TOPICAL REVIEW • OPEN ACCESS

The exploration of hot and dense nuclear matter: introduction to relativistic heavy-ion physics

To cite this article: Hannah Elfner and Berndt Müller 2023 *J. Phys. G: Nucl. Part. Phys.* **50** 103001

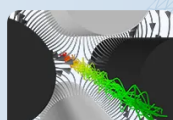View the article online for updates and enhancements.

Topical Review

The exploration of hot and dense nuclear matter: introduction to relativistic heavy-ion physics

Hannah Elfner[1,2,3,4,*] and **Berndt Müller**[5]

[1] GSI Helmholtzzentrum für Schwerionenforschung, Planckstr. 1, D-64291 Darmstadt, Germany
[2] Institut für Theoretische Physik, Goethe Universität, Max-von-Laue-Strasse 1, D-60438 Frankfurt am Main, Germany
[3] Frankfurt Institute for Advanced Studies, Ruth-Moufang-Strasse 1, D-60438 Frankfurt am Main, Germany
[4] Helmholtz Research Academy Hesse for FAIR (HFHF), GSI Helmholtz Center, Campus Frankfurt, Max-von-Laue-Straße 12, D-60438 Frankfurt am Main, Germany
[5] Department of Physics, Duke University, Durham, NC 27708, United States of America

E-mail: elfner@itp.uni-frankfurt.de

Abstract

This article summarizes our present knowledge about nuclear matter at the highest energy densities and its formation in relativistic heavy ion collisions. We review what is known about the structure and properties of the quark-gluon plasma and survey the observables that are used to glean information about it from experimental data.

Keywords: relativistic heavy ion collisions, quark-gluon plasma, hot nuclear matter

(Some figures may appear in colour only in the online journal)

1. Introduction

This section gives an overview of the goals of exploring nuclear matter at high energy density and outlines the recent and ongoing program with relativistic heavy ions.

---

* Author to whom correspondence should be addressed.

What are the properties of the matter that permeated our universe [1] during the first roughly 20 μs of its existence when its temperature exceeded 2×10^{12} K? Theoretical considerations tell us that it was a quark-gluon plasma, i.e. matter in which the quarks were not confined into color singlet objects, summarily called hadrons. Numerical simulations of quantum chromodynamics (QCD) on a Euclidean space-time lattice (lattice-QCD) have established the presence of a rapid, but smooth cross-over transition in the properties of QCD matter in the temperature range 140 MeV $< T <$ 170 MeV from a phase that is well described as a gas of hadrons and resonances to a phase, called quark-gluon plasma, in which the color force between quarks is screened and the spontaneously broken chiral symmetry of the QCD vacuum is restored [2, 3]. The near coincidence of these two transitions (quark deconfinement and chiral symmetry restoration) is due to their mutual reinforcement: absent the action of a confining color force, single quarks are easier to excite than whole hadrons, which suppresses the quark condensate and, in turn, enhances the density of quarks available to screen the color force. The connection is effectively modeled by the Polyakov–Nambu–Jona-Lasinio (PNJL) model [4, 5].

The experimental exploration of nearly baryon number-free hot QCD matter commenced in the year 2000 with the start of operations of the relativistic heavy ion collider (RHIC). In the past two decades, RHIC has produced and studied hot QCD matter in collisions of complex nuclei ranging from ^{16}O to ^{238}U, and it has performed baseline measurements in p +p and $p(d, {}^{3}\text{He})$ + Au collisions. Hot QCD matter with a significant excess of baryons over antibaryons ('baryon excess') was studied in fixed target experiments at the CERN-SPS and the BNL–AGS, and it has become a focus also at RHIC with the exploratory and high-statistics beam energy scans (BES-I and BES-II). At the lower energy end, SIS-18 at GSI is running and data are collected for very dense but moderate temperature systems. The conditions reached in low energy heavy-ion collisions are similar to the ones in neutron star mergers and provide complementary insights into nuclear matter at extreme densities (see [6] for a review on properties of dense nuclear matter).

Figure 1 shows a sketch of the phase diagram of strongly interacting matter. Putting firm constraints on the properties of matter is one of the major goals of heavy-ion research. The quark-gluon plasma phase is situated at high temperatures and/or densities with the cross-over transition to the hadron gas along the temperature axis. This is where the early universe has evolved and high energy heavy-ion experiments explore this net baryon number-free regime. At positive net baryon densities and low temperatures, there are multiple features in the phase diagram that have either been observed or are predicted. Among these is the liquid–gas transition between bound nuclear matter and a nucleon gas [7, 8], as well as phases possibly occurring in the cores of neutron stars, in which quark matter behaves as a color superconductor [9]. At finite net baryon density and high temperature a first-order phase transition with a critical endpoint is expected [10], which is explored by heavy-ion experiments at intermediate and low beam energies.

The results obtained during the first five years of the RHIC experiments have been summarized in 'white papers' [11–14] published by the four experimental collaborations (BRAHMS, PHENIX, PHOBOS, and STAR). The results provided compelling evidence that a novel form of thermal hot QCD matter was produced in nuclear collisions at RHIC, which clearly differs from hadronic gas and has the properties of a nearly inviscid liquid [15]. Since a low viscosity implies strong coupling among the constituents of the medium, a natural interpretation of the RHIC data is provided by the hypothesis that the QCD matter produced in the experiments is a strongly coupled quark-gluon plasma [16, 17].

The insights into the dynamical properties of the quark-gluon plasma garnered from the RHIC program were extensively tested and expanded at much higher beam energies at the
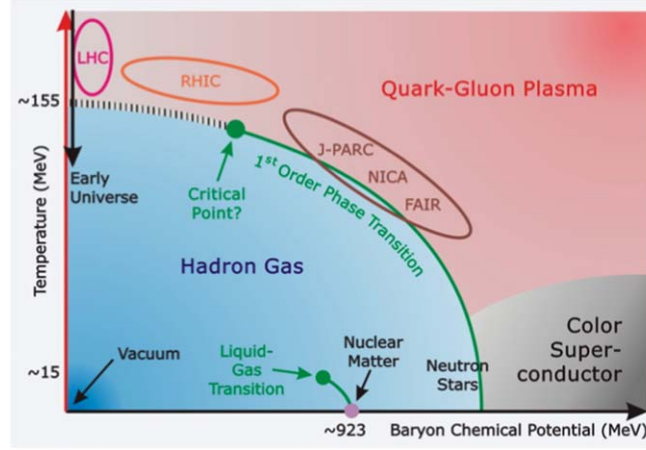
Figure 1. Sketch of the phase diagram of strongly interacting matter as a function of net baryon chemical potential and temperature. (Reproduced with permission from A. Steidl, Frankfurt University).

large hadron collider (LHC) starting in 2010. Since the LHC beam energy reaches up to a factor of 25 higher than that of RHIC, the main question was whether the quark-gluon plasma still exhibits the features of a strongly-coupled gauge liquid at these higher energies. This question was answered in the affirmative by the first reports from the three large experimental collaborations (ALICE [18], ATLAS [19], and CMS [20]). All three collaborations have since published a wealth of data that all confirm this conclusion (see [21] for a first comprehensive assessment). A more narrowly focused review of bulk properties of the QGP based on results from RHIC and LHC can be found in [22]. A high-level summary of established insights and open questions is presented in [23].

As will be reviewed below, various properties of the baryon number-free quark-gluon plasma can now be rigorously and reliably calculated using lattice-QCD simulations, in particular, its equation of state. However, definitive calculations are still limited to thermo-dynamic quantities that can be formulated in terms of static observables. Many properties of hot QCD matter that are of great phenomenological interest, such as its transport coefficients or its excitation spectrum, are still inaccessible to rigorous lattice simulations. Also, the properties of QCD matter containing a large net baryon density cannot be calculated using presently known techniques. For these properties, we still have to rely on tenuous extra-polations of results obtained by means of thermal perturbation theory, on predictions from effective models of QCD, or on answers gleaned from rigorous calculations for strongly coupled gauge theories akin to, but different from, QCD that have known gravity duals [24] (see section 3.3). The results of such model calculations must be regarded as qualitative, but they can provide useful guidance for the experimental program.

This review is structured as follows. Section 2 contains an overview of the connection between experimental observables and fundamental quantities characterizing the quark-gluon plasma. Section 3 covers the foundations of theoretical predictions for the properties of the quark-gluon plasma, based on lattice QCD, thermal perturbation theory, and holographic methods for strongly coupled gauge theories that serve as models for QCD. The basic phenomenological approaches to making this connection that are well established on a quantitative level are described in section 4. In section 5, we review the experimental results

from RHIC and LHC for the observables introduced in section 2 and discuss their interpretation. The review concludes in section 6 with an overview of future opportunities for experimental and theoretical investigations of hot QCD matter.

How to read this article: If the reader is mainly interested in getting an overview of the current state of experimental results, it would be advisable to jump directly to section 5 and then circle back to sections 2, 3, and 4 depending on the nature of one's further interest. Alternatively, one might start with section 2 if the primary interest is in getting an overview of the main experimental observables in relativistic heavy ion collisions, for example, if someone wants to prepare for the first attendance of a quark matter conference and then proceed to either sections 4 or 3. Those most interested in the theoretical motivation and foundation of the field may want to start with section 3 and then proceed to section 4 and the other sections. We assume that any reader will have at least a cursory familiarity with the elements of QCD and its most important phenomenological aspects (asymptotic freedom, quark confinement, hadron spectrum).

Our review does not aim for completeness in citations. The published literature on the quark-gluon plasma and the phenomenology of relativistic heavy ion collisions is just too vast. An excellent source of foundational references is Kapusta *et al*'s annotated reprint collection [25]. Concise presentations of many of the underlying concepts and phenomenological aspects can be found in the textbooks by Yagi *et al* [26] and by Letessier and Rafelski [27]. Apart from specific citations of quoted experimental or theoretical results, we have mainly cited targeted reviews and publications that provide an overview of a specific area.

We hope that this citation strategy will be of value to the readers at whom our review is primarily aimed: Graduate students and postdocs, scientists from other areas of nuclear and particle physics, who want to familiarize themselves with the status of this field of physics because they are considering contributing to it, and all those scientists who would like to get a concise overview of the questions, methods and results of this area of physics. We refrain on purpose from historical accounts of events and rather report the current status of understanding of hot and dense QCD matter as explored in heavy ion collisions.

Notations: As we just explained, our review is not primarily intended as a self-contained survey of the field of relativistic heavy ion physics as you would find in a textbook, but as an introduction and overview that helps guide aspiring practitioners towards the major trends of research. For this reason, we have made an attempt in each subsection to use the notation that is commonly found in the pertinent research literature. As sometimes the same symbol is used to denote different quantities in different areas—e.g. σ can denote the electrical or color conductivity, or the cross section, depending on the context—our approach necessarily leads to a certain amount of inconsistency across sections and subsections. In order to avoid confusing the reader, we have made an attempt to be clear about the meaning of symbols at the point where they are used, but we have refrained from providing a comprehensive list of symbols.

Another possible source of confusion is the difference between physical units used to report data—with allowance for the use of nonstandard units used for quantities in the nuclear and subnuclear domain, e.g. MeV and GeV instead of Joules for energy—and 'natural' units that are widely used in theoretical calculations, which use the convention $\hbar = c = 1$. Most of the formulas in our review make use of natural units, but we make an attempt to report data in the units that are customarily used in the respective area.

One instance of widespread confusion is the difference of units for energy and momentum, which are the same in natural units but differ in physical units (momentum is measured in GeV/c; energy is measured in GeV). The difference is subtle, because some particle detectors

(e.g. tracking detectors in a magnetic field) measure particle momentum, while others (calorimeters) measure the kinetic energy of a particle. Similarly, the center-of-mass collision energy in a particle collider is usually reported in units of energy, while the beam energy for fixed target experiments is usually stated as the beam momentum. We have made an attempt to be clear about these differences in notation and apologize to the attentive reader who detects remaining inconsistencies.

2. Observables

This section describes basic quantities characterizing high-energy density QCD matter and their relation to experimental observables. This section contains very few references, since the details are discussed in section 5 where references, mainly to review articles, are provided.

2.1. Overview

The approaches described in the previous section allow theorists to relate various fundamental properties of the quark-gluon plasma to specific experimental observables:

  (i) The equation of state of the quark-gluon plasma encodes the dependence of the pressure p on the two intensive variables energy density ε and net baryon density $n = n_b - n_{\bar{b}}$, where n_b ($n_{\bar{b}}$) denotes the density of baryons (antibaryons). It is reflected in the final particle spectra, collective flow properties, and the propagation of initial-state density fluctuations into the final state. The speed of sound, $c_s = \sqrt{\partial p/\partial \varepsilon}$, determines how the pressure gradients propagate during the expansion of the system.
 (ii) The viscosities η (shear) and ζ (bulk), which may be expressed as autocorrelation functions of the components of the stress tensor T_{ij} [28],

$$\eta = \frac{1}{T} \int d^4x \, \langle T_{xy}(x) T_{xy}(0) \rangle, \tag{1}$$

$$\zeta = \frac{1}{9T} \sum_{i,j} \int d^4x \, \langle T_{ii}(x) T_{jj}(0) \rangle. \tag{2}$$

influence the collective transverse flow pattern, quantified by Fourier decomposition of the azimuthal distribution of particles in momentum space. The anisotropy of the transverse flow is especially sensitive to the dimensionless ratio η/s, where s denotes the entropy density.
(iii) Momentum transport coefficients, including the so-called jet quenching parameter $\hat{q}$, are expressed as correlation functions of the color force experienced by a fast parton in the quark-gluon plasma:

$$\hat{q} \propto \int dx^- \langle \, \mathrm{Tr} \, [U_{0,x} F^{+i}(x) U_{x,0} F_i^+(0)] \rangle. \tag{3}$$

Here F denotes the field strength tensor and U are so-called gauge links that ensure the gauge invariance of $\hat{q}$. The minus index in x^- indicates that the correlation is sampled along the light-cone. Together with the coefficient $\hat{e}$, which quantifies the rate of longitudinal momentum loss due to collisions with medium constituents [29], $\hat{q}$ determines the rate of energy loss of a fast parton traversing the quark-gluon plasma. Experimentally, these coefficients can be determined by measuring the suppression of hadrons emitted with high momentum transverse to the beam direction.

