”

$$t_f^{-1} \equiv E_{\mathbf{p}} + E_{\mathbf{k}} - E_{\mathbf{p}+\mathbf{k}} \underset{k \ll p}{\approx} \frac{(p+k)^2}{2E_{\mathbf{p}}} = E_k(1 - \cos \theta) \underset{\theta \ll 1}{\approx} \frac{1}{2} E_k \theta^2. \quad (47)$$

This time can also be interpreted as the time necessary for the wave-packet of the new gluon to separate sufficiently from that of the emitter. Indeed, the transverse separation between the emitter and the gluon grows with time according to $\Delta r_{\perp} \approx \theta t$. On the other hand, the wavelength of the gluon, projected on the plane orthogonal to the emitter, reads

$$\lambda_{\perp} = \frac{1}{E_{\mathbf{k}} \sin \theta} \underset{\theta \ll 1}{\approx} \frac{1}{E_{\mathbf{k}} \theta}. \quad (48)$$

We see that the formation time is also given by the condition $\Delta r_{\perp} \Big|_{t_f} = \lambda_{\perp}$.

Parton energy loss The simplest of the observables sensitive to the medium modification of parton splitting consist in comparing inclusive hadron spectra measured in nucleus-nucleus collisions with the same spectra measured in proton-proton collisions at the same energy. However, a direct comparison of these two spectra is no meaningful because these yields come

from system of vastly different volumes. In order to account for this trivial geometrical effect, it is customary to define the so-called “nuclear modification factor”,

$$R_{AA} \equiv \frac{\frac{dN}{d^3\mathbf{p}} \Big|_{AA}}{N_{\text{coll}} \times \frac{dN}{d^3\mathbf{p}} \Big|_{pp}}, \quad (49)$$

where N_{coll} is the number of binary nucleon-nucleon collisions. For hard objects whose production is obtained as an incoherent sum of nucleon-nucleon collisions, scaling by this factor effectively eliminates the trivial volume dependence (in other words, the ratio R_{AA} should be equal to one in the absence of final state medium effects). Note that N_{coll} is not a directly measurable quantity, since its determination requires a modelling of the nucleon distribution inside the nucleus of interest. The ratio R_{AA} , called the nuclear modification factor, is displayed in the case of charged hadrons in the figure 22, for central lead-lead collisions at the LHC. The most obvious feature of the ratio R_{AA} is that it is below unity




Figure 22: Suppression factor of charged hadrons in nucleus-nucleon collisions. See [193, 194].

over a large momentum range, and slowly reaches one at large momenta. This means that the spectra of charged hadrons (at high momentum compared to the QCD scale Λ_{QCD}) are depleted in nucleus-nucleon collisions compared to proton-proton collisions, which can be explained by an increased fragmentation in the presence of a surrounding medium. Following Baier-Dokshitzer-Mueller-Peigné-Schiff [195, 196, 197] and Zakharov [198, 199, 200, 201, 202], this medium-induced radiation can be understood semi-quantitatively as follows. The formation time of a gluon of energy $E_{\mathbf{k}}$ and transverse momentum k_{\perp} is $t_f \approx E_{\mathbf{k}}/k_{\perp}^2$ (this was derived in eq. (47)). The main difference in the presence of a medium is that the transverse momentum of the emitted gluon is altered by the scatterings that occur within its formation time. Assuming that these scatterings act as random and independent kicks, the increase of the transverse momentum is a diffusion process in transverse momentum space. Therefore, the transverse momentum accumulated during the time t_f is given by $k_{\perp}^2 \approx \hat{q} t_f$, where \hat{q} is a constant that quantifies the mean Δk_{\perp}^2 per unit length (it is a quantity that depends on the properties of the medium: density/temperature, Debye screening). Consistency between these two relations implies that $t_f \approx \sqrt{E_{\mathbf{k}}/\hat{q}}$. Based on this, the induced emissions can be divided into three regimes, depending on the energy of the radiated gluon

- $E_{\mathbf{k}} < \omega_{BH}$, with $\omega_{BH} \equiv \hat{q}\lambda^2$ (λ is the mean free path of the emitter in the medium). In this regime, the radiation occurs coherently over path lengths shorter than the mean

free path. Successive collisions, separated by λ (therefore there are L/λ of them), contribute incoherently to the total radiation spectrum. In this regime, known as the Bethe-Heitler regime, the emission spectrum behaves as $dI/dE \sim E^{-1}$.

- $\omega_{BH} < E_{\mathbf{k}} < \omega_c$, with $\omega_c \equiv \hat{q}L^2$ (L is the length traveled by the emitter before exiting the medium). In this energy range, the formation time is larger than the mean free path, and shorter than the medium size. Several scatterings must happen in order to produce one emission, which reduces the total yield compared to the Bethe-Heitler regime (this relative suppression is known as the Landau-Pomeranchuk-Migdal effect). In this regime, the emission spectrum behaves as

$$\frac{dI}{dE} \approx \underbrace{\frac{dI}{dE}}_{\alpha_s E^{-1}} \Big|_{\text{single}} \times \underbrace{\frac{L}{t_f}}_{\sqrt{\omega_c/E}} \sim E^{-3/2}. \quad (50)$$

(For $E_{\mathbf{k}} < \omega_{BH}$, the factor L/t_f was replaced by the constant L/λ .)

- $\omega_c < E_{\mathbf{k}}$. In this regime, the formation time of the gluon is larger than the size of the medium, and induced radiation is effectively suppressed. One may thus view ω_c as an upper limit of the energy of a radiated gluon (for instance, for $\hat{q} = 2 \text{ GeV/fm}^2$ and $L = 5 \text{ fm}$, this cutoff is $\omega_c = 50 \text{ GeV}$). By combining the emission spectra in the three regimes, we can obtain the mean radiated energy by a parton: $\langle E_{\mathbf{k}} \rangle \propto \alpha_s \omega_c$. In other words, energy loss is most often due to a single emission (occurring with probability $\sim \alpha_s$) close to the upper limit. Note that a more dilute (or cooler) medium has a smaller \hat{q} and therefore a smaller cutoff energy ω_c , and a smaller mean energy loss.

Note that two approaches have been used in the literature for implementing the scattering-induced radiation of a hard parton. One option is to perform an opacity expansion, starting from a dilute medium [203, 204, 205, 206, 207]. In this approach, a very small number of scatterings are considered, but treated with exact kinematics. Another option is to resum multiple scatterings, in an approximation where the momentum transfer in these scatterings is soft compared to the parton momentum, more in line with the BDMPS-Z approach (see for instance [208, 209]).

Concrete implementations in the context of heavy ion collisions are of course more complex than this qualitative discussion suggests. Firstly, one needs a realistic modeling of the medium and its evolution, usually taken from hydrodynamics. Another source of difficulty lies in the fact that the momentum scales characteristic of the medium (for instance, the Debye screening mass) can be rather soft at the temperatures reached in heavy ion collisions at present energies, and in particular not large enough to confidently apply perturbative QCD at those scales. This has motivated hybrid approaches [211] that combine perturbative techniques for the hard scales and holography-inspired inputs for the softer medium scales (note that a pure strong coupling approach leads to an energy loss proportional to L^3 instead of L^2 , and is ruled out by data).

Let us close this subsection by mentioning that heavy quarks such as charm are also suppressed in heavy ion collisions, as one can see on the figure 23. One can see that D mesons are as suppressed as the charged pions, suggesting that the fragmentation of c quarks (at the temperature scales relevant in the LHC experimental conditions) is comparable to that of light quarks and gluons. The quark mass starts playing a visible role for b quarks, since the medium suppression of mesons containing b quarks is significantly less important. This is consistent with the *dead cone effect* [212], a kinematical effect that prevents gluon radiation inside a cone of opening m/E centered on the emitting quark of mass m . Note that, since the radiative losses are suppressed by the dead cone effect for heavy quarks, their energy loss through elastic scatterings (collisional energy loss) become relatively more important and must be included.




Figure 23: Open charm and non-prompt J/Ψ nuclear suppression factor, compared to that of pions. From [210].

Vacuum antenna pattern Until this point, we have discussed the medium modification of the spectra of single hadrons. Another possibility is to consider similar observables for jets instead of individual hadrons. One advantage of jets is that jet cross-sections are in principle perturbative since they do not rely on the details of the hadronization process. Moreover, jets provide another handle to probe the loss of energy due to the surrounding medium, since their opening angle may be chosen at will in their definition.

Before we discuss jet modifications in heavy ion collisions, let us recall the main feature of gluon radiation from a jet in vacuum. Compared to radiation by a single parton, the radiation from a jet is obtained by first considering a parton splitting process $a \rightarrow bc$ (where a, b, c could be quarks, antiquarks, gluons, photons). Since the partons b, c are produced from a common ancestor a , their colors are correlated, and so is the soft radiation they produce. In particular, the emission of an additional soft gluon after the splitting $a \rightarrow bc$ depends crucially on whether the new gluon is emitted inside or outside of the cone formed by the partons b and c . This effect can be understood semi-quantitatively as follows in the case where the angular opening θ_{bc} between b and c is small. As we have seen earlier, the formation time of an additional gluon of energy $E_{\mathbf{k}}$ at an angle θ is given by $t_f^{-1} \approx E_{\mathbf{k}}\theta^2$. By this time, the pair bc has grown to a transverse size

$$r_{\perp} = t_f \theta_{bc}. \quad (51)$$

Moreover, the gluon wavelength ($E_{\mathbf{k}}^{-1}$), projected on the plane orthogonal to the pair momentum, is $\lambda_{\perp} \approx (E_{\mathbf{k}}\theta)^{-1}$. When the transverse wavelength of the emitted gluon is smaller than the size of the pair, the gluon resolves the individual constituents b and c of the pair and the emission is the sum of the separate emissions from b and c . Otherwise, the gluon cannot resolve the pair, and sees only its total charge, which is the charge of the parent a . This condition reads

$$\frac{1}{E_{\mathbf{k}}\theta} \lesssim r_{\perp} \Leftrightarrow \theta \lesssim \theta_{bc}. \quad (52)$$

In other words, inside the cone, the emission is the incoherent sum of the emissions from b and c , while the emission outside the cone is equal to that of the total charge, i.e., a . In

particular, if the ancestor is a color neutral object (a virtual photon), then there is no soft radiation outside the cone of the pair. This property is known as *angular ordering* in the context of QCD. It can be used in order to formulate gluon emission in the form of “parton cascades” in which each quark, antiquark or gluon has a certain probability of emitting a new gluon. The effect of quantum interferences, that would in principle preclude such a probabilistic description, is taken into account by vetoing emissions at large angles.

In-medium jet energy loss A striking observation regarding jets is that there are events with very imbalanced pairs of jets in heavy ion collisions, as shown in the figure 24. In proton-




Figure 24: Event display of a very asymmetrical two-jet event. From [35].

proton collisions, events with two jets in the final states have jets with nearly balanced transverse momenta, as imposed by momentum conservation. In the above event display, one jet has almost three times the energy of the second jet (or, put differently, about 130 GeV appear to have been “lost”), and it is not possible to see directly on this histogram where the energy has gone (the only conclusion one may draw by eye is that this energy has been degraded in the form of soft enough fragments that are lost in the background of this calorimetric display).

A first medium effect affecting jets is the fact that each parton in the jet may suffer from medium-induced radiation, in the way discussed earlier for single particle spectra. The only modification imposed by the jet definition is that only the radiation that goes outside of the jet cone must be counted as a loss (thus, this effect is more pronounced for narrow jets than for wider jets). We have seen in the previous part that the mean radiative energy loss of a parton is due to emissions close to the maximum $\omega_c = \hat{q}L^2$. For those, the typical gluon emission angle is $\theta \approx k_{\perp}/E_k \approx (L\omega_c)^{-1/2} \ll 1$. The consequence of this observation is that these rather hard induced emissions do not alter the energy of a jet for typical jet opening angles (usually $0.3 \leq R \leq 0.8$), since the radiated gluon stays within this cone (the jet substructure is altered, but this is not picked by the jet reconstruction algorithm). Therefore, the observed asymmetry between the energy of the pair of jets can be explained by softer gluons emitted at large angles.

Another important effect affecting a jet as it propagates through a medium is that the multiple scatterings of the constituents of a pair of partons eventually lead to the loss of their color coherence [213, 214, 215, 216, 217, 218, 219]. Thanks to this decoherence, the vacuum-like emissions are no longer forbidden outside of the jet cone. This happens when the two partons scatter off external color fields that are uncorrelated, i.e., when the transverse separation r_{\perp} between them is larger than the coherence length of the color field. Since r_{\perp} is proportional to the opening angle of the pair, the decoherence time is shorter for a pair with a larger opening angle. The conclusion of this qualitative argument is that small jets are more robust against in-medium vacuum-like energy loss than wider jets.

Moreover, it has been shown that the energy emitted outside of the jet cone is rapidly degraded into partons whose typical energy is close to the temperature of the surrounding medium [220, 221, 215, 222]. Experimentally, this has been qualitatively confirmed, since one recovers the missing jet energy in the form of many softer particles outside of the jet cone.