(iv) The Debye screening length λ_D of the color force between a heavy quark-antiquark pair in the quark-gluon plasma, which is determined by the correlator of the static gauge potential between two points separated by a spatial distance x:

$$\lambda_D^{-1} = -\lim_{|x|\to\infty} \frac{\langle \text{Tr}\,[U_{0,x}A^0(x)U_{x,0}A^0(0)]\rangle}{|x|}, \tag{4}$$

controls the ability of a pair of heavy quarks to form a bound state. When λ_D is less than the size of the bound state, the state dissolves in the medium. Experimentally, this phenomenon is expected to be revealed by a strong suppression of charmonium ($c\bar{c}$) and bottomonium ($b\bar{b}$) states.

(v) The electromagnetic current-current correlation function

$$C_{\text{em}}^{\mu\nu}(q) = \int d^4x\,e^{iq\cdot x}\langle j^\mu(x)j_\nu(0)\rangle \tag{5}$$

encodes the response of the quark-gluon plasma (or hadronic medium) to electromagnetic fields. It is directly only sensitive to quarks, but indirectly also to gluons, because quarks can be off-shell due to their interactions with thermal gluons. Observables carrying information about $C^{\mu\nu}(q)_{\text{em}}$ are photons, sensitive to $q^2 = 0$, and dileptons, sensitive to time-like $q^2 > 4m^2$ with m being the lepton mass.

These well-defined quantities summarize the main properties of the hot and dense QCD matter. Their linkage to experimental observables lays the ground for robust theoretical and experimental research programs in heavy ion physics. In making this linkage it is important to keep in mind that only the momentum space properties of the matter are accessible to experiments, because the fireball is too small and short-lived to allow direct spatial of temporal resolution. It is possible to infer some information about the space- and time-dependence of the collision from two-particle correlations, which will be discussed in section 2.4 under the topic 'HBT Femtoscopic Radii'[6], but even these indirect measurements are limited to the last stage of the collision and require detailed modeling of the reaction.

From a practical viewpoint, on the experimental side, it is important that precision measurements of the relevant observables are made over a wide kinematic range so that theoretical models of the collision can be constrained. On the theoretical side, precise and reliable calculations of the transport coefficients under given conditions are needed in combination with realistic simulations of the dynamical evolution that connect the matter properties to the observables.

The following sections introduce the basic concepts underpinning the most commonly considered observables and point out their generally accepted connections to the properties of hot and dense QCD matter.

2.2. Single-particle observables

Single-particle observables are all the observables that can be measured by a detector that records only individual particles. Below we discuss a few specific observables of this type that are of interest in heavy ion physics.

Particle yields and multiplicities

The most basic analysis in a heavy ion collision is to count the number of particles of a certain species that are produced. One can either restrict the measurement of the yield to a

---

[6] The acronym "HBT" stands for the names Hanbury-Brown and Twiss, who first proposed identical particle density interferometry as a probe of the size of remote objects (stars) [30]. 'Femtoscopy' [31] refers to the measurement of the size of objects on the femtometer (1 fm = 10^{-15} m) scale.

region in momentum space around mid-rapidity or, if detector acceptance allows, extrapolate the number to the full multiplicity including the whole phase-space (4π yields). By fitting the ratios of particle multiplicities to a grand canonical thermal distribution one can infer information about the conditions at the time of chemical freeze-out, after which the particle yields do not change anymore (see section 5.1).

Rapidity distributions

Rapidity distributions $\mathrm{d}N/\mathrm{d}y$ are sensitive to the initial energy deposition and thus to the stopping power of the colliding nuclei. Extreme limits of the charged particle rapidity distribution are associated with Landau (Gaussian shape) or Bjorken (flat shape) hydrodynamic behavior. The Landau limit corresponds to full stopping followed by explosive expansion. The Bjorken model corresponds to complete transparency for the valence quarks resulting in longitudinal boost invariance—this approximation is only strictly valid in the infinite collision energy limit.

Transverse momentum spectra

Counting the particles in bins of transverse momentum p_T or transverse mass $m_T = \sqrt{m^2 + p_T^2}$ conveys information about the kinetic decoupling temperature and the flow profile

$$\frac{1}{m_T}\frac{\mathrm{d}N}{\mathrm{d}m_T} \propto \exp^{-m_T/T_{\mathrm{eff}}}. \tag{6}$$

By fitting the particle spectra to a thermal distribution one can extract an effective 'temperature' T_{eff}, which is correctly called the *slope parameter*. If the emitted hadrons originate from a locally equilibrated and collectively flowing momentum distribution, the slope parameter represents a blue-shifted temperature. The outward radial flow that the fireball develops during the dynamical evolution generally results in a flatter shape of the spectra of higher mass particles, e.g. baryons, at low transverse momentum (see section 5.2).

Nuclear modification factor R_{AA}

The nuclear modification factor R_{AA} is the ratio of the yield at a given transverse momentum p_T and rapidity y measured in a heavy ion collision (AA) normalized by the yield measured in the corresponding number of independent nucleon–nucleon collisions (p+p)

$$R_{AA}(p_T, y; \langle b\rangle) = \frac{\mathrm{d}^2N/\mathrm{d}p_T\,\mathrm{d}y_{AA}}{\langle N_{coll}\rangle\,\mathrm{d}^2N/\mathrm{d}p_T\,\mathrm{d}y_{p+p}}. \tag{7}$$

Here $\langle N_{\mathrm{coll}}\rangle$ denotes the average number of binary nucleon–nucleon collisions in heavy ion collisions within a certain centrality window; $\langle b\rangle$ is the average impact parameter associated with this window. In any given collision N_{coll} reflects the geometric overlap of the colliding nuclei but fluctuates around the average value by virtue of quantum fluctuations in the nuclear densities and the nucleon–nucleon scattering probability. Sometimes good reference data from p+p collisions do not exist. In such cases, one often takes peripheral heavy ion collisions as a reference point and uses the similarly defined binary collision-weighted ratio of yields measured in central and peripheral collisions, denoted as R_{CP}.

If there are final-state interactions with the medium that degrade the transverse momentum of energetic partons, R_{AA} falls below unity, which explains the alternative name of this observable: nuclear suppression factor.[7] This phenomenon is referred to as jet quenching. For

---

[7] Note that R_{AA} can also exceed unity, e.g. if multiple collisions of a single nucleon impart additional transverse momentum (Cronin effect) or if the emitted particles are pushed to higher-pT by the collective flow.

light-heavy collision systems, such as $p + A$ or $d + A$, the $R_{p(d)A}$ is sensitive to initial-state effects that can be attributed to the behavior of cold nuclear matter relative to the reference of proton–proton collisions. For heavy-heavy collision systems the nuclear modification factor is mainly sensitive to final-state effects. The first determination of the energy loss parameter $\hat{q}$ is based on measurements of the nuclear modification factor (see section 5.9). R_{AA} is also used to quantify the nuclear suppression or enhancement (with respect to p–p collisions) of many other observables in heavy ion collisions, such as the yields of heavy quarkonia, heavy-flavor hadron, photons, etc.

Photon spectra

Direct photon spectra reflect the cumulative emission of thermal and prompt photons from all stages of a heavy ion reaction. The definition of a 'direct photon' is that it does not originate from a particle decay, such as $\pi^0 \to \gamma\gamma$. In the initial non-equilibrium evolution high-p_T photons are produced in hard processes. Thermal emission from the quark-gluon plasma and the subsequent hadron gas phase are added during later collision stages. Since the mean free path of photons is larger than the fireball size, once photons are produced they end up in the detector and, different from strongly interacting probes, are minimally distorted by rescattering processes. By detailed dynamical modeling and matching of the various contributions to the photon spectrum, the goal is to extract information about the initial temperature and the lifetime of the QGP fireball from this observable. Photon spectra are also sensitive to the equation of state, since the emission depends strongly on the lifetime of the system that is substantially longer when a first-order phase transition takes place (see section 5.11).

Dilepton spectra

Another electromagnetic probe are lepton pairs, commonly referred to as dileptons. These originate from virtual photons that decay into a e^+e^- or $\mu^+\mu^-$ pair. The theoretical advantage is that leptons do not interact strongly and therefore can be measured from all stages of the heavy ion reaction. On the other hand, the cross-sections for leptonic pair production are orders of magnitude smaller than for hadronic emission. Since some hadron decays (e.g. $\pi^0 \to \gamma e^+e^-$, $D^0 \to K^-\mu^+\nu$) produce leptons which generates a large combinatorial background of uncorrelated lepton pairs, dilepton spectra constitute challenging measurements. Vector mesons, such as the ρ and J/ψ mesons, can decay into lepton pairs and show up as peaks in the dilepton invariant mass spectrum. By carefully analyzing the height and width of these peaks, one hopes to learn something about the influence of the medium on the resonance properties. Changes in the spectral function of the ρ meson are associated with chiral symmetry restoration and deconfinement (see section 5.11).

Charmonium suppression

One telltale signature for the formation of quark-gluon plasma is the suppression of charmonia (J/Ψ, Υ, etc). While the number of heavy quark pairs created in initial hard interactions is deemed to be independent of the final-state nuclear medium by virtue of QCD factorization (see section 4.2), they are predicted to be unable to form bound meson states in a quark-gluon plasma due to color screening. Lattice calculations for the heavy quark potential and the spectral function predict a sequential melting of charmonia states with increasing medium temperature. Therefore, the measurement of the ratio of charm spectra in nucleus-nucleus collisions divided by the spectrum expected from scaled proton–proton collisions is thought to serve as a direct way to extract the deconfinement temperature. This simplistic picture is complicated due to recombination of heavy quarks during hadronization and other effects (see section 5.10).

Heavy flavor

The mass dependence of medium interactions can be probed by measuring the spectra of hadrons containing unpaired c or b-quarks (open charm or beauty hadrons) or those of single leptons (electrons or muons) that originate from the decay of these hadrons. High statistics measurements of D-mesons identified by their hadronic decays have become the most precise tool for this purpose. Due to the dead-cone effect the expectation is that energetic charm and bottom quarks lose less energy in the quark-gluon plasma than light quarks. Measurements of the azimuthal angular distribution of open heavy-flavor mesons can provide insight into the question whether heavy quarks thermalize similar to light quarks (see section 5.10).

2.3. Two- and few-particle observables

Two- and few-particle correlations are coincidence measurements of the probability of finding a particle with a certain property given that the event contains one particle with another (or the same) property.

Photon-hadron correlations

Photon-hadron correlations use a direct photon as the trigger particle for the correlation function. This has the advantage that the energy of the original parton in the hard process is constrained more tightly by the photon energy and, therefore, allows to access information about the fragmentation function of energetic partons, which measures the fraction z of the primary parton momentum carried by a secondary hadron. This observable is quite clean for theoretical calculations, but poses a big challenge to experiments due to the contamination of the photon yield by decay photons (mainly from π^0) (see section 5.9).

Di-hadron correlations

Hard di-hadron correlations are a sensitive tool for studying the path length dependence of jet energy loss. The trigger particle is a high-p_T particle (usually with $p_T > 5\,\mathrm{GeV/c}$), whereas the associate particle is chosen in a lower p_T range. Di-hadron correlations are called 'hard', if both p_T ranges are $3\,\mathrm{GeV/c}$ and higher, so that a perturbative QCD (pQCD) description is applicable. These correlations are often quantified as a distribution in the difference of the azimuthal angle between trigger and associated particle. One can then separate the near-side and the away-side structures, by an angle difference of roughly $180°$ that originate from the back-to-back hard parton scattering. Given a near-side high-p_T trigger particle the ratio of the yield of particles on the away-side in heavy ion collisions to the one in proton–proton collisions is called I_{AA} in analogy to R_{AA}. Since I_{AA} is a conditional yield, it potentially contains more detailed information about the energy loss, but also is subject to the so-called 'trigger bias' (see section 5.9).

Triggered correlations

In addition to looking for correlations between two high-p_T particles, jet-medium inter-actions can be investigated by studying correlations between only one high-p_T particle and an associated particle of lower p_T. The high-p_T trigger particle preferentially emerges from the surface of the medium, having lost little energy on its way out, and therefore the backwards emitted parton experiences above average modification by the medium. Also, the medium itself might be modified by the traversing high energy parton resulting in characteristic excitation patterns like a Mach cone. To measure the effect of high p_T partons on the medium, multi-particle correlations over various p_T ranges and in longitudinal ($\Delta\eta$) and azimuthal ($\Delta\phi$) phase space have been studied.

Untriggered correlations

The last option for two- or three-particle angular correlations is to measure so-called untriggered correlations in minimum bias events, where one does not require one of the hadrons to have high transverse momentum. These $\mathrm{d}^2 N / \mathrm{d}\Delta\phi\Delta\eta$ correlations have revealed interesting structure: there is an elongated enhancement in pseudorapidity on the near-side $\Delta\phi \approx 0$) that is often referred to as the 'ridge'. This structure is associated with higher Fourier coefficients in the azimuthal distribution with respect to the reaction plane (see below under collective flow) and will be discussed further in section 5.5.

Charge asymmetry correlations

Another 3-particle measurement is the charge asymmetry with respect to the reaction plane. Because up and down quarks carry different electric charges ($+2e/3$ and $-e/3$, respectively) they could be separated by the very strong transient magnetic field that is generated by the colliding nuclei. The chiral magnetic effect (see section 5.13 for more details) is predicted to result in a charge asymmetry of the yield of charged particles with respect to the reaction plane. This asymmetry can be detected by observables of the type:

$$\gamma^{\pm,\pm} = \langle \cos(\phi_\alpha^\pm + \phi_\beta^\pm - 2\Psi_{\mathrm{RP}}) \rangle, \tag{8}$$

where ϕ_α, ϕ_β denotes the azimuthal emission angles of the two particles, Ψ_{RP} denotes the reaction plane defined by the impact parameter vector, and $\pm$ indicates the sign of the electric charge of the detected hadron. The observable (8) requires the detection of at least three particles, one of which is used to identify the reaction plane. A similar asymmetry of charged particle production with respect to the reaction plane can also be generated by a combination of balance functions and elliptic flow (see section 5.14), which means that the observable has a large non-specific background.

2.4. Multi-particle and collective observables

This section contains all of the observables that are based on multi-particle measurements and potentially reflect the collective behavior of the system.

Anisotropic flow

One of the most important observables in heavy ion reactions is collective flow of the particles. The Fourier coefficients v_n of the azimuthal distribution of the final particles in momentum space are used to quantify anisotropic flow.

$$E\frac{\mathrm{d}^3 N}{\mathrm{d}p^3} = \frac{1}{2\pi} \frac{\mathrm{d}^2 N}{p_t \, \mathrm{d}p_t \, \mathrm{d}y} + \sum_{n=1}^{\infty} 2v_n \cos[n(\phi - \Psi_n)], \tag{9}$$

whee ϕ denotes the azimuthal angle around the beam axis, and Ψ_n stands for the angle defined by the maximum of the nth Fourier component of the angular distribution in a given event.

The most prominent and well-studied of these coefficients is v_2, commonly called elliptic flow, which arises due to the different pressure gradients in the transverse plane in non-central collisions. The initial almond shape in the coordinate space is translated by hydrodynamics into a momentum space anisotropy. As it is driven by pressure gradients, elliptic flow is highly sensitive to the shear viscosity and the equation of state of the expanding medium. Its collision energy dependence can, therefore, serve as a signature of the transition from hadronic matter to the quark-gluon plasma. Whereas v_2 can be measured by an event average, higher-order flow coefficients are non-zero require event-by-event measurements. They can be used to constrain the initial state profile of the fireball. The v_n values can be determined by

different analysis techniques like the event-plane method, the cumulant method, the flow vector method, or by means of Lee-Yang zeros (see section 5.4).

HBT femtoscopic radii

Hanbury-Brown–Twiss (HBT) correlations of identical particles allow to infer the spatial volume, lifetime, and outward flow velocity of the fireball. The HBT correlations are caused by quantum (Bose or Fermi) interference between the two identical particles. The measurements are sometimes referred to as density interferometry. The measured correlation functions are usually fitted with a Gaussian that allows to extract three different radii

$$C(\mathbf{q}) = 1 + \lambda \mathrm{e}^{-(q_o^2 R_{\mathrm{out}}^2 + q_s^2 R_{\mathrm{side}}^2 + q_l^2 R_{\mathrm{long}}^2)}, \tag{10}$$

where R_{long} is aligned with the beam direction, R_{out} points along the center-of-momentum of the particle pair, and R_{side} is perpendicular to both. The vector $\mathbf{q}$ denotes the difference between the momenta of the two identical particles, and q_o, q_s, q_l are its components with respect to these three orthogonal directions. The HBT radii for different particle species contain information about the coherence region of the emission at kinetic freeze-out. The ratio of $R_{\mathrm{out}}/R_{\mathrm{side}}$ is indicative of the lifetime of the fireball and therefore sensitive to certain aspects of the equation of state, such as the nature of the QCD phase transition. Results and more details are discussed in section 5.3.

Balance functions

Balance functions are a tool to study charge correlations in heavy ion reactions. They are similar to observables used to investigate hadronization in jets produced in $p\bar{p}$ or e^+e^- collisions. The balance function describes the conditional probability that a particle in the momentum bin p_1 will be accompanied by a particle of opposite charge in the momentum bin p_2

$$\begin{aligned} B_{ab}(p_2|p_1) = \frac{1}{2}(&\rho(b, p_2|a, p_1) - \rho(b, p_2|b, p_1) \\ + &\rho(a, p_2|b, p_1) - \rho(a, p_2|a, p_1), \end{aligned} \tag{11}$$

where $\rho(b, p_2|a, p_1)$ is the conditional probability of observing a particle of type b in bin p_2 given the existence of a particle of type a in bin p_1. Balance functions can be defined for any conserved quantum number.

Reconstructed jets

Another observable sensitive to parton propagation in the quark-gluon plasma are fully reconstructed jets. For p+p collisions with usually only a single hard interaction the outgoing hard partons fragment into hadrons, and the energies of all fragment particles are measured in calorimeters. Applying sophisticated clustering algorithms the full jet can be reconstructed. Events containing a jet are identified online using a dedicated jet trigger. In the high-multiplicity environment of a heavy ion collision these measurements are much more challenging due to the underlying event background. Reconstructed jets allow comprehensive studies of jet energy loss and modifications of the internal jet structure (fragmentation functions, jet shape, etc) by the comparison with predictions from jet shower Monte-Carlo algorithms tuned to jets in p+p collisions (see section 5.9).

Energy flow and track functions

Calorimeters measure energy flow carried by hadrons (hadronic calorimeter) or by photons and electrons (electromagnetic calorimeter) in a given direction. The most general realization of this concept is the energy flow operator [32]

$$\mathcal{E}(\vec{n}) = \lim_{r \to \infty} \int_0^\infty \mathrm{d}t \ r^2 n^i T_i^0(t, r\vec{n}),\tag{12}$$

which measures energy flow into the direction $\vec{n}$. The T_i^0 are the components of the energy-momentum tensor T_ν^μ related to momentum or energy flow. (For photons these are the components of the Poynting vector.) Observables of interest that provide a measure of the energy correlators of the form $\langle \mathcal{E}(\vec{n}_1) \cdots \mathcal{E}(\vec{n}_N) \rangle$, where one looks at the correlation of energy flow within a jet into different directional domains with an opening solid angle ΔR [33].

A related concept is that of track functions $T_{q \to h}(x)$, [34, 35] which measure the probability that the total (light-cone) momentum fraction x of a quark-initiated jet is carried by a certain type of hadron h. The track functions are similar to fragmentation functions, but instead of considering the momentum fraction of a single hadron, they sum over all hadrons of a given type within the jet. For a recent analysis of the QCD evolution of track function moments and a comparison with LHC data, see [36].

Event-by-event fluctuations

Event-by-event fluctuations of conserved quantum numbers, such as net electrical charge Q, net baryon number B, or net strangeness S, are a prominent signal for the phase transition to the quark-gluon plasma. Since quarks and gluons have different elementary units of the charges the predictions for event-by-event fluctuations are very different based on the assumption that a quark-gluon plasma or a hadron gas is present. The fluctuations of the mean transverse momentum can be regarded as the analog to the temperature fluctuations in the cosmic microwave background, which are vestiges of quantum fluctuations in the initial state. The susceptibilities that are directly related to the fluctuations of conserved quantum numbers or their higher moments such as the skewness and the kurtosis are also calculable on the lattice and offer a direct connection between data and QCD predictions (see section 5.7).

3. Theoretical foundations

This section reviews our present theoretical understanding of the structure and properties of the quark-gluon plasma, based on lattice-QCD simulations, thermal effective field theory, and exactly solvable strong coupling models.

3.1. Lattice QCD

The quantitative *ab initio* calculation of the thermodynamic properties of QCD matter requires the model-independent evaluation of the functional integral that defines the quantum field theory. The only known rigorous method that allows to do this starts with the discretization of the QCD Lagrangian on a large space-time lattice and then evaluates the very high, but finite-dimensional functional integral by Monte-Carlo methods. In order to study the properties of QCD at non-zero temperature, periodic boundary conditions (anti-periodic for quarks) are imposed on the imaginary time coordinate with period $\hbar/T$. The functional integral defining the quantum field theory is evaluated by means of Monte-Carlo sampling techniques that are used to generate a sufficient number (hundreds or thousands) of statistically independent, representative field configurations. Observables of interest are then evaluated by averaging over these stored configurations. The numerically most expensive part of the calculation is the integration over the fermion fields, which requires the evaluation of the determinant of a very large matrix. Much of the progress that has occurred over the past decade consists of finding improved ways to represent the QCD Lagrangian on a discrete lattice and speed up the evaluation of the fermion determinant. Other improvements in some

simulations concern the implementation of manifest chiral symmetry by means of the domain wall or overlap fermion algorithms [37, 38].

Lattice QCD simulations have come a long way since the first calculations that demonstrated quark liberation in the high-temperature phase of pure SU(2) lattice gauge theory [39, 40]. State-of-the-art lattice calculations include both, SU(3) gauge fields and dynamical *u*, *d*, and *s* quarks (sometimes even c-quarks) with physical masses, employ improved lattice actions, extrapolate to the continuum limit (lattice distance $a \to 0$), and investigate the convergence of the results as a function of the number N_t of Euclidean time slices. Fully converged lattice QCD simulations for baryon-symmetric QCD matter ($\mu_B = 0$) have been made for the QCD equation of state and various other thermodynamic properties [41–43].

Rigorous algorithms for the evaluation of the functional integral for thermal QCD are known only for $\mu_B = 0$, because the integrand is not positive definite for non-zero real values of μ_B. Nevertheless, attempts have been made to explore the QCD phase diagram away from the $\mu_B = 0$ axis by means of lattice simulations using a number of indirect approaches. One method evaluates a certain number of derivatives of an observable with respect to μ_B at $\mu_B = 0$ and then reconstructs the value at $\mu_B \neq 0$ from the truncated Taylor series [44]. Another method evaluates the functional integral for imaginary values of μ_B, where the integrand is well-behaved, and then attempts an analytic continuation to real values of μ_B [45]. A third method uses field configurations generated for $\mu_B = 0$, but reweights them to simulate the μ_B-dependence of the lattice action [46]. The 'holy grail' of such efforts is to locate a possible critical point in the QCD phase diagram. In spite of considerable progress in obtaining reliable results for $\mu_B/T < 2$, none of these approaches currently gives a positive indication for the possible location of such a critical point. Even the existence of a critical point in the $T - \mu_B$ plane, although extremely plausible on the basis of physical arguments and predicted by a multitude of QCD models, has not been established with absolute certainty [47].

3.1.1. Pseudocritical temperature. It has now been firmly established that the transition in baryon-antibaryon symmetric ($\mu_B = 0$) QCD matter is a smooth cross-over [3]. As a consequence, the value of the pseudocritical transition temperature T_c cannot be defined unambiguously. Usually, one identifies T_c as the location of the peak or inflection point of a relevant thermodynamic quantity, but even this definition is ambiguous, because the value of T_c depends upon the choice of this quantity. The transition temperature defined by the inflection point of the renormalized quark condensate $\Delta_{l,s}$ [48] is $T_c^{(\chi)} \approx 158$ MeV [49] with an uncertainty of less than 1 MeV (see figure 2).