γ, W^\pm, Z -jet correlations Because of momentum conservation, jets are dominantly produced in pairs in hard processes, and more rarely (with a suppression of order α_s) in a 3-jet configuration. After their initial production, the fate of these jets in the surrounding medium depends crucially on the location of the production point with respect to the bulk, and on the direction of motion of the jets. Loosely speaking, a jet that goes inward loses more energy than a jet that travels outwards, because it must travel a longer route through the medium. Except when the production point is very close to the outer boundary of the medium, both jets lose energy to some degree, making it difficult to infer event-by-event how much energy was lost by each of them.

A more direct access to the jet energy loss is possible in situations where a single jet is produced in conjunction with another object that interacts only via electroweak interactions, such as a photon, a W^\pm or a Z^0 boson (for instance, in a process such as $qg \rightarrow q\gamma$). These events are less frequent because of the electroweak coupling involved at the production vertex, but they offer the advantage that the weakly interacting object can escape from the medium without further interactions. Thus, its measurement provides an unaltered reference for the initial energy of the partner jet.

6 Thermometric probes

In this last section, we consider observables that are sensitive to the local temperature of the medium. In principle, these quantities could tell if the temperature reached in a heavy ion collision is above the deconfinement temperature. The actual situation is of course a bit more complicated, since the temperature of the medium is not spatially homogeneous, and because the measured quantities result from the entire history of the system, through which the temperature is not constant.

6.1 Electromagnetic radiation from the quark-gluon plasma

General considerations A first quantity which is quite sensitive to the plasma temperature is the spectrum of photons emitted by the plasma. Let us clarify here a possible paradox: in a large medium (larger than the photon mean free path), the electromagnetic radiation would be in thermal equilibrium with the quarks and the gluons, with a spectrum given by a Bose-Einstein distribution at the local temperature (but only photons emitted within one mean free path of the surface would escape). However, this is not the case in heavy ion collisions, where the size of the medium is considerably smaller than the mean free path of the photons. In this case, there is a net production of photons, and they escape from the medium without further interactions.

Experimentally, the interpretation of photon measurements is quite challenging, because the detected photons can come from several sources. It is customary to divide the observed photons into decay photons (produced from the decay of light hadrons, predominantly neutral pions) and direct photons (produced directly from partonic interactions). Direct photons themselves have several sources: some are produced in hard partonic collisions at the time of the impact of the two nuclei, some are produced by the pre-equilibrium medium, some

are produced by the interaction of a hard parton and the medium it traverses [223], some are produced by the QGP and some are produced by the hot hadron gas after the confinement transition. Among all these sources, the photons produced by the QGP and by the hot hadron gas are the most directly sensitive to the properties of the medium produced in heavy ion collisions, but disentangling them unambiguously from the overall spectrum is nearly impossible. In the region of low to intermediate photon energies, the observed spectrum has an exponential shape in $\exp(-p_{\perp}/T_{\text{eff}})$, suggestive of emissions by a thermalized medium. However, as we shall see, the parameter T_{eff} does not have the direct interpretation of the temperature of the producing medium.

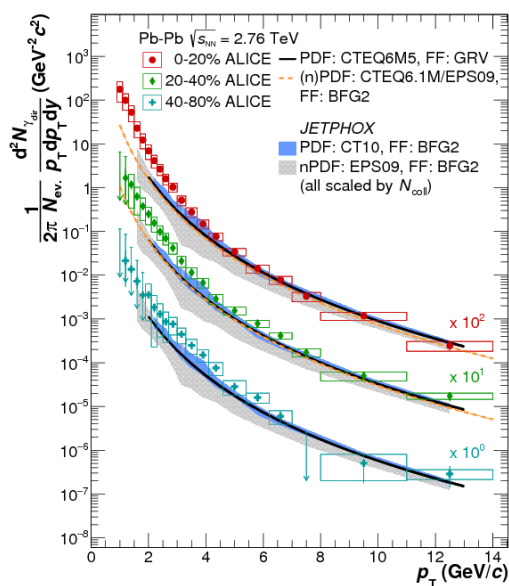


Figure 25: Direct photon spectrum observed by the ALICE experiment, compared to pQCD NLO predictions. Red: exponential fit of the excess. From [224].

Note also that the produced photons can either be real or virtual. In the latter case, they can subsequently decay into a lepton pair (this may be their only decay channel for low invariant masses). Although the decay into a lepton pair requires another electromagnetic coupling and thus reduces the yield, the invariant mass of the pair provides another handle, that may be used to better constrain their possible source (in particular, a non-zero invariant mass may be used to select a region where backgrounds are lower).

Thermal radiation from the QGP The lowest order processes ($qg \rightarrow q\gamma$, $q\bar{q} \rightarrow g\gamma$ [225, 226, 227] and $q\bar{q} \rightarrow \gamma^*$ [228]) have been calculated long ago in an equilibrated quark-gluon plasma. For real photons, the processes with a quark or antiquark exchanged in the t -channel have a logarithmic singularity, which is cured by resumming the appropriate quark hard thermal loop [229, 230]. However, it was soon realized that bremsstrahlung processes, formally of higher order in α_s , suffer from a more severe soft singularity when the photon

has a small invariant mass [231, 232, 233, 234]. This singularity is regularized by the quark in-medium effective mass, but this leads to an enhancement that promotes bremsstrahlung to the same order in α_s as the leading processes.




Figure 26: Top left: Leading Order processes for the production of virtual and real photons. Top right: Enhanced Next-to-Leading Order bremsstrahlung processes. Bottom: multiple scatterings that contribute coherently to the emission of a photon.

This enhancement also occurs for multiple scattering corrections to bremsstrahlung [235]. Like in the discussion of medium induced gluon radiation by a parton, the emission of a photon by a quark or antiquark is affected by the Landau-Pomeranchuk-Migdal effect. The relevant criterion is the comparison between the photon formation time (i.e., how much time is necessary for the photon wave-packet to be sufficiently separated from that of the emitter, which can be estimated to be $t_f^{-1} = E_k(\mathbf{k}_\perp^2 + m^2)/E_p^2$ where k is the photon momentum, p the quark momentum and m its in-medium effective mass) and the mean free path λ between two soft (this is sufficient to induce the emission of a photon) scatterings of a quark in the medium. When $t_f \gtrsim \lambda$, several scatterings are necessary to induce one emission, which effectively reduces the photon yield. Note that, in the weak coupling regime, we have $m \sim gT$ and $\lambda^{-1} \sim g^2T$ (up to logarithms). Therefore, the LPM effect plays a role in two cases: for the production of hard photons emitted at small angle with respect to the quark, or for the production of soft photons. The resummation of these multiple scattering diagrams gives the photon yield at leading order [236, 237, 238, 239]. The next-to-leading order correction has also been calculated more recently [240, 241].

There have also been attempts to extract the photon production rate of a quark-gluon plasma in thermal equilibrium from lattice QCD computations [242, 243, 244, 245, 246, 247], which would in principle be applicable in regimes where the coupling constant may not be small enough for perturbation theory to be reliable. However, since this amounts to computing a spectral function for real energies, there is no direct way to obtain it from an Euclidean lattice formulation. Instead, one can reach it indirectly by unfolding (this is an ill-posed problem, that may be attacked with Bayesian approaches such as the maximal entropy method) the spectral representation of imaginary time correlation functions.

In the context of an actual heavy ion collision, one has also to face the fact that the system may not be in local thermal equilibrium. This is especially true at early times. Firstly, the CGC predicts at the beginning that there are very few quarks compared to the gluons, implying that the system is not yet in chemical equilibrium [248]. It is possible to handle approximately this situation by introducing fugacities for the quark and antiquark distributions, in order to obtain the local photon production rates for a system where the quarks are underpopulated [249]. Moreover, even when quarks and gluons are present in the right proportions, the existence of viscous hydrodynamical corrections implies that their

distributions cannot be the equilibrium ones [250, 251, 252]. These deviations should in principle be taken into account in order to be consistent with the hydrodynamical framework used to describe the bulk evolution of the system.

Thermal radiation from a hot hadron gas Thermal radiation from a hot hadron gas is also obtained from the current-current correlation function, but now it is not accessible to a calculation in terms of the QCD Lagrangian that has quark and gluon degrees of freedom. One may instead use an effective Lagrangian that describes the dynamics and interactions among hadrons at low energy (Such a Lagrangian in general contains some parameters that are constrained by empirical data on lifetimes, cross-sections, etc...) [253, 254, 255, 256, 257, 258, 259]. For dileptons of low mass, the spectral function in the vector channel is dominated by light vector mesons, and is therefore particularly sensitive to thermal modifications of the ρ meson. The main thermal effect is a broadening of the ρ peak (while the center of the peak does not change appreciably), which is even more pronounced when baryons are taken into account. Such a broadening seems in quite good agreement with dilepton measurements in the CERES and NA60 experiments (see Figure 27).

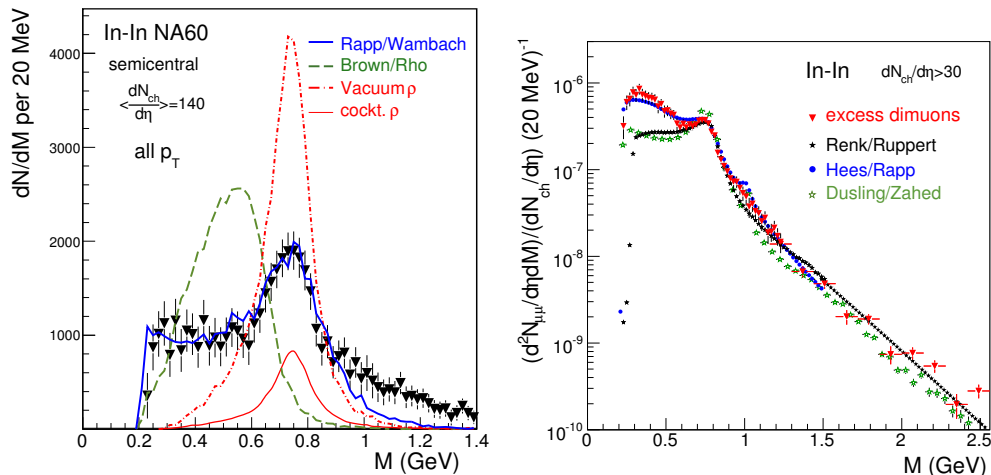


Figure 27: Left: dilepton mass spectrum in the ρ region from the NA60 experiment. Right: invariant mass spectrum of the excess dimuons, integrated over p_T . From [260].

Let us also mention that, besides effective field theories, a more ab initio approach has been used recently in order to extract in-medium spectral functions [261, 262, 263, 264, 265], based on the functional renormalization group (FRG) [140]. The FRG is a functional equation that tracks the evolution of the quantum effective action of a theory as one integrates out the quantum fluctuations in successive layers of momentum, the starting point being the classical action of the theory (i.e., with no quantum fluctuations included). In other words, the FRG is an explicit realization of the renormalization group “a la Wilson”, where a theory is coarse-grained to eliminate its details on short distance scales. In its exact form, there is no practical way to solve the FRG equation, but it is amenable to tractable equations (at least numerically) after some kind of truncation has been performed. When applied to the study of the chiral transition, one observes a temperature dependence of the chiral condensate in quite good agreement with lattice QCD computations, and that the ρ and a_1 mesons become degenerate at high temperature.

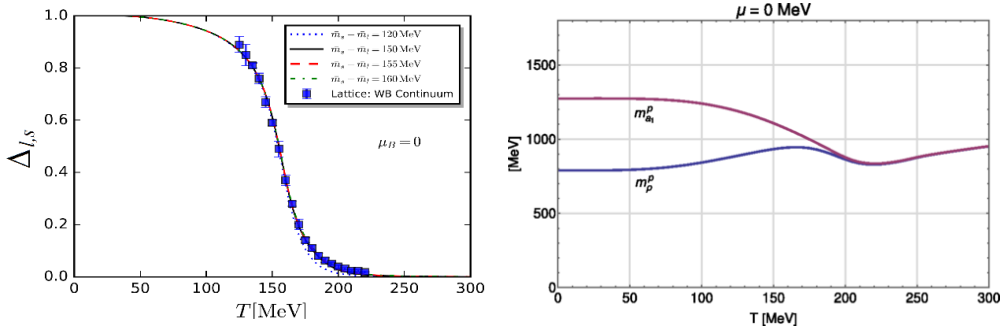


Figure 28: Left: FRG computation of the chiral condensate as a function of temperature, compared to a lattice QCD calculation (From [264]). Right: temperature dependence of the masses of the ρ and a_1 mesons, computed in the FRG framework (From [262]).

Folding with the medium evolution The calculations described above provide a local photon production rate (i.e., the number of photons produced per unit of time and volume, given the local temperature of the plasma). To go from there to a photon spectrum that one may compare with experimental data, this local rate must be integrated over the entire space-time history of the collision (in addition, at each space-time point, one must boost the local spectrum by the 4-velocity of the plasma at this point). The important point to keep in mind is that, even if a fit of the resulting spectrum by an exponential of the form $\exp(-p_\perp/T_{\text{eff}})$ appears to work, the parameter T_{eff} does not have a direct interpretation as the temperature of the plasma. In particular, this effective temperature is blue-shifted by the radial flow of the medium. Moreover, the integration over time is also more sensitive to the stages of the evolution where the system spends more time. With dileptons, one may avoid this blue-shifting effect when looking at the spectrum as a function of the pair invariant mass since the mass is not affected by the flow, but the extracted temperature still reflects a spacetime average rather than an instantaneous temperature. With this caveat in mind, an exponential mass dependence $\exp(-M/T_{\text{eff}})$ of the dilepton spectrum has been reported by the NA60 experiment, with an effective temperature $T_{\text{eff}} \approx 205 \pm 12$ MeV [260].

6.2 Heavy quarkonia in a hot medium

Qualitative aspects Bound states made of heavy quarks can also be viewed as potential “thermometers” [266]. On the theory side, the advantage of considering sufficiently heavy quarks is that they provide a large mass scale (much larger than the QCD non-perturbative scale Λ_{QCD}), and this separation of scales allow the use of effective field theory descriptions such as non-relativistic QCD. In such a non-relativistic framework, one may use the concept of interaction potential between a pair of heavy quarks, in conjunction with a non-relativistic Schrödinger equation. Another theoretical simplification regarding heavy quarks is that their production happens at very early times (of the order of the inverse of their mass), and involve parton distributions at reasonably large momentum fractions, where saturation effects are not important (this assertion should be contrasted in the case of charm quarks at the LHC energy – see [267, 268] for a recent study of J/Ψ production in proton-nucleus collisions; a similar computation in nucleus-nucleus collisions could be done by solving the Dirac equation in the glasma color fields, but is considerably more challenging [269, 270, 271]). Thus, in heavy ion collisions, one is mostly interested in the subsequent fate of the produced heavy quarks, rather than the production itself. As far as experimental measurements are concerned, heavy quarkonia also offer clean signals via their dilepton decay channel.

Consider for instance a heavy meson $Q\bar{Q}$. Loosely speaking, when the thermal excitation energy (i.e., the energy gain provided by the absorption of a gluon from the surrounding thermal bath) is equal or larger to the binding energy of the quark-antiquark pair, the meson may be dissociated. An alternate way of describing this phenomenon is by noting that the interaction potential of the $Q\bar{Q}$ is affected by Debye screening in the presence of a dense medium. When the Debye screening length becomes shorter than the size of the would-be bound state, its dissociation occurs. Afterwards, the quark and the antiquark evolve independently in the medium, and the most likely outcome is that –when the temperature has decreased below the confinement temperature– they bind with one of the light quarks or antiquarks from the surrounding medium, in order to form heavy-light mesons (D and B mesons). Therefore, in the extreme version of this scenario, the yield of $Q\bar{Q}$ mesons would be almost completely suppressed and the produced heavy quarks would all be recovered in the form of open-flavour mesons. Note that in the case of charm quarks, whose production is quite abundant in heavy ion collisions at high energy, the dissociated c and \bar{c} may have a high enough density for accidental recombinations in J/Ψ 's to be important.

Experimentally, a particularly clear observation of this phenomenon has been performed at the LHC for $b\bar{b}$ mesons, as shown in the figure 29. There, one sees a suppression of the




Figure 29: Muon-antimuon invariant mass spectrum in the upsilon region. From [272]. See also [273, 274].

yields of the Υ'' and Υ' mesons, while the ground state Υ has an almost unmodified yield. Given the above discussion, the interpretation of this result is that the gluon thermal energy at the temperatures reached in the collision is sufficient to dissociate the excited states (because they are more weakly bound) but not enough to alter the more tightly bound fundamental state.

In-medium singlet heavy quark bound states In the vacuum (i.e., at zero temperature), the spectrum of color singlet $Q\bar{Q}$ states is typically made of a few (meta)stable states –defined as poles of the propagator in the energy complex plane– occupying the low energy region, and a continuum at higher energies. In the presence of a high temperature medium, several effects may occur:

- Some bound states (starting from the high lying ones) may disappear,
- The lower bound of the continuum of free states may move to a lower energy,
- The surrounding thermal medium can induce transitions between various $Q\bar{Q}$ states, including transitions between singlet and octet states.

The ab initio approach for studying the in-medium modifications of heavy quark states consists in calculating the spectral function in the appropriate channel at non-zero temperature. However, this is not doable in perturbation theory due to the non-perturbative nature of bound states. Although lattice QCD is by construction a non-perturbative approach, spectral functions cannot be computed directly. Instead, one can calculate an Euclidean propagator $G(\tau, \mathbf{p})$ which is related to the corresponding spectral function $\rho(\omega, \mathbf{p})$ as follows:

$$G(\tau, \mathbf{p}) = \int d\omega \rho(\omega, \mathbf{p}) \frac{\cosh(\omega(\tau - 1/2T))}{\sinh(\omega/2T)}. \quad (53)$$

A first obvious difficulty for inverting this relationship is that the propagator is computed only at the finite set of Euclidean times τ that exist in the employed lattice setup, while the expected spectral function depends on a continuous energy ω . But even if the propagator was known at all the real τ 's in the range $[0, 1/T]$, this inversion is a mathematically ill-posed problem, because the linear mapping from ρ to G has zero modes (i.e., functions ρ that give zero when inserted in the integral in the previous equation). Thus, even in the ideal situation where the propagator would be known exactly, the inversion can only be performed up to a linear combination of these zero modes.

A possible strategy is to remove this ambiguity by imposing additional (but generic enough so that they do not bias the outcome in unphysical directions) constraints on the expected spectral function. A minimal constraint that helps disambiguate the answer is to request the positivity of the spectral function. In practice, this can be implemented by using the *maximal entropy method*, which is a Bayesian method for finding the most likely spectral function consistent with the computed values of the propagator and the additional constraints. When using this approach, it is necessary to have very accurate lattice data for a robust extraction of the spectral function (otherwise, the extracted spectral function may be dominated by the additional constraints imposed on the solution). Another limitation is that it is practically impossible to be sensitive to excited states, as this would require an exponentially large statistics. With these caveats in mind, the general trend observed for heavy bound states is a sequential melting of states, starting with the high lying ones, and a trend towards negative medium-induced mass shifts [275, 276, 277, 278].

An alternative to Bayesian methods is to model the spectral functions $\rho(\omega, \mathbf{p})$ with a few free parameters and to perform a standard fit to adjust these parameters in order to reproduce the computed propagator $G(\tau, \mathbf{p})$. Besides the location and width of the lowest lying peak, the model may contain parameters that describe the transport properties of heavy quarks, or excited states. However, one should keep in mind that the propagator may be very weakly sensitive to these additional features of the spectral function [279], and that an unrealistic modeling may introduce a strong model dependence on the outcome. A model of the spectral function may be obtained from effective field theory descriptions, such as non-relativistic QCD (NRQCD – obtained from QCD by integrating out the heavy quark mass scale m_Q) or even potential non-relativistic QCD (pNRQCD – obtained from NRQCD by further integrating out the softer scale $m_Q v$, where v is the heavy quark velocity) [280, 281, 282], that rely on the heavy mass of the quarks, both compared to the QCD non-perturbative scale Λ_{QCD} and to the typical energy scale of the surrounding medium.

In fact, the pNRQCD effective theory also provides a proper connection between QCD and the non-relativistic Schrödinger equation used in order to study the $Q\bar{Q}$ bound states (see for instance [283, 284]). Indeed, one of the parameters in pNRQCD is the $Q\bar{Q}$ interaction potential (that should in principle be obtained by a matching to the underlying field theory), and its equation of motion in the singlet sector is a Schrödinger equation whose discrete energy levels correspond to the $Q\bar{Q}$ singlet bound states. The singlet potential $V_s(r)$ can be

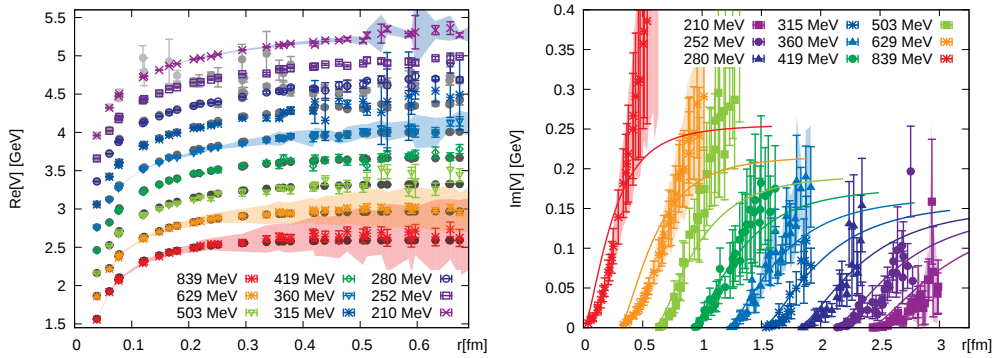


Figure 30: Real and imaginary parts of the singlet $Q\bar{Q}$ potential at various temperatures. From [285]. See also [286].

obtained as

$$V_s(r) = \lim_{t \rightarrow \infty} \frac{\partial_t W(r, t)}{W(r, t)}, \quad (54)$$

where $W(r, t)$ is a rectangular Wilson loop of spatial extent r and temporal extent t . Note however that t is here the Minkowski time, and a direct evaluation of this quantity in lattice QCD is therefore not possible. A possible strategy is to start from an Euclidean rectangular Wilson loop, whose associated spectral function can be used to express $V_s(r)$, thus allowing to constrain the potential via Bayesian methods as shown in Figure 30. Note that the real part of this potential behaves similarly but is not identical to the potential sometimes inferred from the logarithm of the free energy of a singlet $Q\bar{Q}$ pair [287]. (The imaginary part, due to Landau damping and transitions from singlet to octet states, is not present in the potential defined from the free energy.)

In-medium dynamical evolution In the presence of a high temperature medium, not only the spectrum of singlet quarkonia states is modified, but the surrounding medium can also induce transitions between various singlet states, and between singlet and octet states. In equilibrium, there are as many transitions in either direction, and the density matrix of the system is time independent. The situation is far more complicated out-of-equilibrium, since the density matrix is now time dependent. Various forms of kinetic or stochastic equations have been employed to describe the evolution of heavy quarks and quarkonia embedded in a medium [288, 289, 290].

The main physical effect one would like to study here is the possible recombination of the heavy quarks and antiquarks, that may have a sizeable probability of close encounter when their density is large enough (which appears to be the case for charm quarks). When this is the case, the final yield of quarkonia bound states is enhanced compared to what one would get with the assumption that all the bound states that are dissociated go into open-heavy flavor mesons [291, 292, 293, 294, 295]. This phenomenon can be seen for charm quarks




Figure 31: Nuclear modification factor of J/ψ as a function of the number of participants, for RHIC and LHC energies. From [296].

by comparing the nuclear modification factor of J/ψ at RHIC and LHC energies, as shown in Figure 31. On this plot, one sees that J/ψ 's are less suppressed in central collisions at higher energy, which can be explained by the more abundant production of c and \bar{c} quarks at the LHC, which favors their recombination into charmonium.

More recently, there has been an effort to derive these description from a more fundamental starting point, shedding some light on the sequence of approximations necessary to obtain them. The strategy, based on the framework of *open quantum systems* [297, 298, 299, 300, 301, 302], consists in starting from the system made of the $Q\bar{Q}$ pairs and their environment. This system is closed and evolves unitarily, with a density matrix ρ that obeys the von Neumann equation (also known as the quantum Liouville equation),

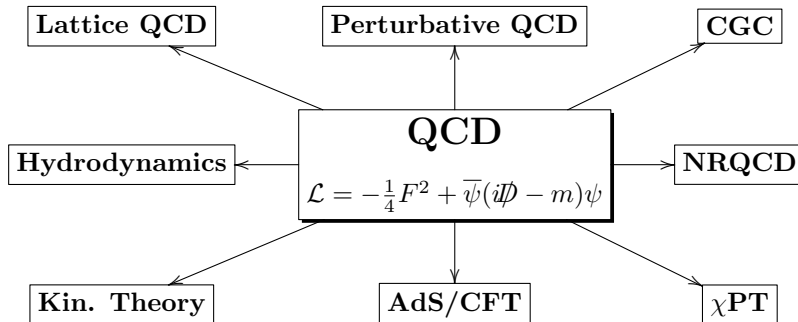
$$\frac{d\rho}{dt} = -i[H, \rho]. \quad (55)$$

The density matrix ρ contains a lot of information regarding the surrounding medium in which the heavy quarks are embedded. A *reduced density matrix* describing specifically the heavy quark degrees of freedom can be obtained by performing a partial trace over environment degrees of freedom, $\rho_Q \equiv \text{tr}_{\text{env}}(\rho)$. However, after doing this, the evolution equation for the reduced matrix density is no longer a von Neumann equation, and is in general dissipative and non unitary (this is just a consequence of the fact that the subsystem made of the heavy quarks is an open system). With the assumption that the environment relaxes more quickly than the heavy quarks, it has been shown that ρ_Q obeys a Lindblad equation. Under the additional assumption that the transitions between the various $Q\bar{Q}$ states are faster than their adjustment to changes of the environment, this Lindblad equation can lead to a Boltzmann equation (and a rate equation is the momenta are integrated out).