On the other hand, the transition region of quantities that are more direct measures of quark liberation, such as the expectation value of the Polyakov loop $\langle L \rangle$ or the strange quark number susceptibility is considerably broader. The expectation value of the Polyakov loop measures the interaction energy E_Q of an isolated heavy, static quark with the thermal bath of gluons via the relation $E_Q = -T \ln \langle L \rangle$. In the absence of light, dynamical quarks, it vanishes in the confined phase of QCD, which implies that an isolated quark would have infinite energy. At high temperature, in the deconfined phase, the Polyakov loop approaches unity because the interaction energy is small compared with the temperature. This behavior is illustrated in figure 3, which shows the temperature dependence of the strange quark susceptibility (left panel) and of the renormalized Polyakov loop (right panel). Both observables yield values of the pseudocritical temperature, defined by the location of the inflection point, of $T_c(L) \approx T_c^{(s)} \approx 175$ MeV. (The higher value may be simply due to the greater width of the transition for these quantities.) The transition occurs over a range of about 20 MeV for the strange quark condensate and an even wider range for the Polyakov loop,
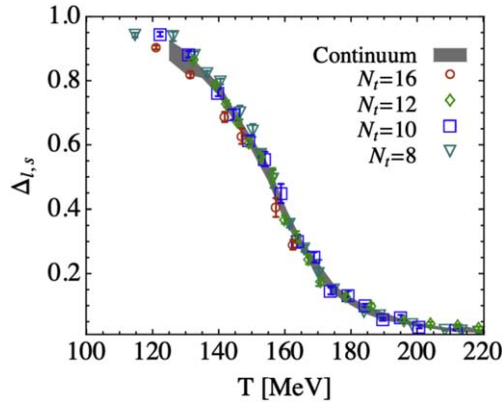
Figure 2. Temperature dependence of the renormalized light quark condensate $\Delta_{l,s}(T)$, divided by the value of the condensate in the vacuum, $\langle \bar{\psi}\psi \rangle_0$. The quark condensate, which becomes the order parameter of the chiral phase transition for massless quarks, serves as a measure of spontaneous chiral symmetry breaking. Reproduced from [48], with permission from Springer Nature.
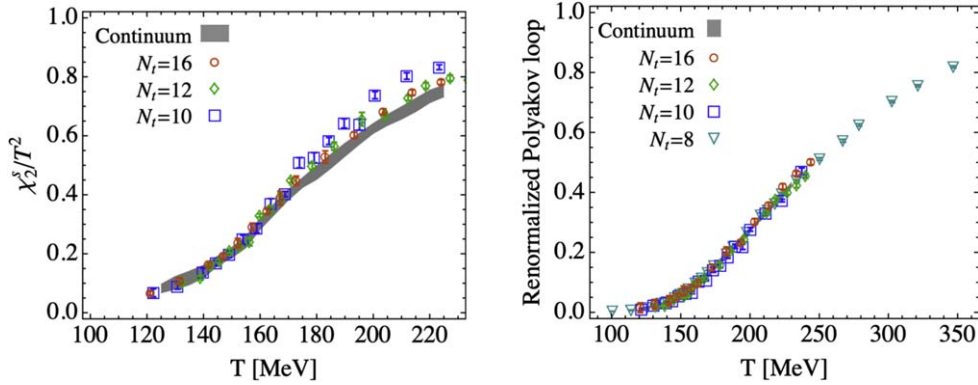
Figure 3. Left panel: temperature dependence of the strange quark number susceptibility $\chi_2^{(s)}/T^2$. Right panel: temperature dependence of the expectation value of the renormalized Polyakov loop. Both quantities provide measures of quark liberation. Reproduced from [48], with permission from Springer Nature.

which indicates that quark deconfinement is occurring gradually as the quark–gluon plasma is heated beyond T_c.

The transition is seen first, and over the narrowest interval, in the chiral quark condensate. This suggests that for the physical values of the parameters of QCD (the confinement scale Λ_{QCD} and the quark masses m_u, m_d, m_s) the transition is driven by the restoration of approximate chiral symmetry for the light quarks. The broad range of the transition seen in the strange quark susceptibility indicates that strange quark degrees of freedom thaw more gradually and at slightly higher temperature. The even more gradual transition of the renormalized Polyakov loop may be an indication that remnants of quark confinement in terms of correlations among quarks and antiquarks persist over a larger
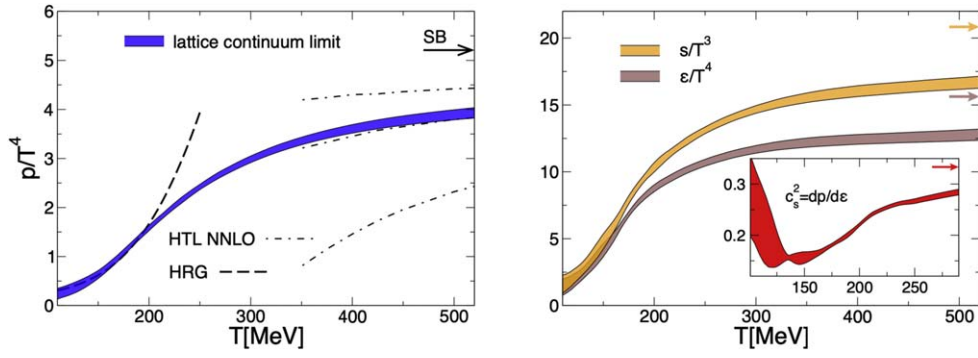
Figure 4. Continuum extrapolated equation of state of (2+1)-flavor hot QCD matter at $\mu_B = 0$. Left: Pressure $P(T)/T^4$; right panel: Energy density $\varepsilon(T)/T^4$ and entropy density $s(T)/T^3$. The insert on the right panel shows the speed of sound $c_s(T)^2$. The arrows indicate the Stefan–Boltzmann limit for a non-interacting gas of massless quanta. For an explanation of the theoretical curves labeled HRG and HTL NNLO, see text. Reproduced from [42]. CC BY 3.0.

temperature range. This behavior is also evident in the QCD equation of state, which we discuss next.

3.1.2. Equation of state. The QCD equation of state at $\mu_B = 0$ exhibits a rapid rise in the ratio of all thermodynamic quantities to the Stefan–Boltzmann (SB) limit for massless quanta over the temperature range $100\,\mathrm{MeV} < T < 500\,\mathrm{MeV}$, as shown in figure 4 [48]. Perhaps the clearest manifestation of this behavior is in the effective number of degrees of freedom d_{eff} defined via the formula for the entropy density $s(T)$ of a massless, noninteracting gas of particles:

$$s(T) = \frac{45}{2\pi^2}\, d_{\mathrm{eff}} T^3. \tag{13}$$

As seen in the right panel of figure 4, the effective number of degrees of freedom in hot QCD matter rises steadily over this temperature range from $d_{\mathrm{eff}} \approx 8.5$ at $T = 140\,\mathrm{MeV}$ to $d_{\mathrm{eff}} \approx 38$ at $T = 500\,\mathrm{MeV}$, which corresponds to 80% of the Stefan–Boltzmann limit. The remaining deviation from the SB limit can be understood as the effect of perturbative color interactions in the quark-gluon plasma.

Effective descriptions of the low-temperature and high-temperature domains are indicated in the left panel of figure 4. At low temperatures, the hadron resonance gas (HRG) provides a rather good description, although in detail the equation of state indicates the existence of hadron resonances beyond those listed in the Particle Data Book [50, 51]. The HRG gives a good description for most quantities up to $T \approx 150\,\mathrm{MeV}$. Some especially sensitive quantities may not only depend on the hadron mass spectrum, but also on the widths of resonances [52].

At high temperature ($T > 350\,\mathrm{MeV}$), hard-thermal loop resummed perturbation theory gives a good description of the equation of state and related observables. As indicated in the left panel of figure 4 the next-to-next-to-leading order (NNLO) calculations in this scheme [53, 54] match the lattice results for the pressure quite well, albeit with a large uncertainty deriving from the choice of the renormalization scale μ. The three dash-dotted lines shown in
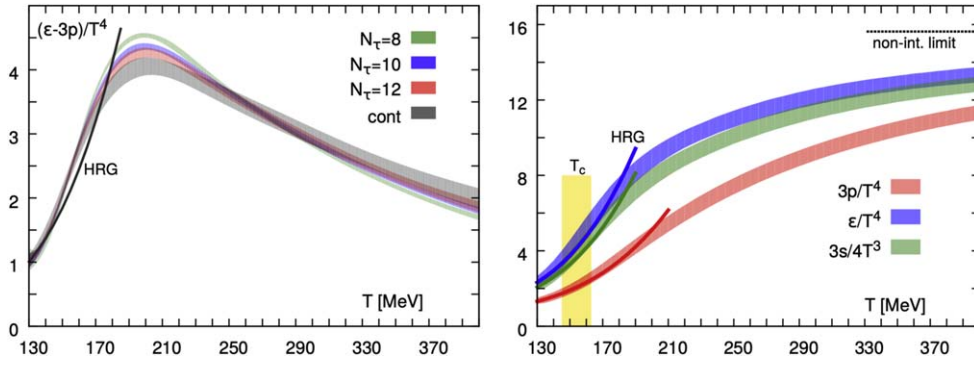
Figure 5. Left panel: continuum extrapolated scale anomaly $(\varepsilon - 3P)/T^4$ for (2+1)-flavor hot QCD matter at $\mu_B = 0$. Right panel: energy density $\varepsilon(T)/T^4$, pressure $P(T)/T^4$, and entropy density $s(T)/T^3$. The results from the Hot QCD collaboration shown here agree with those of the Wuppertal-Budapest (WB) Collaboration shown in figure 4. Reprinted (figure) with permission from [43], Copyright (2014) by the American Physical Society.

the figure correspond to $\mu = \pi T$ (top), $2\pi T$ (center), $4\pi T$ (bottom); the central curve obviously matches best.

An especially interesting feature is the broad peak in the QCD interaction measure, defined as $I(T) = (\varepsilon - 3P)/T^4$, shown in the left panel of figure 5, which measures the manifest violation of scale invariance in the QCD equation of state. The nonperturbative 'interaction measure' $I(T)$ is distinct from the perturbative violation of scale invariance in the light quark and gauge sector of QCD that derives from the temperature dependence of the running coupling constant $\alpha_s(T)$. There have been multiple attempts to interpret this nonperturbative feature of hot QCD in terms of some quasi-particle picture [55, 56], but the strong coupling among the degrees of freedom has made it difficult to draw definite conclusions. The scale anomaly remains large throughout the entire temperature domain ($T \leqslant 400$ MeV) that is probed in relativistic heavy ion collisions.

For $T < T_c$, the growth of $I(T)$ can be understood in terms of the hadron resonance gas model, which predicts that more massive states are being excited with increasing temperature. As the figure shows, the lattice result slightly exceeds the resonance gas model, which accounts for all experimentally known resonances, below 175 MeV. A possible explanation for this excess is the excitation of a large number of unknown resonances [50]. Although the microscopic structure of these states is not known, many of them must contain internal gluon excitations, either as hybrid states (hadrons that contain both valence quarks and valence gluons) or as glueballs (hadrons composed solely of valence gluons). If this is so, the question of what happens to the gluons when the quark-gluon plasma hadronizes, has a simple resolution: the thermally excited gluons become internal gluonic excitations of short-lived hadrons. As the hadron gas expands, these internal excitations will quickly decay into light hadrons, mostly pions, which absorb the additional entropy.

3.2. Thermal effective field theory

Lattice gauge theory currently does not permit to perform reliable computations of real-time processes. For this reason, insight into transport processes in the hot QCD matter presently relies either on thermal perturbation theory or on QCD-inspired models of strongly coupled

gauge theories with a holographic gravity dual. Holographic models will be discussed in the next subsection; here we focus on thermal perturbation theory.

Thermal perturbation theory is usually formulated in terms of the hard thermal loop (HTL) effective theory [57–59]. The HTL formalism resumes the leading thermal contributions to the gluon and quark self-energies into momentum-dependent effective masses. The effective theory relies on the separation of the scales $T \gg gT \gg g^2 T$, where $g \equiv g(T)$ is the QCD coupling constant at the thermal scale. The scale separation requires weak coupling (formally $g \ll 1$ corresponding $\alpha_s = g^2/4\pi \ll 0.1$). QCD is not so weakly coupled at any physically relevant temperature, but it turns out that the power counting implied by the HTL scheme appears to work even when the scale ordering is not strictly realized, as it often is the case for effective field theories.

Gauge invariance dictates that the thermal modifications of the propagators are balanced by vertex corrections that satisfy the Ward-Takahashi identities. In the HTL effective theory the gauge field develops a collective longitudinal mode, the plasmon, with mass

$$m_{\mathrm{pl}}^2 = \frac{2N_c + N_f}{6}(gT)^2 + O(g^4 T^2), \tag{14}$$

for a plasmon at rest in the medium. Here N_c and N_f denote the number of colors and light quark flavors, respectively. m_{pl} represents the additional mass scale that enters into the HTL effective Lagrangian and controls the thermal properties of the gauge theory. For example, the HTL formalism predicts that the static chromo-electric field around a heavy color charge in a QGP with N_f light quark flavors is screened with a Debye mass $m_D = \sqrt{3}\, m_{\mathrm{pl}} = gT\sqrt{1 + N_f/6}$. Furthermore, all HTL gauge field modes are Landau damped in the space-like domain at the same scale.

An estimate of the domain of applicability of HTL perturbation theory can be obtained by calculating and comparing quantities that can be computed reliably on the lattice. An example is the QCD equation of state, which has been calculated up to three-loop (NNLO) order for both, $\mu_B = 0$ [53, 54] and $\mu_B \neq 0$ [60], in the resummed HTL formalism and found to agree well with the lattice results for $T > 350 - 400$ MeV (see figure 4). It is also clear from these results that the HTL effective theory has very limited applicability in the temperature range $T < 2T_c$ that is most relevant to the heavy ion experiments.

Many properties of the quark-gluon plasma have been calculated in the HTL effective theory, mostly at leading order, but NLO results have increasingly become available. Examples include the shear viscosity η [61–63], the so-called jet quenching parameter $\hat{q}$ [64–66], the rate of collisional energy loss of a fast parton $\hat{e}$ [67], the heavy-quark diffusion constant D [68] and the rates of thermal photon [69, 70] and lepton-pair [71–73] production.

3.3. Strongly coupled gauge theories

The discovery of the Anti-de Sitter/Conformal Field Theory (AdS/CFT) duality by Maldacena in 1997 fundamentally changed theorists' ability to perform exact calculations for strongly coupled gauge theories [24, 74–76]. The AdS/CFT duality is based on the notion that conformal field theories in $(d + 1)$-dimensional space-time can be mapped onto quantum gravity in $(d + 2)$-dimensional Anti-de Sitter space (AdS$_{d+2}$), which has $(d + 1)$-dimensional space-time as a boundary at infinity. The AdS space is known as the 'bulk'. Because the mapping is between theories defined on space-times with a different number of dimensions, one speaks of a *holographic* duality.