7 Conclusions, open questions for the future

What should first come to mind after this brief survey of the theoretical aspects of heavy ion collisions is the difficulty of describing a large, time-dependent, non-equilibrated system in terms of a (rather complicated) underlying microscopic theory (QCD). Since an ab-initio description in terms of QCD of these collisions is not practically feasible, most approaches

are based on effective descriptions that capture the relevant dynamics at scales larger than the typical QCD scales. Thus, instead of a unique theory from which everything would be derived, theoretical works in this field use instead a large variety of tools that are more or less connected to QCD:



Although we have not spent much time discussing this, we should also stress the fact that many observables in heavy ion collisions depend on a number of mundane aspects of low energy nuclear physics, namely the shape and size of the nuclei, and the distribution of the nucleons inside a nucleus and the fluctuations thereof. These properties, that are not the main targets of the heavy-ion collision program, play nevertheless an important role when trying to uncover some property of QCD from experimental data.

By a combination of experimental and theoretical efforts, many properties of the quark-gluon plasma have been uncovered:

- the QGP is a nearly perfect fluid,
- its shear viscosity to entropy ratio is in the range $[1, 2.5]$ (in units of $\hbar/4\pi$), making it the substance with the smallest ratio so far,
- its equation of state is consistent with lattice QCD expectations, and with the deconfinement of the color degrees of freedom,
- the yield of “light” partons, including charm quarks, is significantly suppressed compared to rescaled proton-proton collisions,
- the suppression of bottom quarks is less pronounced, in agreement with theoretical expectations (dead-cone effect due to the mass of the emitter),
- the studies of energy loss can now be supplemented by direct observations of reconstructed jets. This has allowed to determine that a large amount of energy is radiated by soft emissions at large angle,
- a sequential pattern has been observed in the disappearance of $b\bar{b}$ bound states, consistent with the theoretical understanding of the dissociation phenomenon,
- at the highest energies, the production of charm quarks is copious enough to lead to the formation of J/ψ bound states by recombination of uncorrelated quarks and antiquarks.

The picture that emerges from these observations is that the matter produced in heavy ion collisions is a very “opaque”, strongly interacting, fluid, in rather sharp contrast with

the ethereal quark-gluon *plasma* that was the common point of view before the RHIC experiment. Despite many progresses, it is also clear that extracting the underlying QCD properties from the outcome of heavy ion collisions is extremely difficult, since in several instances the comparisons with QCD have remained rather qualitative although the experimental measurements were quite detailed. Another source of complication is that some of these studies are done by comparing the outcome of nucleus-nucleus collisions with that of rescaled proton-nucleus or proton-proton collisions. However, high energy proton-nucleus and proton-proton collisions have turned out to display some features close to those observed in nucleus-nucleus collisions, casting some doubts on their use as “references” to compare with in order to pinpoint effects specific to the quark-gluon plasma. If flow is confirmed to occur even in proton-proton collisions, one will have to learn to live without such a reference for certain observables.

It is of course hard to predict where the next advances will happen, but given the areas that have received most attention in the past years, some improvements are probably within reach in a reasonable future in the following directions:

- Determine the temperature dependence of the shear viscosity,
- Obtain a better determination of the bulk viscosity,
- Better disentangle the mechanisms of energy loss, especially in the case of jets,
- Characterize *when* heavy quark bound states are formed,
- estimate the initial temperature from thermal photons and the melting of quarkonia,
- Clarify to what extent the concept of flow applies to the system formed in proton-proton collisions. This entails a more robust control over the state of the system immediately after the collision, in order to disentangle initial flow from the flow hydrodynamically generated later on.

For this to be possible, besides some improvements to the calculation of the elementary relevant phenomena, an important aspect (and difficulty) is to merge as seamlessly as possible tools that have originally been developed independently. This is especially true for observables that depend on the interactions between some probe and the surrounding medium, for which it is crucial to use a modeling of the background and its evolution which is as realistic as possible.

Acknowledgements: The author is supported by the Agence Nationale de la Recherche through the project 11-BS04-015-01.

References

- [1] David J. Gross and Frank Wilczek. Ultraviolet Behavior of Nonabelian Gauge Theories. *Phys. Rev. Lett.*, 30:1343–1346, 1973.
- [2] H.David Politzer. Reliable Perturbative Results for Strong Interactions? *Phys. Rev. Lett.*, 30:1346–1349, 1973.
- [3] M. Tanabashi et al. Review of Particle Physics. *Phys. Rev. D*, 98(3):030001, 2018.
- [4] Kenneth G. Wilson. Confinement of Quarks. pages 45–59, 2 1974.

- [5] Philippe de Forcrand and Massimo D’Elia. Continuum limit and universality of the Columbia plot. *PoS, LATTICE2016:081*, 2017.
- [6] P. Hasenfratz and F. Karsch. Chemical Potential on the Lattice. *Phys. Lett. B*, 125:308–310, 1983.
- [7] Mark G. Alford, Anton Kapustin, and Frank Wilczek. Imaginary chemical potential and finite fermion density on the lattice. *Phys. Rev. D*, 59:054502, 1999.
- [8] Simon Hands, John B. Kogut, Maria-Paola Lombardo, and Susan E. Morrison. Symmetries and spectrum of SU(2) lattice gauge theory at finite chemical potential. *Nucl. Phys. B*, 558:327–346, 1999.
- [9] Z. Fodor and S.D. Katz. A New method to study lattice QCD at finite temperature and chemical potential. *Phys. Lett. B*, 534:87–92, 2002.
- [10] S. Gupta. Lattice QCD with chemical potential: Evading the fermion-sign problem. *Pramana*, 63:1211–1224, 2004.
- [11] Mark G. Alford, Krishna Rajagopal, and Frank Wilczek. QCD at finite baryon density: Nucleon droplets and color superconductivity. *Phys. Lett. B*, 422:247–256, 1998.
- [12] Juergen Berges and Krishna Rajagopal. Color superconductivity and chiral symmetry restoration at nonzero baryon density and temperature. *Nucl. Phys. B*, 538:215–232, 1999.
- [13] D.T. Son. Superconductivity by long range color magnetic interaction in high density quark matter. *Phys. Rev. D*, 59:094019, 1999.
- [14] Mark G. Alford, Andreas Schmitt, Krishna Rajagopal, and Thomas Schäfer. Color superconductivity in dense quark matter. *Rev. Mod. Phys.*, 80:1455–1515, 2008.
- [15] Ulrich W. Heinz and Maurice Jacob. Evidence for a new state of matter: An Assessment of the results from the CERN lead beam program. 1 2000.
- [16] M.C. Abreu et al. Evidence for deconfinement of quarks and gluons from the J / psi suppression pattern measured in Pb + Pb collisions at the CERN SPS. *Phys. Lett. B*, 477:28–36, 2000.
- [17] E. Andersen et al. Strangeness enhancement at mid-rapidity in Pb Pb collisions at 158-A-GeV/c. *Phys. Lett. B*, 449:401–406, 1999.
- [18] G. Agakichiev et al. Low mass e+ e- pair production in 158/A-GeV Pb - Au collisions at the CERN SPS, its dependence on multiplicity and transverse momentum. *Phys. Lett. B*, 422:405–412, 1998.
- [19] I. Bearden et al. Strange meson enhancement in Pb Pb collisions. *Phys. Lett. B*, 471:6–12, 1999.
- [20] F. Siklér et al. Hadron production in nuclear collisions from the NA49 experiment at 158-GeV/c/A. *Nucl. Phys. A*, 661:45–54, 1999.
- [21] Sonja Kabana et al. Impact parameter dependence of pi+-, K+-, p, anti-p, d and anti-d production in Pb + Pb collisions at 158-A-GeV. *Nucl. Phys. A*, 661:370–373, 1999.
- [22] M. M. Aggarwal et al. Freezeout parameters in central 158/A-GeV Pb-208 + Pb-208 collisions. *Phys. Rev. Lett.*, 83:926–930, 1999.

- [23] John Adams et al. Experimental and theoretical challenges in the search for the quark gluon plasma: The STAR Collaboration’s critical assessment of the evidence from RHIC collisions. *Nucl. Phys. A*, 757:102–183, 2005.
- [24] K. Adcox et al. Formation of dense partonic matter in relativistic nucleus-nucleus collisions at RHIC: Experimental evaluation by the PHENIX collaboration. *Nucl. Phys. A*, 757:184–283, 2005.
- [25] I. Arsene et al. Quark gluon plasma and color glass condensate at RHIC? The Perspective from the BRAHMS experiment. *Nucl. Phys. A*, 757:1–27, 2005.
- [26] B.B. Back et al. The PHOBOS perspective on discoveries at RHIC. *Nucl. Phys. A*, 757:28–101, 2005.
- [27] K Aamodt et al. Elliptic flow of charged particles in Pb-Pb collisions at 2.76 TeV. *Phys. Rev. Lett.*, 105:252302, 2010.
- [28] K. Aamodt et al. Suppression of Charged Particle Production at Large Transverse Momentum in Central Pb-Pb Collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Phys. Lett. B*, 696:30–39, 2011.
- [29] K. Aamodt et al. Higher harmonic anisotropic flow measurements of charged particles in Pb-Pb collisions at $\sqrt{s_{NN}}=2.76$ TeV. *Phys. Rev. Lett.*, 107:032301, 2011.
- [30] Betty Abelev et al. Long-range angular correlations on the near and away side in p -Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV. *Phys. Lett. B*, 719:29–41, 2013.
- [31] Georges Aad et al. Observation of a Centrality-Dependent Dijet Asymmetry in Lead-Lead Collisions at $\sqrt{s_{NN}} = 2.77$ TeV with the ATLAS Detector at the LHC. *Phys. Rev. Lett.*, 105:252303, 2010.
- [32] Georges Aad et al. Measurement of the azimuthal anisotropy for charged particle production in $\sqrt{s_{NN}} = 2.76$ TeV lead-lead collisions with the ATLAS detector. *Phys. Rev. C*, 86:014907, 2012.
- [33] Georges Aad et al. Observation of Associated Near-Side and Away-Side Long-Range Correlations in $\sqrt{s_{NN}}=5.02$ TeV Proton-Lead Collisions with the ATLAS Detector. *Phys. Rev. Lett.*, 110(18):182302, 2013.
- [34] Vardan Khachatryan et al. Observation of Long-Range Near-Side Angular Correlations in Proton-Proton Collisions at the LHC. *JHEP*, 09:091, 2010.
- [35] Serguei Chatrchyan et al. Observation and studies of jet quenching in PbPb collisions at nucleon-nucleon center-of-mass energy = 2.76 TeV. *Phys. Rev. C*, 84:024906, 2011.
- [36] Serguei Chatrchyan et al. Observation of Long-Range Near-Side Angular Correlations in Proton-Lead Collisions at the LHC. *Phys. Lett. B*, 718:795–814, 2013.
- [37] L.V. Gribov, E.M. Levin, and M.G. Ryskin. Semihard Processes in QCD. *Phys. Rept.*, 100:1–150, 1983.
- [38] Alfred H. Mueller and Jian-wei Qiu. Gluon Recombination and Shadowing at Small Values of x . *Nucl. Phys. B*, 268:427–452, 1986.
- [39] J.P. Blaizot and Alfred H. Mueller. The Early Stage of Ultrarelativistic Heavy Ion Collisions. *Nucl. Phys. B*, 289:847–860, 1987.
- [40] Abhay Deshpande, Rolf Ent, and Richard Milner. The EIC’s route to a new frontier in QCD. *CERN Cour.*, 49N9:13–15, 2009.

- [41] Larry D. McLerran and Raju Venugopalan. Computing quark and gluon distribution functions for very large nuclei. *Phys. Rev. D*, 49:2233–2241, 1994.
- [42] Larry D. McLerran and Raju Venugopalan. Gluon distribution functions for very large nuclei at small transverse momentum. *Phys. Rev. D*, 49:3352–3355, 1994.
- [43] Edmond Iancu, Andrei Leonidov, and Larry D. McLerran. Nonlinear gluon evolution in the color glass condensate. 1. *Nucl. Phys. A*, 692:583–645, 2001.
- [44] Edmond Iancu, Andrei Leonidov, and Larry D. McLerran. The Renormalization group equation for the color glass condensate. *Phys. Lett. B*, 510:133–144, 2001.
- [45] Elena Ferreiro, Edmond Iancu, Andrei Leonidov, and Larry McLerran. Nonlinear gluon evolution in the color glass condensate. 2. *Nucl. Phys. A*, 703:489–538, 2002.
- [46] Edmond Iancu, Andrei Leonidov, and Larry McLerran. The Color glass condensate: An Introduction. In *Cargese Summer School on QCD Perspectives on Hot and Dense Matter*, pages 73–145, 2 2002.
- [47] Heribert Weigert. Evolution at small $x(bj)$: The Color glass condensate. *Prog. Part. Nucl. Phys.*, 55:461–565, 2005.
- [48] Francois Gelis, Edmond Iancu, Jamal Jalilian-Marian, and Raju Venugopalan. The Color Glass Condensate. *Ann. Rev. Nucl. Part. Sci.*, 60:463–489, 2010.
- [49] T. Lappi. Wilson line correlator in the MV model: Relating the glasma to deep inelastic scattering. *Eur. Phys. J. C*, 55:285–292, 2008.
- [50] Francois Gelis and Raju Venugopalan. Particle production in field theories coupled to strong external sources. *Nucl. Phys. A*, 776:135–171, 2006.
- [51] Francois Gelis, Tuomas Lappi, and Raju Venugopalan. High energy scattering in Quantum Chromodynamics. *Int. J. Mod. Phys. E*, 16:2595–2637, 2007.
- [52] Alex Kovner, Larry D. McLerran, and Heribert Weigert. Gluon production from nonAbelian Weizsacker-Williams fields in nucleus-nucleus collisions. *Phys. Rev. D*, 52:6231–6237, 1995.
- [53] Alex Krasnitz and Raju Venugopalan. Nonperturbative computation of gluon minijet production in nuclear collisions at very high-energies. *Nucl. Phys. B*, 557:237, 1999.
- [54] Alex Krasnitz and Raju Venugopalan. The Initial energy density of gluons produced in very high-energy nuclear collisions. *Phys. Rev. Lett.*, 84:4309–4312, 2000.
- [55] Alex Krasnitz and Raju Venugopalan. The Initial gluon multiplicity in heavy ion collisions. *Phys. Rev. Lett.*, 86:1717–1720, 2001.
- [56] Alex Krasnitz, Yasushi Nara, and Raju Venugopalan. Coherent gluon production in very high-energy heavy ion collisions. *Phys. Rev. Lett.*, 87:192302, 2001.
- [57] T. Lappi. Production of gluons in the classical field model for heavy ion collisions. *Phys. Rev. C*, 67:054903, 2003.
- [58] Francois Gelis, Tuomas Lappi, and Raju Venugopalan. High energy factorization in nucleus-nucleus collisions. *Phys. Rev. D*, 78:054019, 2008.
- [59] Francois Gelis, Tuomas Lappi, and Raju Venugopalan. High energy factorization in nucleus-nucleus collisions. II. Multigluon correlations. *Phys. Rev. D*, 78:054020, 2008.

- [60] Alejandro Ayala, Jamal Jalilian-Marian, Larry D. McLerran, and Raju Venugopalan. Quantum corrections to the Weizsacker-Williams gluon distribution function at small x . *Phys. Rev. D*, 53:458–475, 1996.
- [61] Jamal Jalilian-Marian, Alex Kovner, Larry D. McLerran, and Heribert Weigert. The Intrinsic glue distribution at very small x . *Phys. Rev. D*, 55:5414–5428, 1997.
- [62] Francois Gelis, Tuomas Lappi, and Raju Venugopalan. High energy factorization in nucleus-nucleus collisions. 3. Long range rapidity correlations. *Phys. Rev. D*, 79:094017, 2009.
- [63] Jean-Paul Blaizot, Edmond Iancu, and Heribert Weigert. Nonlinear gluon evolution in path integral form. *Nucl. Phys. A*, 713:441–469, 2003.
- [64] Kari Rummukainen and Heribert Weigert. Universal features of JIMWLK and BK evolution at small x . *Nucl. Phys. A*, 739:183–226, 2004.
- [65] T. Lappi. Gluon spectrum in the glasma from JIMWLK evolution. *Phys. Lett. B*, 703:325–330, 2011.
- [66] Adrian Dumitru, Jamal Jalilian-Marian, Tuomas Lappi, Bjoern Schenke, and Raju Venugopalan. Renormalization group evolution of multi-gluon correlators in high energy QCD. *Phys. Lett. B*, 706:219–224, 2011.
- [67] T. Lappi and H. Mäntysaari. Proposal for a running coupling JIMWLK equation. *Nucl. Phys. A*, 932:69–74, 2014.
- [68] Ian Balitsky and Giovanni A. Chirilli. Rapidity evolution of Wilson lines at the next-to-leading order. *Phys. Rev. D*, 88:111501, 2013.
- [69] Alex Kovner, Michael Lublinsky, and Yair Mulian. Jalilian-Marian, Iancu, McLerran, Weigert, Leonidov, Kovner evolution at next to leading order. *Phys. Rev. D*, 89(6):061704, 2014.
- [70] I. Balitsky. Operator expansion for high-energy scattering. *Nucl. Phys. B*, 463:99–160, 1996.
- [71] Yuri V. Kovchegov. Small x $F(2)$ structure function of a nucleus including multiple pomeron exchanges. *Phys. Rev. D*, 60:034008, 1999.
- [72] Yuri V. Kovchegov and Heribert Weigert. Triumvirate of Running Couplings in Small- x Evolution. *Nucl. Phys. A*, 784:188–226, 2007.
- [73] Einan Gardi, Janne Kuokkanen, Kari Rummukainen, and Heribert Weigert. Running coupling and power corrections in nonlinear evolution at the high-energy limit. *Nucl. Phys. A*, 784:282–340, 2007.
- [74] Ian Balitsky and Giovanni A. Chirilli. Next-to-leading order evolution of color dipoles. *Phys. Rev. D*, 77:014019, 2008.
- [75] E. Iancu, J.D. Madrigal, A.H. Mueller, G. Soyez, and D.N. Triantafyllopoulos. Resumming double logarithms in the QCD evolution of color dipoles. *Phys. Lett. B*, 744:293–302, 2015.
- [76] T. Lappi and H. Mäntysaari. Next-to-leading order Balitsky-Kovchegov equation with resummation. *Phys. Rev. D*, 93(9):094004, 2016.
- [77] Javier L. Albacete and Cyrille Marquet. Gluon saturation and initial conditions for relativistic heavy ion collisions. *Prog. Part. Nucl. Phys.*, 76:1–42, 2014.

- [78] Tolga Altinoluk, Nestor Armesto, Guillaume Beuf, Alex Kovner, and Michael Lublinsky. Single-inclusive particle production in proton-nucleus collisions at next-to-leading order in the hybrid formalism. *Phys. Rev. D*, 91(9):094016, 2015.
- [79] T. Lappi and H. Mäntysaari. Direct numerical solution of the coordinate space Balitsky-Kovchegov equation at next to leading order. *Phys. Rev. D*, 91(7):074016, 2015.
- [80] E. Iancu, J.D. Madrigal, A.H. Mueller, G. Soyez, and D.N. Triantafyllopoulos. Collinearly-improved BK evolution meets the HERA data. *Phys. Lett. B*, 750:643–652, 2015.
- [81] T. Lappi and L. McLerran. Some features of the glasma. *Nucl. Phys. A*, 772:200–212, 2006.
- [82] Kenji Fukushima and Francois Gelis. The evolving Glasma. *Nucl. Phys. A*, 874:108–129, 2012.
- [83] Berndt Muller and Atanas Trayanov. Deterministic chaos in nonAbelian lattice gauge theory. *Phys. Rev. Lett.*, 68:3387–3390, 1992.
- [84] T.S. Biro, C. Gong, Berndt Muller, and A. Trayanov. Hamiltonian dynamics of Yang-Mills fields on a lattice. *Int. J. Mod. Phys. C*, 5:113–149, 1994.
- [85] Teiji Kunihiro, Berndt Muller, Akira Ohnishi, Andreas Schafer, Toru T. Takahashi, and Arata Yamamoto. Chaotic behavior in classical Yang-Mills dynamics. *Phys. Rev. D*, 82:114015, 2010.
- [86] Paul Romatschke and Raju Venugopalan. Collective non-Abelian instabilities in a melting color glass condensate. *Phys. Rev. Lett.*, 96:062302, 2006.
- [87] F. Gelis. Color Glass Condensate and Glasma. *Int. J. Mod. Phys. A*, 28:1330001, 2013.
- [88] Thomas Epelbaum and Francois Gelis. Fluctuations of the initial color fields in high energy heavy ion collisions. *Phys. Rev. D*, 88:085015, 2013.
- [89] Thomas Epelbaum and Francois Gelis. Pressure isotropization in high energy heavy ion collisions. *Phys. Rev. Lett.*, 111:232301, 2013.
- [90] J. Berges, K. Boguslavski, S. Schlichting, and R. Venugopalan. Turbulent thermalization process in heavy-ion collisions at ultrarelativistic energies. *Phys. Rev. D*, 89(7):074011, 2014.
- [91] Juergen Berges, Kirill Boguslavski, Soeren Schlichting, and Raju Venugopalan. Universal attractor in a highly occupied non-Abelian plasma. *Phys. Rev. D*, 89(11):114007, 2014.
- [92] J.M. Luttinger and John Clive Ward. Ground state energy of a many fermion system. 2. *Phys. Rev.*, 118:1417–1427, 1960.
- [93] Gordon Baym and Leo P. Kadanoff. Conservation Laws and Correlation Functions. *Phys. Rev.*, 124:287–299, 1961.
- [94] Juergen Berges. Introduction to nonequilibrium quantum field theory. *AIP Conf. Proc.*, 739(1):3–62, 2004.
- [95] Hendrik van Hees and Joern Knoll. Renormalization in selfconsistent approximations schemes at finite temperature. 1. Theory. *Phys. Rev. D*, 65:025010, 2002.

- [96] Juergen Berges, Szabolcs Borsanyi, Urko Reinosa, and Julien Serreau. Nonperturbative renormalization for 2PI effective action techniques. *Annals Phys.*, 320:344–398, 2005.
- [97] Yoshitaka Hatta and Akihiro Nishiyama. Towards thermalization in heavy-ion collisions: CGC meets the 2PI formalism. *Nucl. Phys. A*, 873:47–67, 2012.
- [98] A.H. Mueller and D.T. Son. On the Equivalence between the Boltzmann equation and classical field theory at large occupation numbers. *Phys. Lett. B*, 582:279–287, 2004.
- [99] Sangyong Jeon. The Boltzmann equation in classical and quantum field theory. *Phys. Rev. C*, 72:014907, 2005.
- [100] Thomas Epelbaum, Francois Gelis, Naoto Tanji, and Bin Wu. Properties of the Boltzmann equation in the classical approximation. *Phys. Rev. D*, 90(12):125032, 2014.
- [101] Thomas Epelbaum, Francois Gelis, Sangyong Jeon, Guy Moore, and Bin Wu. Kinetic theory of a longitudinally expanding system of scalar particles. *JHEP*, 09:117, 2015.
- [102] Peter Brockway Arnold, Guy D. Moore, and Laurence G. Yaffe. Effective kinetic theory for high temperature gauge theories. *JHEP*, 01:030, 2003.
- [103] Alekski Kurkela and Yan Zhu. Isotropization and hydrodynamization in weakly coupled heavy-ion collisions. *Phys. Rev. Lett.*, 115(18):182301, 2015.
- [104] Jean-Paul Blaizot and Li Yan. Onset of hydrodynamics for a quark-gluon plasma from the evolution of moments of distribution functions. *JHEP*, 11:161, 2017.
- [105] Jean-Paul Blaizot and Li Yan. Fluid dynamics of out of equilibrium boost invariant plasmas. *Phys. Lett. B*, 780:283–286, 2018.
- [106] Jean-Paul Blaizot and Li Yan. Emergence of hydrodynamical behavior in expanding ultra-relativistic plasmas. *Annals Phys.*, 412:167993, 2020.
- [107] Jean-Paul Blaizot and Li Yan. Analytical attractor for Bjorken expansion. 6 2020.
- [108] Derek A. Teaney. *Viscous Hydrodynamics and the Quark Gluon Plasma*, pages 207–266. 2010.
- [109] Paul Romatschke. New Developments in Relativistic Viscous Hydrodynamics. *Int. J. Mod. Phys. E*, 19:1–53, 2010.
- [110] Paul Romatschke and Ulrike Romatschke. *Relativistic Fluid Dynamics In and Out of Equilibrium*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 5 2019.
- [111] Wojciech Florkowski, Michal P. Heller, and Michal Spalinski. New theories of relativistic hydrodynamics in the LHC era. *Rept. Prog. Phys.*, 81(4):046001, 2018.
- [112] Charles Gale, Sangyong Jeon, and Bjoern Schenke. Hydrodynamic Modeling of Heavy-Ion Collisions. *Int. J. Mod. Phys. A*, 28:1340011, 2013.
- [113] Michal P. Heller, Alekski Kurkela, Michal Spaliński, and Viktor Svensson. Hydrodynamization in kinetic theory: Transient modes and the gradient expansion. *Phys. Rev. D*, 97(9):091503, 2018.
- [114] W. Israel and J.M. Stewart. Transient relativistic thermodynamics and kinetic theory. *Annals Phys.*, 118:341–372, 1979.