The most ubiquitous case for relativistic heavy ion physics is the original duality between $\mathcal{N} = 4$ super-Yang-Mills theory in $(3 + 1)$ dimensions, which can be considered as a 'cousin'
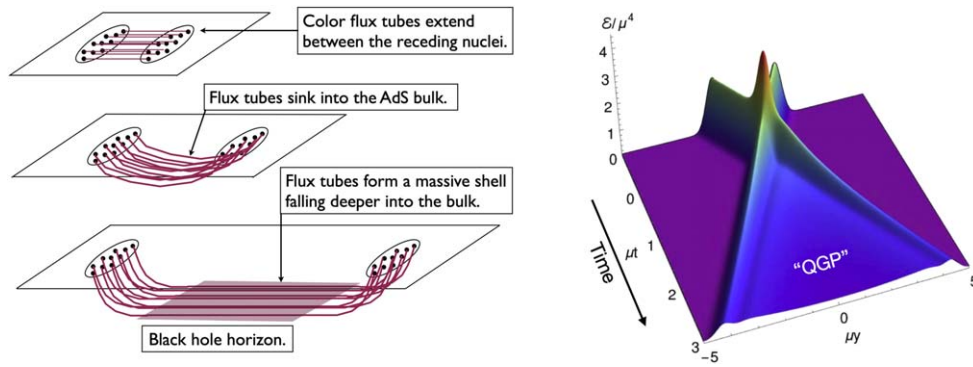
Figure 6. Left panel: schematic view of thermalization in a strong coupled gauge theory as formation of a black hole in the AdS bulk. Color flux tubes extending between the receding nuclei are idealized as a massive shell that falls into the depth of AdS space and eventually forms a black hole. Right panel: collision of shock waves in the dual AdS gravity theory. The energy density extending between the receding shocks is the holographic dual of the hot gauge field plasma. Reproduced from [86]. CC BY 4.0.

of QCD, and superstring theory on the space $AdS_5 \times S_5$. In the limit of an infinite number of colors N_c, expressed as the limit of strong 't Hooft coupling ($\lambda = g^2 N_c \rightarrow \infty$), the holographic dual is reduced to (super-)gravity on AdS_5. While QCD is not conformally symmetric—as we already mentioned, the violation of scale invariance is quite large in the temperature range of greatest interest to relativistic heavy ion physics—there are a number of modifications of the original holographic model that incorporate scale symmetry breaking in a fundamental (see [77] for an overview) or phenomenological way (see [78, 79]).

An especially attractive feature of AdS/CFT duality is that it enables rigorous calculations of dynamical processes at strong coupling, which is not possible in the framework of Euclidean lattice gauge theory [80]. In particular, it is possible to study thermalization at strong coupling by mapping thermalization onto the process of formation of a black hole in the bulk [81, 82], as illustrated in the left panel of figure 6.

The dynamics of space-time near the event horizon of black holes have long been known to be governed by viscous relativistic hydrodynamics, a phenomenon known as the membrane paradigm [83]. In the context of AdS space with a black hole, this translates into properties of viscous hydrodynamics in the gauge theory on the AdS boundary, some of which are deemed universal for any strongly coupled nonabelian gauge theory [84, 85]. The best known of these properties is the ratio of the shear viscosity to the entropy density, $\eta/s = 1/(4\pi)$ characteristic of a 'perfect' fluid.

Other experimentally relevant properties of strongly coupled gauge theories have also been calculated using holographic techniques, e.g. the energy loss of heavy quarks [87] and light quarks [88, 89]. Quite generally, holographic approaches lend themselves most easily to the study of energy and momentum transport as these are directly mapped onto the dynamics of the gravitational field in the bulk. Modeling of other processes associated with quark flavor or spin is possible, but requires adding additional fields or geometric objects, such as D-branes, to the model thereby rendering the conclusions less universally valid.

The holographic approach makes it possible to model the approach to hydrodynamics from arbitrary initial conditions without arbitrary approximations. Under a wide range of initial conditions, boost invariant (Bjorken) hydrodynamics with $\eta/s = 1/(4\pi)$ emerges as asymptotic description of the dynamical evolution [88, 90–92]. The complete process of energy
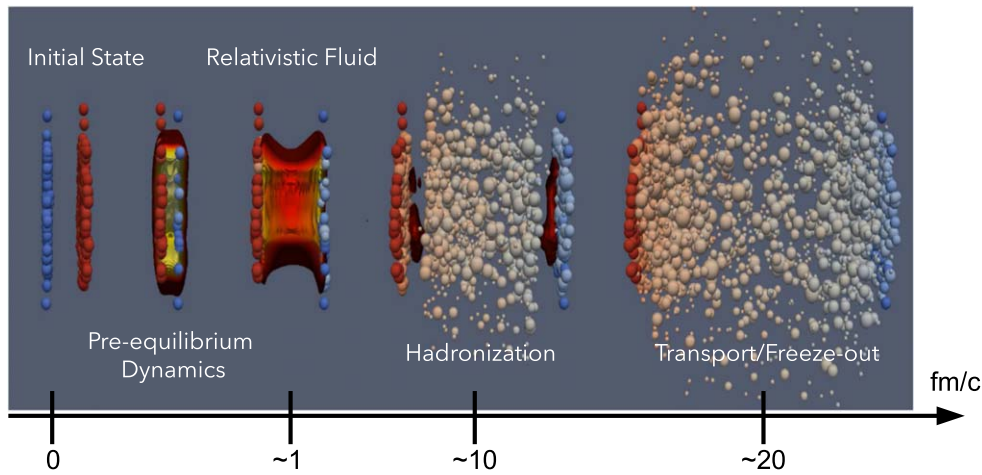
Figure 7. Visualization of the different stages of heavy-ion collisions in a hybrid approach based on hydrodynamics and hadronic transport for the initial and final stages. Reproduced with permission from G. Denicol.

deposition, thermalization, and hydrodynamic expansion can be modeled in terms of collisions of two gravitational shock waves. Such shock wave collisions in AdS space can be rigorously solved using numerical techniques [93–95]. By varying the height and width of the shock waves the whole range from transparency to full stopping [96] including asymmetric collisions [97, 98] can be explored. By adding a scalar field to AdS gravity, nonconformal models can be studied, which leads to hydrodynamic behavior including bulk viscosity [99]. When one combines holographic modeling with other descriptions that are more appropriate for the late stages of a nuclear collision, a seamless end-to-end description of heavy ion collisions can be implemented [100].

4. Phenomenological approaches

This section describes the basics of effective approaches that are currently used to connect theoretical input from fundamental theories to observables.

4.1. Bulk evolution

The dynamical evolution of heavy-ion collisions can be divided into several stages (see figure 7) that are described in detail in the following subsections. Initially, there are two Lorentz contracted nuclei approaching each other. Right after the impact there is a period characterized by non-equilibrium evolution. Eventually relativistic viscous fluid dynamics becomes applicable for the hot and dense stage. When the system dilutes enough, the fluid is converted to hadrons, which decay and interact with each other until the freeze-out.

4.1.1. Initial conditions. The initial conditions are an indispensable component of any dynamical description of heavy ion reactions. We have rather extensive knowledge of the structure of the colliding nuclei, but we lack a first-principle treatment on the basis of QCD of the evolution from the colliding nuclear ground states to the initial conditions that are used as input for hydrodynamics.

The simplest picture for the initial geometry of the overlap region in the transverse plane is based on the Glauber model. While the original Glauber model provides a general formalism for multiple scattering, its application to relativistic heavy ion physics commonly makes use of the eikonal limit, where all particles are assumed to travel along the straight line trajectories. (For a simplified introduction to the formalism, see [101]; for an overview of its various uses in relativistic heavy ion physics, see [102].) In this eikonalized Glauber model, the two nuclei are described by thickness functions $T_A(\vec{s})$, which represent Woods-Saxon density profiles denoted by ρ_A that are integrated over the longitudinal direction: $\hat{T}_A(\vec{s}) = \int \rho_A(\vec{s}, z)\,\mathrm{d}z$. $\vec{s}$ is the transverse position with respect to the center-of-mass of the nucleus A. The joint probability for the overlap of nucleus A with nucleus B colliding at an impact parameter $\vec{b}$ can be expressed as

$$\hat{T}_{AB}(\vec{b}) = \int \hat{T}_A(\vec{s})\hat{T}_B(\vec{s} - \vec{b})\,\mathrm{d}^2s. \tag{15}$$

Participant or 'wounded' nucleons are defined as the nucleons that interact with at least one nucleon of the opposite nucleus, where the interaction probability is obtained from the inelastic nucleon–nucleon cross section (σ_{NN}) at a given energy:

$$N_{\mathrm{part}}(\vec{b}) = A \int \hat{T}_A(\vec{s})\{1 - [1 - \hat{T}_B(\vec{s} - \vec{b})\sigma_{\mathrm{NN}}]^B\}\,\mathrm{d}^2s \tag{16}$$

$$+B \int \hat{T}_B(\vec{s} - \vec{b})\{1 - [1 - \hat{T}_A(\vec{s})\sigma_{\mathrm{NN}}]^A\}\,\mathrm{d}^2s, \tag{17}$$

where A and B denote the nuclear mass numbers. The nucleons outside the interaction region are called spectators and are not considered further. The number of binary nucleon–nucleon collisions (N_{coll}) can be calculated as well.

$$N_{\mathrm{coll}}(\vec{b}) = A\,B\,\hat{T}_{AB}(\vec{b})\sigma_{\mathrm{NN}}. \tag{18}$$

In some heuristic models [103] of the centrality dependence of the charged particle multiplicity binary collisions are assumed to contribute a fraction $\alpha < 1$ and while a fraction $(1 - \alpha)$ is attributed the participants:

$$N_{AA}(\vec{b}) = \left[\alpha N_{\mathrm{coll}}(\vec{b}) + \frac{1 - \alpha}{2} N_{\mathrm{part}}(\vec{b}) \right] N_{pp}, \tag{19}$$

where N_{pp} is the number of particles produced in a proton–proton collision at the same energy. The Glauber model is usually implemented as a Monte-Carlo process and therefore provides different initial condition profiles for individual events. Figure 8 shows the distribution of participant nucleons (solid circles) and spectator nucleons (dotted circles) in a randomly chosen Au+Au collision event at top RHIC energy. Many improvements of this simplistic scheme have been studied over the years. These include treatments of nucleon–nucleon correlations in the initial state [104], the finite extent and internal structure of nucleons, fluctuations in the energy deposition by individual nucleon–nucleon collision, and much more (see [105] for a recent review).

The *Color Glass Condensate* (CGC) model [107] aims to describe the colliding nuclei directly in terms of their quark-gluon content and is thus limited to collisions at sufficiently high-energy. The CGC model assumes that energy deposition at midrapidity is dominated by the liberation of gluons from the colliding nuclei, and that the relevant gluon content of the nuclei can be described by semiclassical gauge fields. These fields are generated by sources (primarily valence quarks) that move at near-projectile rapidity and whose dynamics during the collision can be ignored because of time dilatation. In the simplest version of the CGC, the
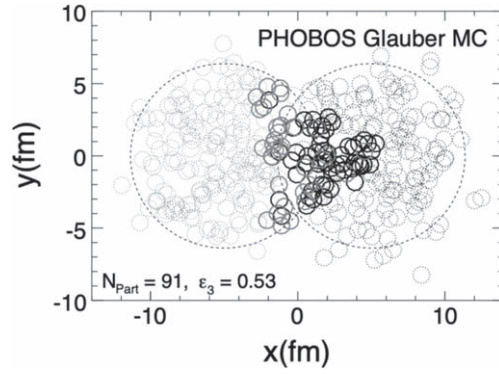
Figure 8. Transverse distribution of participant nucleons (solid circles) and spectator nucleons (dotted circles) in a random Au+Au collision event at $\sqrt{s_{\mathrm{NN}}} = 200$ GeV. The large dashed circles outline the transverse positions of the two nuclei. The event was selected for its nearly triangular distribution of the participant nucleons. Reprinted (figure) with permission from [106], Copyright (2010) by the American Physical Society.

McLerran-Venugopalan model [108, 109], the colliding nuclei are modeled as Gaussian distribution of color currents $J^{a\pm}(x) = \rho^a(\vec{x}_\perp)\delta(x^\mp)n^\pm$ along the light-cone. The distribution of these color sources is given by

$$W[\rho^a] = N \exp\left(-\frac{1}{2g^2\mu^2} \int \mathrm{d}^2 x_\perp (\rho^a(\vec{x}_\perp))^2\right), \tag{20}$$

where μ is a scale setting parameter that is related to the saturation scale $Q_s \approx 0.6 g^2 \mu$ [110]. The semiclassical color fields are obtained as solutions of the classical Yang-Mills equation $D_\mu^{ab} F^{a\mu\nu} = gJ^{a\nu}$, where D^{ab} is the gauge covariant derivative and $J^{a\,\mu}$ is the color current due to quarks.

The classical approximation is only warranted when the occupation number of gauge field modes is of the order of $\alpha_s(Q_s)^{-1} \gg 1$. This condition requires that $Q_s \gg \Lambda_{\mathrm{QCD}}$. Because the saturation scale Q_s scales as $Q_s^2 \sim A^{1/3}x^{-\lambda}$ with $\lambda \approx 0.3$, the condition can be fulfilled either for a large nuclear mass number $A \gg 1$ or for very small momentum fraction $x \ll 1$. Figure 9 gives a graphical representation of this scaling. Small values of x are only accessible at high collision energies, because the transverse momentum p_T of a produced hard probe is given by the relation $p_T^2 = x_1 x_2 s_{\mathrm{NN}}$, where x_1 and x_2 are the momentum fractions of the participating partons in the two colliding nuclei and $\sqrt{s_{\mathrm{NN}}}$ is the center-of-mass collision energy per nucleon pair.

The dependence of the color sources $\rho^a(x_\perp)$ on the momentum fraction x is governed by renormalization group (RG) evolution. At a given value of x_0 all sources of color fields residing at $x > x_0$ contribute (valence quarks for all x_0 and gluons and sea quarks for small values of x_0). The RG evolution can either be expressed as a differential equation for the functional $W[\rho^a]$, the so-called JIMWLK equation [111], or as a nonlinear differential equation for so-called dipole density $N_{\mathbf{xy}}$, the Balitsky-Kovchegov (BK) equation [112]. Whereas the scaling with A and x can be calculated from geometry and QCD renormalization group arguments, the overall scale is nonperturbative. This means that the range of applicability of the CGC model is still under debate and awaits precise data from the future electron-ion collider (EIC).
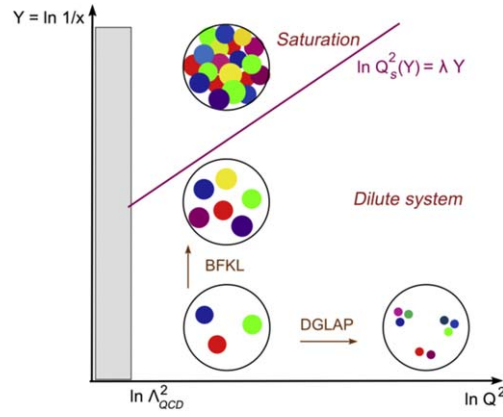
Figure 9. Schematic view of different domains of the gluon distribution in nuclei. The Color Glass Condensate (CGC) model or saturation model applies above the *x*-dependent saturation scale $Q_s(x)$ shown by the straight line. Dilute parton systems at low Bjorken-*x* are described by the Balitsy–Fadin–Kuraev–Lipatov (BFKL) equation; dilute and weakly interacting parton systems at large virtuality Q^2 are described by the Dokshitzer–Gribov–Lipatov–Altarelli–Parisi (DGLAP) equations. Reproduced with permission from [107].
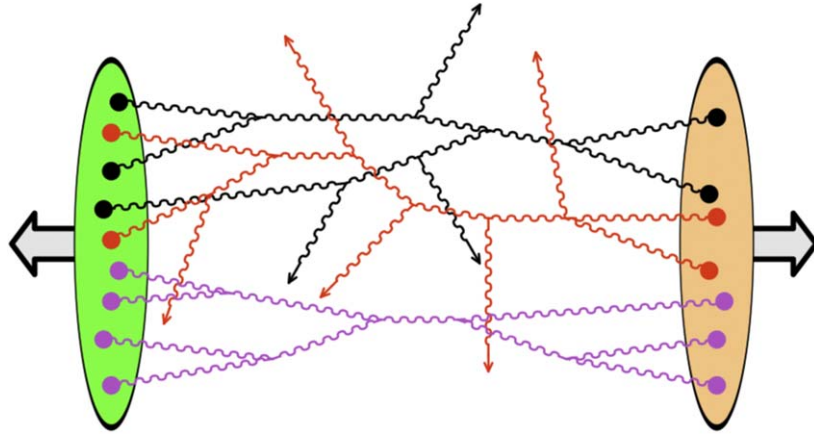
Figure 10. The not yet thermalized distribution of gluons liberated by scattering from the CGC fields of colliding nuclei is called glasma. Reproduced with permission from [107].

The color fields generated in this way are virtual and remain an integral part of the nuclei until an interaction occurs. During the collision of two heavy ion the CGC fields of the nuclei interact, and a large fraction of the virtual gluons are scattered on-shell, as sketched in figure 10. The resulting nonequilibrium state is known as glasma. Because the CGC color fields are transversely polarized, the resulting energy-momentum tensor is far off equilibrium and cannot be directly used as an initial condition for the hydrodynamic evolution.

In order to describe individual heavy ion collision events, it is necessary to not only model the average energy deposition, but also the event-by-event fluctuations. There are two
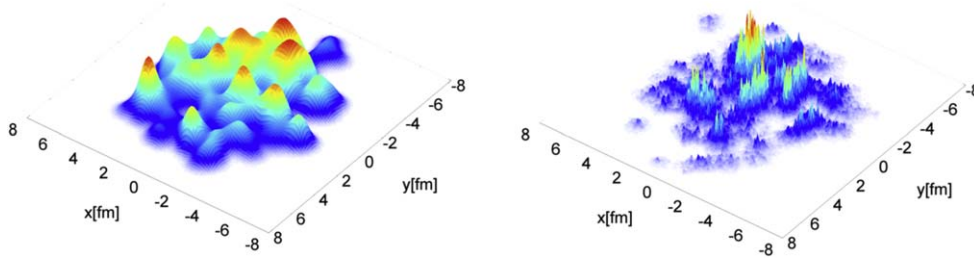
Figure 11. Fluctuating initial energy density in the transverse $x - y$ plane for a single event of a Au+Au collision at top RHIC energy. Left panel: Nucleon-based Glauber model; right panel: IP-glasma model. In the nucleon-based Glauber model the fluctuations occur on the length scale of the nucleon radius ($\sim$0.8 fm). In the IP-glasma model the fluctuations occur on the length scale of the inverse saturation scale $Q_s^{-1} \sim 0.1$ fm, and result in a much more spiky distribution. Reprinted (figure) with permission from [116], Copyright (2012) by the American Physical Society.

main sources of such initial-state fluctuations. One are the quantum fluctuations in the distribution of nucleons in the colliding nuclei or, if one instead looks at the collision at the parton level, in the distribution of valence quarks in the nuclei. The other source of fluctuations are the quantum fluctuations of particle production that are encoded in the S-matrix of a nucleon–nucleon or quark-quark collision. In the standard Glauber model, whether it is implemented at the nucleon or the valence quark level, the second type of fluctuations is commonly parametrized by a negative binomial distribution with the generating function [113]

$$ F(t; k) = \frac{k^k}{\Gamma(k)} t^{k-1} e^{-kt}, \tag{21} $$

where t is a continuous variable and the parameter k (with $1 < k < 4$) depends on the energy. The probability distribution of the number of produced particles is given by

$$ P(n) = \int_0^\infty dt \, F(t; k) e^{-\bar{n}t} \frac{(\bar{n}t)^n}{n!}, \tag{22} $$

where $\bar{n}$ is the average number of produced particles.

For the CGC model, one can obtain fluctuating initial conditions using a formalism similar to the Glauber model, where the positions of nucleons are chosen randomly from the nuclear density distribution and the color sources are distributed within the individual nucleons. This implementation of the CGC model is known as impact parameter saturation (IP-sat) model [114]. When applied to nuclear collisions, the IP-sat model generates fluctuating initial conditions for the glasma (IP-glasma), which naturally result in a negative binomial distribution for the initial energy density [115, 116]. The energy densities obtained with these two alternative approaches are depicted in figure 11, where the nucleon-based Glauber model (left panel) is seen to result in a much smoother initial distribution than the parton-based IP-sat model (right panel).

The initial pattern of energy or entropy deposition in these models does not lend itself directly to be used as initial conditions for viscous hydrodynamics, because they do not correspond to an energy-momentum tensor near thermal equilibrium and the spatial gradients are too large in the case of event-by-event fluctuations. In the case of the standard Glauber model, this problem is commonly resolved by a short period of free streaming evolution,

which smoothes out the spatial gradients. Alternatively, one widely used version of a general Glauber-type initial condition generator, the TRENTO model [117], uses the entropy density to map the initial conditions generated by the model directly onto the initial conditions for viscous hydrodynamics.

For CGC initial conditions, the initial time evolution of the glasma is described by the nonlinear classical Yang-Mills equation; the resulting energy-momentum tensor is then mapped onto initial conditions for hydrodynamics [116]. The microscopic thermalization process is thought to follow a scenario known as 'bottom-up' thermalization [118]. In this scenario semi-hard gluons contained in the colliding nuclei radiate soft gluons which form a thermal bath that drains energy from the semi-hard gluons via elastic scattering and thereby thermalizes them. An effective kinetic description of this thermalization process has recently been developed and applied to the initial condition of hydrodynamics in relativistic heavy ion collisions [119]. There are many additional details and possible refinements of this process, which are under investigation, but no generally accepted approach has emerged yet.

4.1.2. The QGP phase. Due to asymptotic freedom, the quark-gluon plasma was originally expected to behave as a weakly interacting gas with almost massless degrees of freedom, the quarks and gluons, for which collective flow would be difficult to establish on the length scale of the fireball. However, when the first measurements of anisotropic flow in Au+Au collisions at $\sqrt{s_{\mathrm{NN}}} = 200$ GeV at RHIC were published, a surprisingly good agreement with predictions from ideal fluid dynamic calculations was observed [120, 121], if the initial conditions were matched to fit the overall particle production yields and transverse momentum spectra. The elliptic flow is then a natural consequence of collective behavior for a geometrically anisotropic initial state.

Fluid dynamics is based on the equations for the conservation laws

$$\partial_\mu T^{\mu\nu} = 0, \quad \text{and} \quad \partial_\mu N_i^\mu = 0 \tag{23}$$

of the energy-momentum tensor $T^{\mu\nu}$ and the current density N_i^μ of any conserved quantum number (e.g. *B*, *S*, *Q*). If the fluid is in local equilibrium, the energy-momentum tensor of an ideal fluid takes the form

$$T^{\mu\nu} = (\epsilon + P)u^\mu u^\nu + Pg^{\mu\nu}, \tag{24}$$

with ϵ being the energy density, *P* the pressure and u^μ denoting the flow 4-velocity. $g^{\mu\nu}$ is the metric tensor and typically chosen as $(+, -, -, -)$ in special relativistic applications. The quantum number current density has the simple form $N^\mu = nu^\mu$, where *n* is the density measured in the local rest frame.

In viscous hydrodynamics, the energy-momentum tensor is decomposed as

$$T^{\mu\nu} = \varepsilon u^\mu u^\nu - \Delta^{\mu\nu}(p + \Pi) + \pi^{\mu\nu}, \tag{25}$$

with Π denoting the bulk viscous contribution and $\pi^{\mu\nu}$ the shear viscous tensor. First order viscous hydrodynamics is not causal, since it permits signal propagation that is faster than the speed of light. Therefore, second order viscous fluid dynamics must be applied with the following relaxation equations for shear and bulk viscosity including the coupling between the two

$$D\Pi = \frac{-\zeta\theta - \Pi}{\tau_\Pi} - \frac{\delta_{\Pi\Pi}}{\tau_\Pi}\Pi\theta + \frac{\lambda_{\Pi\pi}}{\tau_\Pi}\pi^{\mu\nu}\sigma_{\mu\nu} \tag{26}$$
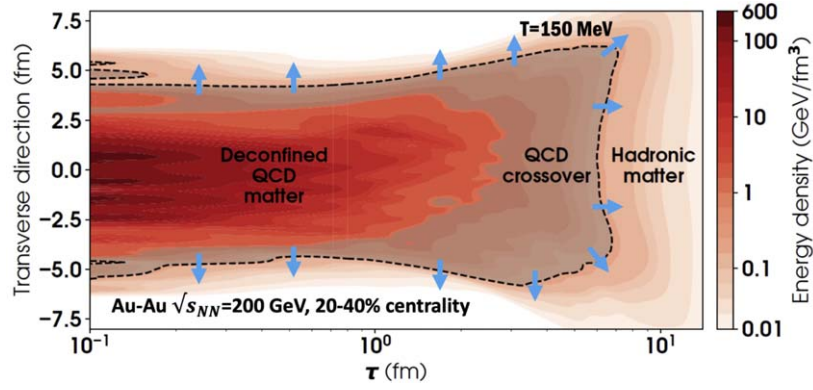
Figure 12. Contour plot of the evolution of the energy density in a midcentral Au+Au collision at the highest RHIC energy $\sqrt{s_{\mathrm{NN}}} = 200$ GeV. The horizontal axis shows the local proper time τ, the vertical axis shows one of the transverse coordinates. The black dashed line delineates the hadronization hypersurface $T = T_c = 150$ MeV. Reproduced with permission from [131]. CC BY-NC-ND 4.0.

$$D\pi^{\langle\mu\nu\rangle} = \frac{2\eta\sigma^{\mu\nu} - \pi^{\mu\nu}}{\tau_\pi} - \frac{\delta_{\pi\pi}}{\tau_\pi}\pi^{\mu\nu}\theta + \frac{\phi_7}{\tau_\pi}\pi_\alpha^{\langle\mu}\pi^{\nu\rangle\alpha}$$
$$- \frac{\tau_{\pi\pi}}{\tau_\pi}\pi_\alpha^{\langle\mu}\sigma^{\nu\rangle\alpha} + \frac{\lambda_{\pi\Pi}}{\tau_\pi}\Pi\sigma^{\mu\nu}. \tag{27}$$

$\theta = \partial_\mu u^\mu$ is the expansion scalar, $\sigma^{\mu\nu}$ is the strain tensor, τ_π and τ_Π are the relaxation times for the shear and bulk viscous corrections and $\lambda_{\Pi\pi}$, $\lambda_{\pi\Pi}$, $\delta_{\pi\pi}$, $\tau_{\pi\pi}$ and ϕ_7 are higher-order couplings, whose forms are taken from [122].