- [115] G.S. Denicol, T. Koide, and D.H. Rischke. Dissipative relativistic fluid dynamics: a new way to derive the equations of motion from kinetic theory. *Phys. Rev. Lett.*, 105:162501, 2010.
- [116] Fernando G. Gardim, Giuliano Giacalone, Matthew Luzum, and Jean-Yves Ollitrault. Revealing QCD thermodynamics in ultrarelativistic nuclear collisions. *Nature Phys.*, 16(6):615–619, 2020.
- [117] Eric Braaten and Robert D. Pisarski. Soft Amplitudes in Hot Gauge Theories: A General Analysis. *Nucl. Phys. B*, 337:569–634, 1990.
- [118] J.P. Blaizot, Edmond Iancu, and A. Rebhan. The Entropy of the QCD plasma. *Phys. Rev. Lett.*, 83:2906–2909, 1999.
- [119] J.P. Blaizot, Edmond Iancu, and A. Rebhan. Selfconsistent hard thermal loop thermodynamics for the quark gluon plasma. *Phys. Lett. B*, 470:181–188, 1999.
- [120] J.P. Blaizot, Edmond Iancu, and A. Rebhan. Approximately selfconsistent resummations for the thermodynamics of the quark gluon plasma. 1. Entropy and density. *Phys. Rev. D*, 63:065003, 2001.
- [121] J.P. Blaizot, E. Iancu, and A. Rebhan. On the apparent convergence of perturbative QCD at high temperature. *Phys. Rev. D*, 68:025011, 2003.
- [122] Jens O. Andersen, Eric Braaten, and Michael Strickland. Hard thermal loop resummation of the thermodynamics of a hot gluon plasma. *Phys. Rev. D*, 61:014017, 2000.
- [123] Jens O. Andersen, Eric Braaten, and Michael Strickland. Screened perturbation theory to three loops. *Phys. Rev. D*, 63:105008, 2001.
- [124] Jens O. Andersen, Lars E. Leganger, Michael Strickland, and Nan Su. Three-loop HTL QCD thermodynamics. *JHEP*, 08:053, 2011.
- [125] Najmul Haque, Aritra Bandyopadhyay, Jens O. Andersen, Munshi G. Mustafa, Michael Strickland, and Nan Su. Three-loop HTLpt thermodynamics at finite temperature and chemical potential. *JHEP*, 05:027, 2014.
- [126] Owe Philipsen. The QCD equation of state from the lattice. *Prog. Part. Nucl. Phys.*, 70:55–107, 2013.
- [127] A. Bazavov et al. The QCD Equation of State to $\mathcal{O}(\mu_B^6)$ from Lattice QCD. *Phys. Rev. D*, 95(5):054504, 2017.
- [128] Szabolcs Borsanyi, Zoltan Fodor, Jana N. Guenther, Sandor K. Katz, Attila Pasztor, Israel Portillo, Claudia Ratti, and K.K. Szabó. Towards the equation of state at finite density from the lattice. *Nucl. Phys. A*, 982:223–226, 2019.
- [129] Sangyong Jeon. Hydrodynamic transport coefficients in relativistic scalar field theory. *Phys. Rev. D*, 52:3591–3642, 1995.
- [130] Sangyong Jeon and Laurence G. Yaffe. From quantum field theory to hydrodynamics: Transport coefficients and effective kinetic theory. *Phys. Rev. D*, 53:5799–5809, 1996.
- [131] Peter Brockway Arnold, Guy D. Moore, and Laurence G. Yaffe. Transport coefficients in high temperature gauge theories. 1. Leading log results. *JHEP*, 11:001, 2000.
- [132] Peter Brockway Arnold, Guy D Moore, and Laurence G. Yaffe. Transport coefficients in high temperature gauge theories. 2. Beyond leading log. *JHEP*, 05:051, 2003.

- [133] G. Policastro, Dan T. Son, and Andrei O. Starinets. The Shear viscosity of strongly coupled N=4 supersymmetric Yang-Mills plasma. *Phys. Rev. Lett.*, 87:081601, 2001.
- [134] M. Asakawa, T. Hatsuda, and Y. Nakahara. Maximum entropy analysis of the spectral functions in lattice QCD. *Prog. Part. Nucl. Phys.*, 46:459–508, 2001.
- [135] T. Yamazaki et al. Spectral function and excited states in lattice QCD with maximum entropy method. *Phys. Rev. D*, 65:014501, 2002.
- [136] Kiyoshi Sasaki, Shoichi Sasaki, and Tetsuo Hatsuda. Spectral analysis of excited nucleons in lattice QCD with maximum entropy method. *Phys. Lett. B*, 623:208–217, 2005.
- [137] Colin Morningstar. Bayesian curve fitting for lattice gauge theorists. *Nucl. Phys. B Proc. Suppl.*, 109A:185–191, 2002.
- [138] Yannis Burnier, Olaf Kaczmarek, and Alexander Rothkopf. The Bayesian reconstruction of the in-medium heavy quark potential from lattice QCD and its stability. *AIP Conf. Proc.*, 1701(1):100017, 2016.
- [139] Nicolai Christiansen, Michael Haas, Jan M. Pawłowski, and Nils Strodthoff. Transport Coefficients in Yang–Mills Theory and QCD. *Phys. Rev. Lett.*, 115(11):112002, 2015.
- [140] Juergen Berges, Nikolaos Tetradis, and Christof Wetterich. Nonperturbative renormalization flow in quantum field theory and statistical physics. *Phys. Rept.*, 363:223–386, 2002.
- [141] A. Andronic, P. Braun-Munzinger, K. Redlich, and J. Stachel. Statistical hadronization of charm in heavy ion collisions at SPS, RHIC and LHC. *Phys. Lett. B*, 571:36–44, 2003.
- [142] A. Andronic, P. Braun-Munzinger, K. Redlich, and J. Stachel. Statistical hadronization of heavy quarks in ultra-relativistic nucleus-nucleus collisions. *Nucl. Phys. A*, 789:334–356, 2007.
- [143] J. Stachel, A. Andronic, P. Braun-Munzinger, and K. Redlich. Confronting LHC data with the statistical hadronization model. *J. Phys. Conf. Ser.*, 509:012019, 2014.
- [144] Anton Andronic, Peter Braun-Munzinger, Krzysztof Redlich, and Johanna Stachel. Decoding the phase structure of QCD via particle production at high energy. *Nature*, 561(7723):321–330, 2018.
- [145] R.J. Fries, Berndt Muller, C. Nonaka, and S.A. Bass. Hadronization in heavy ion collisions: Recombination and fragmentation of partons. *Phys. Rev. Lett.*, 90:202303, 2003.
- [146] R.J. Fries, Berndt Muller, C. Nonaka, and S.A. Bass. Hadron production in heavy ion collisions: Fragmentation and recombination from a dense parton phase. *Phys. Rev. C*, 68:044902, 2003.
- [147] Tetsufumi Hirano and Yasushi Nara. Hydrodynamic afterburner for the color glass condensate and the parton energy loss. *Nucl. Phys. A*, 743:305–328, 2004.
- [148] Tetsufumi Hirano, Ulrich W. Heinz, Dmitri Kharzeev, Roy Lacey, and Yasushi Nara. Hadronic dissipative effects on elliptic flow in ultrarelativistic heavy-ion collisions. *Phys. Lett. B*, 636:299–304, 2006.

- [149] Akihiko Monnai and Tetsufumi Hirano. Effects of Bulk Viscosity at Freezeout. *Phys. Rev. C*, 80:054906, 2009.
- [150] Jean-Yves Ollitrault. Anisotropy as a signature of transverse collective flow. *Phys. Rev. D*, 46:229–245, 1992.
- [151] K. Dusling and D. Teaney. Simulating elliptic flow with viscous hydrodynamics. *Phys. Rev. C*, 77:034905, 2008.
- [152] Huichao Song and Ulrich W. Heinz. Causal viscous hydrodynamics in 2+1 dimensions for relativistic heavy-ion collisions. *Phys. Rev. C*, 77:064901, 2008.
- [153] Matthew Luzum and Paul Romatschke. Conformal Relativistic Viscous Hydrodynamics: Applications to RHIC results at $\sqrt{s_{NN}}^{*}(1/2) = 200$ -GeV. *Phys. Rev. C*, 78:034915, 2008. [Erratum: *Phys.Rev.C* 79, 039903 (2009)].
- [154] Sergei A. Voloshin, Arthur M. Poskanzer, and Raimond Snellings. *Collective phenomena in non-central nuclear collisions*, volume 23, pages 293–333. 2010.
- [155] B. Alver and G. Roland. Collision geometry fluctuations and triangular flow in heavy-ion collisions. *Phys. Rev. C*, 81:054905, 2010. [Erratum: *Phys.Rev.C* 82, 039903 (2010)].
- [156] Burak Han Alver, Clement Gombeaud, Matthew Luzum, and Jean-Yves Ollitrault. Triangular flow in hydrodynamics and transport theory. *Phys. Rev. C*, 82:034913, 2010.
- [157] Derek Teaney and Li Yan. Triangularity and Dipole Asymmetry in Heavy Ion Collisions. *Phys. Rev. C*, 83:064904, 2011.
- [158] Berndt Muller, Jurgen Schukraft, and Boleslaw Wyslouch. First Results from Pb+Pb collisions at the LHC. *Ann. Rev. Nucl. Part. Sci.*, 62:361–386, 2012.
- [159] Ulrich Heinz and Raimond Snellings. Collective flow and viscosity in relativistic heavy-ion collisions. *Ann. Rev. Nucl. Part. Sci.*, 63:123–151, 2013.
- [160] Betty Bezverkhny Abelev et al. Elliptic flow of identified hadrons in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *JHEP*, 06:190, 2015.
- [161] D. Everett et al. Phenomenological constraints on the transport properties of QCD matter with data-driven model averaging. *Phys. Rev. Lett.*, 126(24):242301, 2021.
- [162] D. Everett et al. Multisystem Bayesian constraints on the transport coefficients of QCD matter. *Phys. Rev. C*, 103(5):054904, 2021.
- [163] B. Betz, G.S. Denicol, T. Koide, E. Molnar, H. Niemi, and D.H. Rischke. Second order dissipative fluid dynamics from kinetic theory. *EPJ Web Conf.*, 13:07005, 2011.
- [164] G.S. Denicol, H. Niemi, E. Molnar, and D.H. Rischke. Derivation of transient relativistic fluid dynamics from the Boltzmann equation. *Phys. Rev. D*, 85:114047, 2012. [Erratum: *Phys.Rev.D* 91, 039902 (2015)].
- [165] G.S. Denicol, E. Molnár, H. Niemi, and D.H. Rischke. Derivation of fluid dynamics from kinetic theory with the 14-moment approximation. *Eur. Phys. J. A*, 48:170, 2012.
- [166] Azwinndini Muronga. Second order dissipative fluid dynamics for ultrarelativistic nuclear collisions. *Phys. Rev. Lett.*, 88:062302, 2002. [Erratum: *Phys.Rev.Lett.* 89, 159901 (2002)].

- [167] Azwinndini Muronga. Causal theories of dissipative relativistic fluid dynamics for nuclear collisions. *Phys. Rev. C*, 69:034903, 2004.
- [168] Michal P. Heller, Romuald A. Janik, and Przemyslaw Witaszczyk. Hydrodynamic Gradient Expansion in Gauge Theory Plasmas. *Phys. Rev. Lett.*, 110(21):211602, 2013.
- [169] Gabriel S. Denicol and Jorge Noronha. Analytical attractor and the divergence of the slow-roll expansion in relativistic hydrodynamics. *Phys. Rev. D*, 97(5):056021, 2018.
- [170] Michal P. Heller, Romuald A. Janik, and Przemyslaw Witaszczyk. The characteristics of thermalization of boost-invariant plasma from holography. *Phys. Rev. Lett.*, 108:201602, 2012.
- [171] M. Strickland. The non-equilibrium attractor for kinetic theory in relaxation time approximation. *JHEP*, 12:128, 2018.
- [172] Michael Strickland. The non-equilibrium attractor: Beyond hydrodynamics. *Acta Phys. Polon. B*, 50:1243, 2019.
- [173] Lanny Ray. The ridge(s) in STAR. *Nucl. Phys. A*, 854:89–96, 2011.
- [174] Sergei A. Voloshin. Transverse radial expansion in nuclear collisions and two particle correlations. *Phys. Lett. B*, 632:490–494, 2006.
- [175] E.V. Shuryak. On the origin of the 'Ridge' phenomenon induced by jets in heavy ion collisions. *Phys. Rev. C*, 76:047901, 2007.
- [176] Adrian Dumitru, Francois Gelis, Larry McLerran, and Raju Venugopalan. Glasma flux tubes and the near side ridge phenomenon at RHIC. *Nucl. Phys. A*, 810:91–108, 2008.
- [177] Piotr Bozek and Wojciech Broniowski. Correlations from hydrodynamic flow in p-Pb collisions. *Phys. Lett. B*, 718:1557–1561, 2013.
- [178] Piotr Bozek and Wojciech Broniowski. Collective dynamics in high-energy proton-nucleus collisions. *Phys. Rev. C*, 88(1):014903, 2013.
- [179] Adrian Dumitru, Kevin Dusling, Francois Gelis, Jamal Jalilian-Marian, Tuomas Lappi, and Raju Venugopalan. The Ridge in proton-proton collisions at the LHC. *Phys. Lett. B*, 697:21–25, 2011.
- [180] Kevin Dusling and Raju Venugopalan. Azimuthal collimation of long range rapidity correlations by strong color fields in high multiplicity hadron-hadron collisions. *Phys. Rev. Lett.*, 108:262001, 2012.
- [181] Kevin Dusling and Raju Venugopalan. Explanation of systematics of CMS p+Pb high multiplicity di-hadron data at $\sqrt{s_{NN}} = 5.02$ TeV. *Phys. Rev. D*, 87(5):054014, 2013.
- [182] Kevin Dusling and Raju Venugopalan. Comparison of the color glass condensate to dihadron correlations in proton-proton and proton-nucleus collisions. *Phys. Rev. D*, 87(9):094034, 2013.
- [183] Bjoern Schenke, Chun Shen, and Prithwish Tribedy. Hybrid Color Glass Condensate and hydrodynamic description of the Relativistic Heavy Ion Collider small system scan. *Phys. Lett. B*, 803:135322, 2020.
- [184] Björn Schenke, Chun Shen, and Prithwish Tribedy. Bulk properties and multi-particle correlations in large and small systems. *Nucl. Phys. A*, 1005:121756, 2021.