There are two typical choices for the local rest frame, the Landau frame which moves with the energy density and the Eckart frame which moves with the conserved density, if one exists. More details on viscous relativistic hydrodynamics and its application to heavy ion physics can be found in [123–125].

Since the set of hydrodynamics equations listed above is insufficient to determine all independent components of the quantities of interest, one needs an additional equation as an input. This is the equation of state, which encodes the properties of the fluid under consideration. In the context of heavy ion physics this has the major benefit that hydrodynamics provides a controlled handle on the phase transition between a hadron gas and the quark-gluon plasma phase. One can study different equations of state to explore its influence on the dynamics of the system and various experimentally measurable observables.

The simplest equation of state is that of a relativistic ideal gas composed of massless particles:

$$P = \frac{\varepsilon}{3}. \tag{28}$$

At vanishing baryon chemical potential the equation of state is known from lattice QCD calculations and can be parametrized as originally done in [126]. Currently, there are many other possibilities including the extension to multiple conserved currents (see e.g. equations of state by the BEST collaboration [127], the BNL group [128] or from holography [129]). At finite density and vanishing temperatures there are constraints from the mass-radius relation of neutron stars. Also any equation of state should take the constraints based on stable nuclear matter into account.

At high collision energies, hydrodynamics predicts the evolution of the fireball at midrapidity to be approximately boost invariant [130]. The hydrodynamical variables then become functions of the transverse coordinates and the local proper time τ. The evolution of the fireball can be visualized by contour plots of the energy density, as shown in figure 12 for a Au+Au collision at top RHIC energy. At lower beam energies the framework has to be extended to include finite net baryon charges and three spatial dimensions, since the assumption of a boost invariant longitudinal expansion breaks down.

The second input into hydrodynamics calculations, which contains information about the microscopic properties of the medium, are the transport coefficients. Main consideration has been given to the shear and bulk viscosity (η and ζ), but recently calculations have begun to include quantum number diffusion coefficients and even their cross-talk [132]. The most interesting dimensionless quantities are η/s and ζ/s, where s is the entropy density. It is extremely challenging to compute transport coefficients from fundamental lattice QCD, therefore one has to rely on models and approximations as well as limits for the high temperature–high density region (perturbative QCD) and the low temperature–low density region (hadron gas models). One of the major goals of the heavy ion program is to infer the values of the QCD transport coefficients from experimental data by detailed model-to-data comparison as explained in section 5.

Fluid dynamics is only applicable when the system is sufficiently close to local equilibrium. Very shortly after the initial collision of the two nuclei, which originally only have longitudinal momentum, this condition is not satisfied. There is some transition time required for the transverse momenta of the medium constituents to become of similar magnitude as the longitudinal momenta. The initial state and associated thermalization process has been discussed in section 4.1.1. In the later stages of the reaction, the increasingly dilute medium will drop out of equilibrium, when the expansion rate is faster than the collision rate. This stage is the subject of the next section 4.1.3. The boundary conditions for the differential equations that govern the fluid dynamic behavior, are very important for a complete understanding of the dynamics.

Conservation equations can be solved using a variety of algorithms. It is important to ensure that they also apply to relativistic hydrodynamics and event-by-event calculations with potentially spiky initial conditions without creating numerical artefacts. For this purpose a suite of tests of several scenarios, for which an analytic solution is known, was developed within the TEC-HQM collaboration.[8] Some widely used implementations of viscous fluid dynamics for heavy ion collisions are: MUSIC [133], vHLLE [134], VISH2 + 1 [135], and CLVisc [136].

On the theoretical side, there are activities aimed at rooting relativistic viscous hydrodynamics in kinetic theory (see e.g. [137]). Formalisms extending the applicability of fluid dynamics are also being developed. One example is anisotropic fluid dynamics that allows for larger differences between the longitudinal and transverse pressure in the system and therefore applies at earlier times than second-order viscous hydrodynamics [138, 139]. Owing to the interest in the chiral magnetic effect and polarization observables, attention has been given to developments of fluid dynamic formulations with spin and or chiral currents (see e.g. [140–142]).

Fluid dynamics is not the only option to describe the hot and dense stage of heavy ion reactions, where the quark-gluon plasma is formed. Another possibility are microscopic transport codes that include partons as their degrees of freedom. Quarks and gluons are either treated as nearly massless particles interacting according to cross sections derived from

---

[8] https://wiki.bnl.gov/TECHQM/index.php/TECHQM_Main_Page (see under *Documents*)

perturbative QCD [143, 144] or as dynamical quasiparticles [145]. These microscopic non-equilibrium calculations have the advantage of providing detailed information about the complete phase-space over the entire dynamical evolution of the fireball at the expense of making more detailed assumptions that may be difficult to test.

4.1.3. Hadron transport and freeze-out. When the matter created in a heavy ion collision becomes more dilute again, the quark-gluon plasma hadronizes and the hadrons subsequently decouple from the system. The moment in time when hadron abundances are fixed and inelastic interactions no longer occur is called chemical freeze-out, while the moment when also elastic scattering ceases is called kinetic freeze-out. The standard approach to the dynamical evolution at high beam energies involves hybrid models, based on (viscous) fluid dynamics as discussed in the previous section 4.1.2 for the hot and dense stage and Boltzmann-type hadronic transport equations for the late dilute stage of the reaction (see [146] for a review).

The transition from partons to hadrons happens within the fluid dynamics calculation via a change of degrees of freedom in the equation of state. When a chosen switching criterion, usually either a certain temperature or a certain energy density, is reached, a three-dimensional hypersurface is constructed [147] within the four-dimensional space-time that serves as the basis for the particlization process, where fluid quantities are mapped into particle properties [148]. An example of such a Cooper–Frye freeze-out hypersurface is shown as the dashed curve in figure 12. The equation of state needs to be identical on both sides of the Cooper–Frye hypersurface and the transition happens according to

$$\frac{\mathrm{d}N_i}{\mathrm{d}^3 p} = \frac{d_i}{(2\pi)^3} \int_{\Sigma} [f_0(x, \vec{p}) + \delta f_{\mathrm{shear}}(x, \vec{p}) + \delta f_{\mathrm{bulk}}(x, \vec{p})] \frac{p^\mu \, \mathrm{d}\Sigma_\mu}{E_{\vec{p}}}. \tag{29}$$

Here Σ is the hypersurface, f_0 the equilibrium distribution function, and δf denotes the viscous corrections to the distribution function. The specific form of the viscous corrections is one of the major uncertainties in the current hybrid approaches [149]. Once the particles are sampled according to their distributions on the hypersurface one can feed them into hadronic transport approaches like UrQMD [150, 151] or SMASH [152].

One generally assumes that the majority of particles escapes outward of this hypersurface but, in principle, there can be also inward directed contributions. These are particles that reenter the hydrodynamic evolution and are therefore very hard to deal within practice [154]. While all other parts of the evolution satisfy exact conservation laws event by event, in the sampling process they are only fulfilled for an ensemble of events. Therefore, very often many such samples and hadronic evolutions are calculated for each hydrodynamic event (oversampling method). To study correlation and fluctuation observables one might need a more careful consideration of conservation laws at the Cooper–Frye transition [155]. For calculations of dilepton emission one needs to include the finite-width spectral functions of the resonances in the sampling process. For this purpose, it is necessary to match the degrees of freedom and their properties in the hadronic transport approach that is chosen for the final rescattering and decays.

The advantage of the microscopic non-equilibrium treatment of the late hadronic stage is that the chemical and kinetic freeze-out are happening automatically, whenever the corresponding interactions cease for each particle species individually. An alternative option is to keep running hydrodynamics and introducing partial chemical equilibrium. In this approach one encounters a kinetic freeze-out hypersurface, an example of which is shown in figure 13. The explicit particle degrees of freedom have the added advantage that the outcome is very similar to what is measured in experiment. When using unified output formats, like
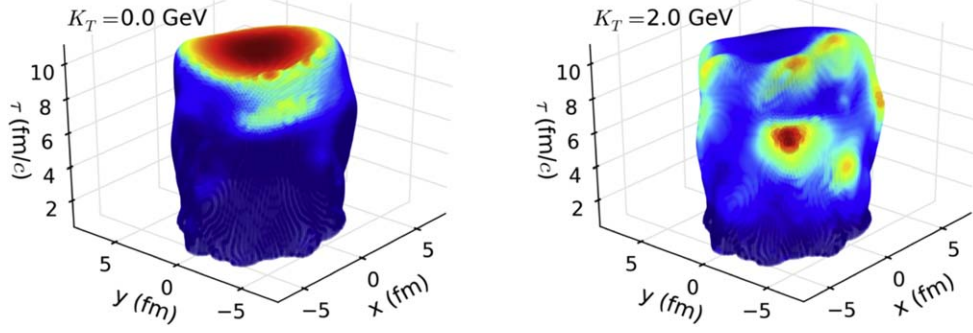
Figure 13. Intensity plots of the emission hypersurface for hadron pairs in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV. Left panel: hadron pairs with total momentum $K_T = 0$, selectively weighting hadrons from bulk hadronization. Right panel: Hadron pairs with $K_T = 2$ GeV/c, selectively weighting hadrons from surface radiation. The emission intensity is highest in the dark red regions and lowest in the dark blue regions. Reproduced with permission from [153].

OSCAR or the HepMC library[9], the same analysis can then be applied to the simulation results and to the experimental data, which allows for more accurate comparisons. Recently, the RIVET library[10] is also getting extended to include heavy ion observables, which could be an interesting option in the future.

The same information on hadronic degrees of freedom and their interactions that one needs for final-stage hadron transport calculations in high-energy heavy ion collisions is employed for the full evolution in low energy heavy ion reactions. Hadronic transport approaches are based on the relativistic Boltzmann equation where, appropriate for low net baryon density, the mean-field term is omitted and only collisions are taken into account. The input for cross-sections and processes is taken from the Particle Data Book, experimental data for elementary reactions, or estimated using effective models. For low energy binary scatterings, the reactions are modeled via resonance excitation and decay, while at higher energies string excitation and fragmentation is used to model inelastic processes. Such microscopic approaches have the advantage that one has access to the complete phase-space information of all particles at all times, and they often serve as event generators for simulations of experimental results.

4.2. Hard and electromagnetic probes

Hard probes of hot and dense QCD matter are defined as those that are, in important parts, perturbatively calculable. This applies to strongly interacting (QCD) probes at high virtuality and to electromagnetic probes. In most instances, hard probes rely on input that is nonperturbative and must be derived from other experimental data. Parton distribution functions (PDFs) $f(x)$ and fragmentation functions $D(z)$ are salient examples.

Generally, hard probes in p+p or A+A collisions rely on the principle of factorization, illustrated in figure 14: A scattering process is factorized into a nonperturbative initial-state matrix element, a perturbatively calculable hard scattering S-matrix element, and a nonperturbative final-state matrix element:
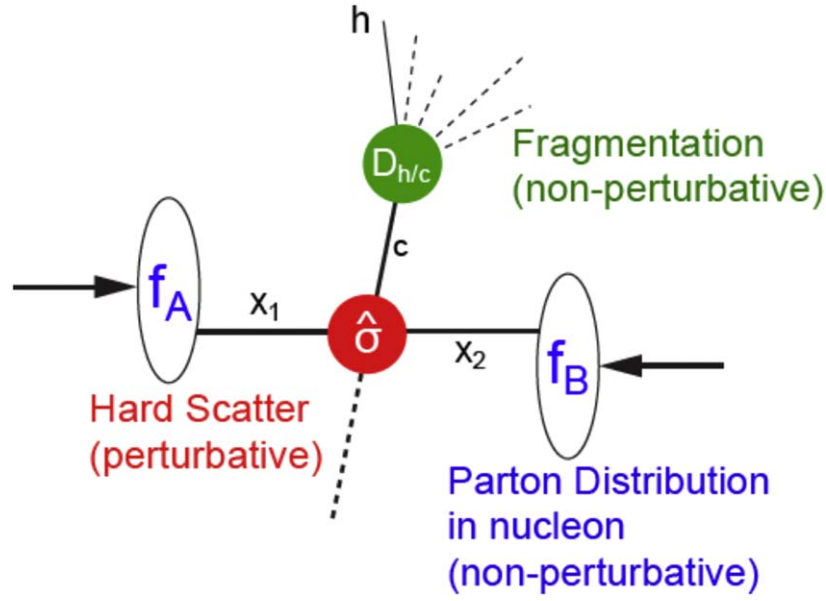
Figure 14. Factorization of a hard scattering process into initial-state, hard scattering, and final state-state processes. (Reproduced with permission from K. Reygers.)

$$d\sigma = f_A(x_1, Q^2) \otimes f_B(x_2, Q^2) \otimes d\hat{\sigma}(Q^2) \otimes D_{h/c}(z, Q^2). \tag{30}$$

The process at the core of this scheme, the hard scattering cross section $d\hat{\sigma}(Q^2)$, is universal and only depends on the hard scale Q^2, but the initial- and final-state matrix elements that give rise to the parton distribution functions $f(x, Q^2)$ and the fragmentation function $D(z, Q^2)$ can be modified by nuclear effects—cold nuclear matter in the case of the initial state and hot QCD matter in the case of the final state.

The deceptively simple sketch in figure 14 hides many of the complexities that arise as a result of the nonperturbative interactions occurring in the initial and final states. This is illustrated in figure 15. Their sensitivity to the nuclear medium is what makes them of core interest to relativistic heavy ion physics. The challenge is to encapsulate these modifications in physics-motivated descriptions that can be used to probe the properties of the nuclear medium. Examples are parton energy loss, diffusion coefficients, saturation scales, and so on.