- [185] Shreyasi Acharya et al. Investigations of Anisotropic Flow Using Multiparticle Azimuthal Correlations in pp, p-Pb, Xe-Xe, and Pb-Pb Collisions at the LHC. *Phys. Rev. Lett.*, 123(14):142301, 2019.
- [186] Matthew D. Sievert and Jacquelyn Noronha-Hostler. CERN Large Hadron Collider system size scan predictions for PbPb, XeXe, ArAr, and OO with relativistic hydrodynamics. *Phys. Rev. C*, 100(2):024904, 2019.
- [187] V.N. Gribov and L.N. Lipatov. Deep inelastic e p scattering in perturbation theory. *Sov. J. Nucl. Phys.*, 15:438–450, 1972.
- [188] Guido Altarelli and G. Parisi. Asymptotic Freedom in Parton Language. *Nucl. Phys. B*, 126:298–318, 1977.
- [189] Yuri L. Dokshitzer. Calculation of the Structure Functions for Deep Inelastic Scattering and e+ e- Annihilation by Perturbation Theory in Quantum Chromodynamics. *Sov. Phys. JETP*, 46:641–653, 1977.
- [190] S. Moch, J.A.M. Vermaseren, and A. Vogt. The Three loop splitting functions in QCD: The Nonsinglet case. *Nucl. Phys. B*, 688:101–134, 2004.
- [191] A. Vogt, S. Moch, and J.A.M. Vermaseren. The Three-loop splitting functions in QCD: The Singlet case. *Nucl. Phys. B*, 691:129–181, 2004.
- [192] J.A.M. Vermaseren, A. Vogt, and S. Moch. The Third-order QCD corrections to deep-inelastic scattering by photon exchange. *Nucl. Phys. B*, 724:3–182, 2005.
- [193] Panagiota Foka and Malgorzata Anna Janik. An overview of experimental results from ultra-relativistic heavy-ion collisions at the CERN LHC: Hard probes. *Rev. Phys.*, 1:172–194, 2016.
- [194] Betty Abelev et al. Centrality Dependence of Charged Particle Production at Large Transverse Momentum in Pb–Pb Collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Phys. Lett. B*, 720:52–62, 2013.
- [195] R. Baier, Yuri L. Dokshitzer, Alfred H. Mueller, S. Peigne, and D. Schiff. Radiative energy loss of high-energy quarks and gluons in a finite volume quark - gluon plasma. *Nucl. Phys. B*, 483:291–320, 1997.
- [196] R. Baier, Yuri L. Dokshitzer, Alfred H. Mueller, S. Peigne, and D. Schiff. Radiative energy loss and p(T) broadening of high-energy partons in nuclei. *Nucl. Phys. B*, 484:265–282, 1997.
- [197] R. Baier, Yuri L. Dokshitzer, Alfred H. Mueller, and D. Schiff. Quenching of hadron spectra in media. *JHEP*, 09:033, 2001.
- [198] B.G. Zakharov. Fully quantum treatment of the Landau-Pomeranchuk-Migdal effect in QED and QCD. *JETP Lett.*, 63:952–957, 1996.
- [199] B.G. Zakharov. Radiative energy loss of high-energy quarks in finite size nuclear matter and quark - gluon plasma. *JETP Lett.*, 65:615–620, 1997.
- [200] B.G. Zakharov. Light cone path integral approach to the Landau-Pomeranchuk-Migdal effect. *Phys. Atom. Nucl.*, 61:838–854, 1998.
- [201] B.G. Zakharov. On the energy loss of high-energy quarks in a finite size quark - gluon plasma. *JETP Lett.*, 73:49–52, 2001.

- [202] R. Baier, D. Schiff, and B.G. Zakharov. Energy loss in perturbative QCD. *Ann. Rev. Nucl. Part. Sci.*, 50:37–69, 2000.
- [203] Miklos Gyulassy, Peter Levai, and Ivan Vitev. Jet quenching in thin quark gluon plasmas. 1. Formalism. *Nucl. Phys. B*, 571:197–233, 2000.
- [204] M. Gyulassy, P. Levai, and I. Vitev. NonAbelian energy loss at finite opacity. *Phys. Rev. Lett.*, 85:5535–5538, 2000.
- [205] M. Gyulassy, P. Levai, and I. Vitev. Reaction operator approach to nonAbelian energy loss. *Nucl. Phys. B*, 594:371–419, 2001.
- [206] Ivan Vitev and Miklos Gyulassy. High p_T tomography of $d + Au$ and $Au+Au$ at SPS, RHIC, and LHC. *Phys. Rev. Lett.*, 89:252301, 2002.
- [207] Miklos Gyulassy, Ivan Vitev, Xin-Nian Wang, and Ben-Wei Zhang. Jet quenching and radiative energy loss in dense nuclear matter. pages 123–191, 2 2003.
- [208] Carlos A. Salgado and Urs Achim Wiedemann. Calculating quenching weights. *Phys. Rev. D*, 68:014008, 2003.
- [209] Jorge Casalderrey-Solana and Carlos A. Salgado. Introductory lectures on jet quenching in heavy ion collisions. *Acta Phys. Polon. B*, 38:3731–3794, 2007.
- [210] Jaroslav Adam et al. Centrality dependence of high- p_T D meson suppression in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *JHEP*, 11:205, 2015. [Addendum: *JHEP* 06, 032 (2017)].
- [211] Jorge Casalderrey-Solana, Doga Can Gulhan, José Guilherme Milhano, Daniel Pablos, and Krishna Rajagopal. A Hybrid Strong/Weak Coupling Approach to Jet Quenching. *JHEP*, 10:019, 2014. [Erratum: *JHEP* 09, 175 (2015)].
- [212] Yuri L. Dokshitzer and D.E. Kharzeev. Heavy quark colorimetry of QCD matter. *Phys. Lett. B*, 519:199–206, 2001.
- [213] Jorge Casalderrey-Solana, Yacine Mehtar-Tani, Carlos A. Salgado, and Konrad Tywoniuk. New picture of jet quenching dictated by color coherence. *Phys. Lett. B*, 725:357–360, 2013.
- [214] Yacine Mehtar-Tani, Jose Guilherme Milhano, and Konrad Tywoniuk. Jet physics in heavy-ion collisions. *Int. J. Mod. Phys. A*, 28:1340013, 2013.
- [215] Jean-Paul Blaizot, Fabio Dominguez, Edmond Iancu, and Yacine Mehtar-Tani. Probabilistic picture for medium-induced jet evolution. *JHEP*, 06:075, 2014.
- [216] Jean-Paul Blaizot and Yacine Mehtar-Tani. Jet Structure in Heavy Ion Collisions. *Int. J. Mod. Phys. E*, 24(11):1530012, 2015.
- [217] Guang-You Qin and Xin-Nian Wang. *Jet quenching in high-energy heavy-ion collisions*, volume 24, page 1530014. 10 2015.
- [218] José Guilherme Milhano and Korinna Christine Zapp. Origins of the di-jet asymmetry in heavy ion collisions. *Eur. Phys. J. C*, 76(5):288, 2016.
- [219] Yacine Mehtar-Tani and Konrad Tywoniuk. Groomed jets in heavy-ion collisions: sensitivity to medium-induced bremsstrahlung. *JHEP*, 04:125, 2017.
- [220] Jean-Paul Blaizot, Fabio Dominguez, Edmond Iancu, and Yacine Mehtar-Tani. Medium-induced gluon branching. *JHEP*, 01:143, 2013.

- [221] Jean-Paul Blaizot, Edmond Iancu, and Yacine Mehtar-Tani. Medium-induced QCD cascade: democratic branching and wave turbulence. *Phys. Rev. Lett.*, 111:052001, 2013.
- [222] P. Caucal, E. Iancu, A.H. Mueller, and G. Soyez. Vacuum-like jet fragmentation in a dense QCD medium. *Phys. Rev. Lett.*, 120:232001, 2018.
- [223] Simon Turbide, Charles Gale, Sangyong Jeon, and Guy D. Moore. Energy loss of leading hadrons and direct photon production in evolving quark-gluon plasma. *Phys. Rev. C*, 72:014906, 2005.
- [224] Jaroslav Adam et al. Direct photon production in Pb-Pb collisions at $\sqrt{s}=2.76$ TeV. *Phys. Lett. B*, 754:235–248, 2016.
- [225] R. Baier, B. Pire, and D. Schiff. Dilepton production at finite temperature: Perturbative treatment at order α_s . *Phys. Rev. D*, 38:2814, 1988.
- [226] T. Altherr, P. Aurenche, and T. Becherawy. On Infrared and Mass Singularities of Perturbative QCD in a Quark - Gluon Plasma. *Nucl. Phys. B*, 315:436–464, 1989.
- [227] T. Altherr and P.V. Ruuskanen. Low mass dileptons at high momenta in ultrarelativistic heavy ion collisions. *Nucl. Phys. B*, 380:377–390, 1992.
- [228] Larry D. McLerran and T. Toimela. Photon and Dilepton Emission from the Quark - Gluon Plasma: Some General Considerations. *Phys. Rev. D*, 31:545, 1985.
- [229] Joseph I. Kapusta, P. Lichard, and D. Seibert. High-energy photons from quark - gluon plasma versus hot hadronic gas. *Phys. Rev. D*, 44:2774–2788, 1991. [Erratum: *Phys.Rev.D* 47, 4171 (1993)].
- [230] R. Baier, H. Nakkagawa, A. Niegawa, and K. Redlich. Production rate of hard thermal photons and screening of quark mass singularity. *Z. Phys. C*, 53:433–438, 1992.
- [231] P. Aurenche, F. Gelis, R. Kobes, and E. Petitgirard. Enhanced photon production rate on the light cone. *Phys. Rev. D*, 54:5274–5279, 1996.
- [232] P. Aurenche, F. Gelis, R. Kobes, and E. Petitgirard. Breakdown of the hard thermal loop expansion near the light cone. *Z. Phys. C*, 75:315–332, 1997.
- [233] P. Aurenche, F. Gelis, R. Kobes, and H. Zaraket. Bremsstrahlung and photon production in thermal QCD. *Phys. Rev. D*, 58:085003, 1998.
- [234] P. Aurenche, F. Gelis, and H. Zaraket. A Simple sum rule for the thermal gluon spectral function and applications. *JHEP*, 05:043, 2002.
- [235] P. Aurenche, F. Gelis, and H. Zaraket. Landau-Pomeranchuk-Migdal effect in thermal field theory. *Phys. Rev. D*, 62:096012, 2000.
- [236] Peter Brockway Arnold, Guy D. Moore, and Laurence G. Yaffe. Photon emission from ultrarelativistic plasmas. *JHEP*, 11:057, 2001.
- [237] Peter Brockway Arnold, Guy D. Moore, and Laurence G. Yaffe. Photon emission from quark gluon plasma: Complete leading order results. *JHEP*, 12:009, 2001.
- [238] Peter Brockway Arnold, Guy D. Moore, and Laurence G. Yaffe. Photon and gluon emission in relativistic plasmas. *JHEP*, 06:030, 2002.
- [239] P. Aurenche, F. Gelis, G.D. Moore, and H. Zaraket. Landau-Pomeranchuk-Migdal resummation for dilepton production. *JHEP*, 12:006, 2002.

- [240] Jacopo Ghiglieri, Juhee Hong, Aleksi Kurkela, Egang Lu, Guy D. Moore, and Derek Teaney. Next-to-leading order thermal photon production in a weakly coupled quark-gluon plasma. *JHEP*, 05:010, 2013.
- [241] M. Laine. NLO thermal dilepton rate at non-zero momentum. *JHEP*, 11:120, 2013.
- [242] F. Karsch, E. Laermann, P. Petreczky, S. Stickan, and I. Wetzorke. A Lattice calculation of thermal dilepton rates. *Phys. Lett. B*, 530:147–152, 2002.
- [243] Frithjof Karsch, S. Datta, E. Laermann, P. Petreczky, S. Stickan, and I. Wetzorke. Hadron correlators, spectral functions and thermal dilepton rates from lattice QCD. *Nucl. Phys. A*, 715:701–704, 2003.
- [244] H.-T. Ding, A. Francis, O. Kaczmarek, F. Karsch, E. Laermann, and W. Soeldner. Thermal dilepton rate and electrical conductivity: An analysis of vector current correlation functions in quenched lattice QCD. *Phys. Rev. D*, 83:034504, 2011.
- [245] Heng-Tong Ding, Olaf Kaczmarek, and Florian Meyer. Thermal dilepton rates and electrical conductivity of the QGP from the lattice. *Phys. Rev. D*, 94(3):034504, 2016.
- [246] J. Ghiglieri, O. Kaczmarek, M. Laine, and F. Meyer. Lattice constraints on the thermal photon rate. *Phys. Rev. D*, 94(1):016005, 2016.
- [247] Bastian B. Brandt, Anthony Francis, Tim Harris, Harvey B. Meyer, and Aman Steinberg. An estimate for the thermal photon rate from lattice QCD. *EPJ Web Conf.*, 175:07044, 2018.
- [248] Jurgen Berges, Klaus Reygers, Naoto Tanji, and Raju Venugopalan. Parametric estimate of the relative photon yields from the glasma and the quark-gluon plasma in heavy-ion collisions. *Phys. Rev. C*, 95(5):054904, 2017.
- [249] F. Gelis, H. Niemi, P.V. Ruuskanen, and S.S. Rasanen. Photon production from nonequilibrium QGP in heavy ion collisions. *J. Phys. G*, 30:S1031–S1036, 2004.
- [250] Chun Shen, Jean-Francois Paquet, Ulrich Heinz, and Charles Gale. Photon Emission from a Momentum Anisotropic Quark-Gluon Plasma. *Phys. Rev. C*, 91(1):014908, 2015.
- [251] Sigtryggur Hauksson, Sangyong Jeon, and Charles Gale. Photon emission from quark-gluon plasma out of equilibrium. *Phys. Rev. C*, 97(1):014901, 2018.
- [252] Anna Schäfer, Juan M. Torres-Rincon, Jonas Rothermel, Niklas Ehlert, Charles Gale, and Hannah Elfner. Benchmarking a nonequilibrium approach to photon emission in relativistic heavy-ion collisions. *Phys. Rev. D*, 99(11):114021, 2019.
- [253] Charles Gale and Joseph I. Kapusta. Dilepton radiation from high temperature nuclear matter. *Phys. Rev. C*, 35:2107–2116, 1987.
- [254] R. Rapp, G. Chanfray, and J. Wambach. Rho meson propagation and dilepton enhancement in hot hadronic matter. *Nucl. Phys. A*, 617:472–495, 1997.
- [255] Ralf Rapp and Edward V. Shuryak. Thermal dilepton radiation at intermediate masses at the CERN - SPS. *Phys. Lett. B*, 473:13–19, 2000.
- [256] R. Rapp. Signatures of thermal dilepton radiation at RHIC. *Phys. Rev. C*, 63:054907, 2001.
- [257] Hendrik van Hees and Ralf Rapp. Dilepton Radiation at the CERN Super Proton Synchrotron. *Nucl. Phys. A*, 806:339–387, 2008.

- [258] Ralf Rapp. Dilepton Spectroscopy of QCD Matter at Collider Energies. *Adv. High Energy Phys.*, 2013:148253, 2013.
- [259] Gojko Vujanovic, Clint Young, Bjoern Schenke, Ralf Rapp, Sangyong Jeon, and Charles Gale. Dilepton emission in high-energy heavy-ion collisions with viscous hydrodynamics. *Phys. Rev. C*, 89(3):034904, 2014.
- [260] Hans J. Specht. Thermal Dileptons from Hot and Dense Strongly Interacting Matter. *AIP Conf. Proc.*, 1322(1):1–10, 2010.
- [261] Ralf-Arno Tripolt, Lorenz von Smekal, and Jochen Wambach. Spectral functions and in-medium properties of hadrons. *Int. J. Mod. Phys. E*, 26(01n02):1740028, 2017.
- [262] Christopher Jung, Fabian Rennecke, Ralf-Arno Tripolt, Lorenz von Smekal, and Jochen Wambach. In-Medium Spectral Functions of Vector- and Axial-Vector Mesons from the Functional Renormalization Group. *Phys. Rev. D*, 95(3):036020, 2017.
- [263] Ralf-Arno Tripolt, Christopher Jung, Naoto Tanji, Lorenz von Smekal, and Jochen Wambach. In-medium spectral functions and dilepton rates with the Functional Renormalization Group. *Nucl. Phys. A*, 982:775–778, 2019.
- [264] Wei-jie Fu, Jan M. Pawłowski, and Fabian Rennecke. QCD phase structure at finite temperature and density. *Phys. Rev. D*, 101(5):054032, 2020.
- [265] Christopher Jung and Lorenz von Smekal. Fluctuating vector mesons in analytically continued functional RG flow equations. *Phys. Rev. D*, 100(11):116009, 2019.
- [266] T. Matsui and H. Satz. J/ψ Suppression by Quark-Gluon Plasma Formation. *Phys. Lett. B*, 178:416–422, 1986.
- [267] Yan-Qing Ma, Raju Venugopalan, and Hong-Fei Zhang. J/ψ production and suppression in high energy proton-nucleus collisions. *Phys. Rev. D*, 92:071901, 2015.
- [268] Yan-Qing Ma, Tomasz Stebel, and Raju Venugopalan. J/ψ polarization in the CGC+NRQCD approach. *JHEP*, 12:057, 2018.
- [269] F. Gelis, K. Kajantie, and T. Lappi. Chemical thermalization in relativistic heavy ion collisions. *Phys. Rev. Lett.*, 96:032304, 2006.
- [270] Francois Gelis and Naoto Tanji. Quark production in heavy ion collisions: formalism and boost invariant fermionic light-cone mode functions. *JHEP*, 02:126, 2016.
- [271] Naoto Tanji and Juergen Berges. Nonequilibrium quark production in the expanding QCD plasma. *Phys. Rev. D*, 97(3):034013, 2018.
- [272] Albert M Sirunyan et al. Measurement of nuclear modification factors of $\Upsilon(1S)$, $\Upsilon(2S)$, and $\Upsilon(3S)$ mesons in PbPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV. *Phys. Lett. B*, 790:270–293, 2019.
- [273] Serguei Chatrchyan et al. Observation of Sequential Upsilon Suppression in PbPb Collisions. *Phys. Rev. Lett.*, 109:222301, 2012. [Erratum: *Phys.Rev.Lett.* 120, 199903 (2018)].
- [274] Vardan Khachatryan et al. Suppression of $\Upsilon(1S)$, $\Upsilon(2S)$ and $\Upsilon(3S)$ production in PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Phys. Lett. B*, 770:357–379, 2017.
- [275] Yannis Burnier and Alexander Rothkopf. Bayesian Approach to Spectral Function Reconstruction for Euclidean Quantum Field Theories. *Phys. Rev. Lett.*, 111:182003, 2013.

- [276] Gert Aarts, Chris Allton, Tim Harris, Seyong Kim, Maria Paola Lombardo, Sinéad M. Ryan, and Jon-Ivar Skullerud. The bottomonium spectrum at finite temperature from $N_f = 2 + 1$ lattice QCD. *JHEP*, 07:097, 2014.
- [277] Atsuro Ikeda, Masayuki Asakawa, and Masakiyo Kitazawa. In-medium dispersion relations of charmonia studied by maximum entropy method. *Phys. Rev. D*, 95(1):014504, 2017.
- [278] Seyong Kim, Peter Petreczky, and Alexander Rothkopf. Quarkonium in-medium properties from realistic lattice NRQCD. *JHEP*, 11:088, 2018.
- [279] Gert Aarts and Jose Maria Martinez Resco. Transport coefficients, spectral functions and the lattice. *JHEP*, 04:053, 2002.
- [280] Nora Brambilla, Antonio Pineda, Joan Soto, and Antonio Vairo. Potential NRQCD: An Effective theory for heavy quarkonium. *Nucl. Phys. B*, 566:275, 2000.
- [281] Nora Brambilla, Antonio Pineda, Joan Soto, and Antonio Vairo. Effective Field Theories for Heavy Quarkonium. *Rev. Mod. Phys.*, 77:1423, 2005.
- [282] Nora Brambilla, Jacopo Ghiglieri, Antonio Vairo, and Peter Petreczky. Static quark-antiquark pairs at finite temperature. *Phys. Rev. D*, 78:014017, 2008.
- [283] M. Laine, O. Philipsen, P. Romatschke, and M. Tassler. Real-time static potential in hot QCD. *JHEP*, 03:054, 2007.
- [284] A. Beraudo, J.-P. Blaizot, and C. Ratti. Real and imaginary-time Q anti-Q correlators in a thermal medium. *Nucl. Phys. A*, 806:312–338, 2008.
- [285] Yannis Burnier, Olaf Kaczmarek, and Alexander Rothkopf. Static quark-antiquark potential in the quark-gluon plasma from lattice QCD. *Phys. Rev. Lett.*, 114(8):082001, 2015.
- [286] David Lafferty and Alexander Rothkopf. Improved Gauss law model and in-medium heavy quarkonium at finite density and velocity. *Phys. Rev. D*, 101(5):056010, 2020.
- [287] Szabolcs Borsányi, Zoltán Fodor, Sándor D. Katz, Attila Pásztor, Kálmán K. Szabó, and Csaba Török. Static $\bar{Q}Q$ pair free energy and screening masses from correlators of Polyakov loops: continuum extrapolated lattice results at the QCD physical point. *JHEP*, 04:138, 2015.
- [288] Guy D. Moore and Derek Teaney. How much do heavy quarks thermalize in a heavy ion collision? *Phys. Rev. C*, 71:064904, 2005.
- [289] Weiyao Ke, Yingru Xu, and Steffen A. Bass. Linearized Boltzmann-Langevin model for heavy quark transport in hot and dense QCD matter. *Phys. Rev. C*, 98(6):064901, 2018.
- [290] Xiaojun Yao, Weiyao Ke, Yingru Xu, Steffen Bass, and Berndt Müller. Quarkonium production in heavy ion collisions: coupled Boltzmann transport equations. *Nucl. Phys. A*, 982:755–758, 2019.
- [291] Robert L. Thews, Martin Schroedter, and Johann Rafelski. Enhanced J/ψ production in deconfined quark matter. *Phys. Rev. C*, 63:054905, 2001.
- [292] Jean-Paul Blaizot, Davide De Boni, Pietro Faccioli, and Giovanni Garberoglio. Heavy quark bound states in a quark-gluon plasma: Dissociation and recombination. *Nucl. Phys. A*, 946:49–88, 2016.

- [293] Xiaojian Du, Ralf Rapp, and Min He. Color Screening and Regeneration of Bottomonia in High-Energy Heavy-Ion Collisions. *Phys. Rev. C*, 96(5):054901, 2017.
- [294] Xiaojun Yao and Berndt Müller. Approach to equilibrium of quarkonium in quark-gluon plasma. *Phys. Rev. C*, 97(1):014908, 2018. [Erratum: *Phys.Rev.C* 97, 049903 (2018)].
- [295] Xiaojun Yao and Berndt Müller. Quarkonium inside the quark-gluon plasma: Diffusion, dissociation, recombination, and energy loss. *Phys. Rev. D*, 100(1):014008, 2019.
- [296] Jaroslav Adam et al. Differential studies of inclusive J/ψ and $\psi(2S)$ production at forward rapidity in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *JHEP*, 05:179, 2016.
- [297] Yukinao Akamatsu. Heavy quark master equations in the Lindblad form at high temperatures. *Phys. Rev. D*, 91(5):056002, 2015.
- [298] Shiori Kajimoto, Yukinao Akamatsu, Masayuki Asakawa, and Alexander Rothkopf. Dynamical dissociation of quarkonia by wave function decoherence. *Phys. Rev. D*, 97(1):014003, 2018.
- [299] Xiaojun Yao and Thomas Mehen. Quarkonium in-medium transport equation derived from first principles. *Phys. Rev. D*, 99(9):096028, 2019.
- [300] Nora Brambilla, Miguel A. Escobedo, Joan Soto, and Antonio Vairo. Heavy quarkonium suppression in a fireball. *Phys. Rev. D*, 97(7):074009, 2018.
- [301] Nora Brambilla, Miguel A. Escobedo, Joan Soto, and Antonio Vairo. Quarkonium suppression in heavy-ion collisions: an open quantum system approach. *Phys. Rev. D*, 96(3):034021, 2017.
- [302] Jean-Paul Blaizot and Miguel Angel Escobedo. Approach to equilibrium of a quarkonium in a quark-gluon plasma. *Phys. Rev. D*, 98(7):074007, 2018.