We have already covered models of the initial state in section 4.1.1. We now discuss models for the final-state modifications. The output of the hard scattering process are two or more energetic partons, which then can interact with the medium created in the collision, often referred to as the underlying event. These interactions can result in energy or momentum loss [67], deflection [156], emission of radiation [157], or transmutation of the parton into another particle [158]. A sequence of such interactions leads to modifications of the jet shower evolution, which manifests itself in various observables (see section 5.9) and can be modeled using Monte-Carlo methods (see e.g. [159, 160]).

The mechanisms causing an energetic parton to lose energy during its propagation through matter can be divided into two categories: elastic and inelastic. Elastic energy loss occurs in two-body scattering, where energy is transferred to the scattering partner. Inelastic energy loss occurs predominantly by radiation of a gluon following the interaction with a scattering center in the matter. At high energies, inelastic energy loss dominates; at low energies elastic energy loss dominates. The threshold at which inelastic energy loss begins to dominate increases
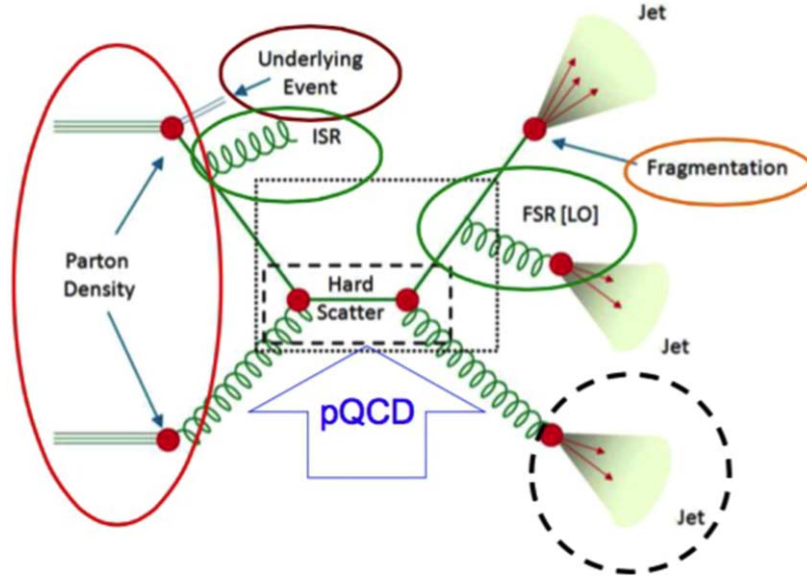
Figure 15. Complexities hiding in the factorization of a hard scattering process into initial-state, hard-scattering, and final-state processes. Reproduced with permission from D. S. Cerci.

with the mass of the energetic parton, because radiation from heavy particles is suppressed by the dead-cone effect [161].

Elastic energy loss is one aspect of momentum change in the medium caused by scattering. For heavy quarks with mass $M \gg T$ multiple scatterings are required to appreciably change the momentum. The process can then be considered as a diffusion process and described by a Langevin equation with a fluctuation constant κ and drag coefficient η.[11] In a thermal medium κ and η are related to the spatial diffusion constant D and to each other by the Einstein relation

$$\kappa = \frac{2T^2}{D}, \quad \eta = \frac{T}{MD} = \frac{\kappa}{2MT}. \tag{31}$$

The diffusion constant can be calculated in thermal perturbation theory [68], on the lattice [162], or by holographic techniques [163]. The Langevin framework can be easily simulated within a fluid dynamical description of the quark-gluon plasma. Rare scatterings with larger momentum transfer are not described by the Langevin approach. These can be treated in the framework of linearized Boltzmann transport in a thermalized, hydrodynamically flowing medium [164]. Both descriptions can be merged into a unified framework that effectively describes all aspects of heavy quark transport [165].

Inelastic energy loss, i.e. energy loss by gluon radiation, is the dominant mode of energy loss for light quarks and gluons. Medium-induced radiation requires off-mass shell scattering of the energetic parton in the medium, followed by gluon emission. In a dense medium, multiple scattering events contribute coherently to a single radiation event. The resulting suppression of the radiation yield is known as Landau–Pomeranchuk–Migdal (LPM) effect.

---

[11] We alert the reader to the multiple use of the symbol η, which can denote the pseudorapidity, shear viscosity or, here, the drag coefficient, depending on the context.

The scattering power of the medium (momentum transfer squared per unit length) is expressed in the transport coefficient $\hat{q}$ defined in (3). Several different formalisms have been developed to describe the radiative energy loss of quarks or gluons. These correspond to different approximations or truncations of the multiple scattering process. For an overview and comparison of existing approaches, see [166]. General overviews of energy loss and jet quenching can be found in [29].

A fraction of the energy lost by an energetic parton traversing the quark-gluon plasma is deposited into the thermal medium. This applies to all energy lost by elastic collisions and to the soft component of the radiated energy which thermalizes. The response of the medium to the injection of energy and momentum along the trajectory of the energetic parton is treated by adding a source term to the hydrodynamic equations. The linear response of the medium can be understood in terms of sound propagation and results in the possible formation of a Mach cone [167]. A more microscopic approach treats the individual interactions within the evolution of a full jet as sources in a numerical hydrodynamics code [168]. For the effect of the medium response on jet observables, see [169].

Electromagnetic (EM) probes, i.e. photons and dileptons (often called real and virtual photons), are special in that they do not suffer from final state interactions. (They may, however, suffer from large backgrounds.) Their production can be related to the photon spectral function $\rho(\omega, \vec{p})$ in the medium [58]. In the vacuum, ρ is a function of the invariant mass $M^2 = \omega^2 - p^2$ only; in a thermal medium, $\rho(\omega, M)$ is a function of, both, M and the frequency measured in the fluid rest frame $\omega = -u_\mu p^\mu$. The rate of photon emission can be expressed as [170]

$$\omega \frac{\mathrm{d}R}{\mathrm{d}^3 p} = \frac{n_\mathrm{B}(\omega)\rho(\omega, 0)}{(2\pi)^3} \tag{32}$$

with $n_\mathrm{B}(\omega) = (e^{\omega/T} - 1)^{-1}$. The dilepton emission rate is given by

$$\frac{\mathrm{d}R}{\mathrm{d}^4 p} = \alpha \frac{n_\mathrm{B}(\omega)\rho(\omega, M)}{12\pi^4 M^2}\left(1 + \frac{2m^2}{M^2}\right)\sqrt{1 - \frac{4m^2}{M^2}}\,\theta(M^2 - 4m^2), \tag{33}$$

where m denotes the lepton mass. Apart from the fact that the measured yields involve an integration over the space-time evolution of the emitting matter, EM probes thus provide direct information about the photon spectral function. A broad overview of electromagnetic probes of relativistic heavy ion collisions can be found in [171]; concise up-to-date reviews are contained in [172, 173]. For a comprehensive recent review of direct photon emission see [174].

The photon spectral function is predicted to change from one that is dominated by vector mesons at $T < T_c$ to a broad continuum corresponding to free quark-antiquark pair excitations at $T > T_c$ (see figure 16). It can be estimated using lattice QCD [176, 177], thermal perturbation theory [72], or using QCD sum rules [178, 179]. In the sum rule approach, the transition of the photon spectral function is closely related to the phenomenon of chiral symmetry restoration. At low frequencies and long wavelengths, the spectral function is related to the electrical conductivity σ of the QGP by

$$\sigma = \frac{1}{2}\lim_{\omega \to 0}\frac{1}{\omega}\rho(\omega, p)\big|_{\omega^2 = p^2} = \frac{1}{3}\lim_{\omega \to 0}\frac{1}{\omega}\rho(\omega, p)\big|_{p=0}. \tag{34}$$

It is a challenge for the future to combine the soft bulk evolution with the hard probes consistently. Many calculations use state-of-the-art descriptions of the bulk matter evolution for the evaluation of hard probes, but including the medium response and developing a fully coupled picture over the different kinematic regions remain challenges for the future.
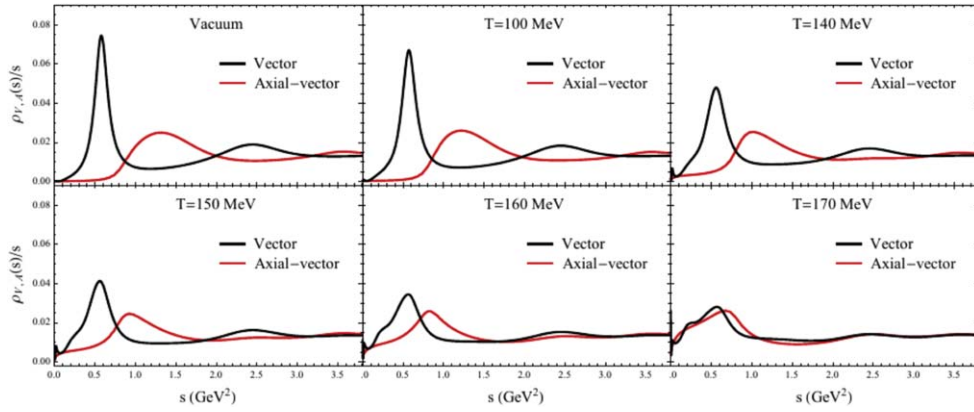
Figure 16. Conjectured evolution of the vector (black curves) and axial vector (red curves) spectral functions with temperature T for $\mu_B = 0$ obtained by solution of QCD sum rules. Reproduced from [175], with permission from Springer Nature.

4.3. Hadronization models

Different kinematic regions of the particles produced in heavy ion collisions are governed by different regimes of quantum chromodynamics. The soft region below ~ 3 GeV is typically described by the bulk dynamic evolution (see section 4.1), while the hard region above ~ 6 GeV is covered by perturbative QCD (see section 4.2). Experimentally, it is only possible to detect final-state hadrons, while we would like to learn about the hot and dense quark-gluon plasma stage of the reaction. Therefore, the hadronization process plays an important role in our interpretation of the data from heavy ion reactions.

Figure 17 shows an example of the nuclear modification factor R_{AA} as a function of transverse momentum for several particle species from the ALICE collaboration in PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. This picture serves nicely as a basis to introduce the hadronization models that are applied in different regions.

Below ~ 2 GeV the particle production is described by the viscous hydrodynamic evolution, and hadronization happens via the change of degrees of freedom in the equation of state. On the Cooper–Frye surface the hadronic fluid is converted to individual particles and those are then propagated by a hadronic cascade. This is one of the big advantages of hydrodynamics, that the transition from partonic to hadronic degrees of freedom is under nice theoretical control as long as local equilibrium can be assumed. In this region the nuclear modification factor shows the typical mass dependence expected from collective behavior. For a more detailed discussion see section 5.2.

Above a transverse momentum of $p_T \sim 8$ GeV one observes a rather universal behavior, this is the region where hadronization is described by fragmentation of color flux tubes. The strings break apart when the leading quark-antiquark (or quark-diquark) pairs are pulled apart and the color field becomes strong enough to create particle-antiparticle pairs out of the vacuum. Subsequently, the newly produced partons connect with each other to form hadrons with well-defined quantum numbers.

The intermediate region $p_T = 2-8$ GeV is governed by the recombination and coalescence models for hadronization and the interaction of hard fragmentation and soft hadrons. Recombination means that the partons from the quark-gluon plasma are clustered into
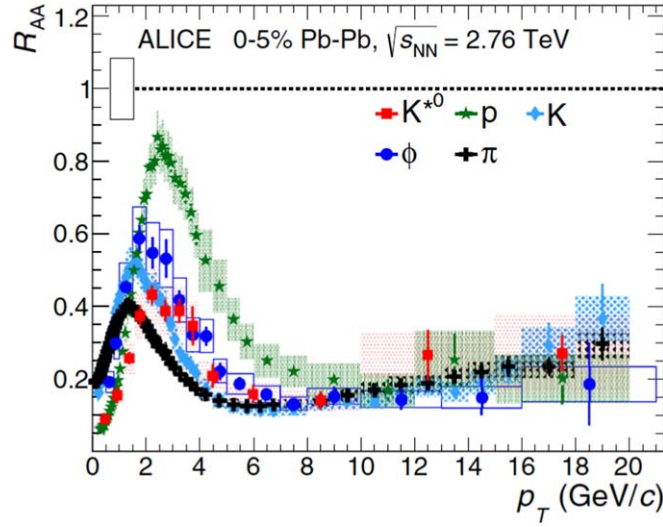
Figure 17. Transverse momentum dependence of the nuclear modification factor of several particle species measured by the ALICE collaboration. Reproduced from [180]. CC BY 4.0.

hadrons when they are close in phase space (see section 5.8 for more details). In this region, unusually large baryon-to-meson ratios can be observed in the nuclear modification factor.

Understanding hadronization on a microscopic level poses a challenge for the future. It is important to advance the understanding on how hadrons emerge from the collective partonic state that is formed in ultra-relativistic heavy ion collisions. This involves understanding QCD on multiple different scales connecting the different regions outlined above in a smooth fashion.

5. Experimental results

This section reviews selected experimental results from the experiments with relativistic heavy ion collisions at RHIC and LHC. Note: the aim of this section is to present an overview of the relevant data, not of theoretical attempts to describe them quantitatively. In many instances, the figures shown in this section contain theoretical curves, which we do not explain in detail as they are extraneous to our discussion. The interested reader is urged to consult the original references.

5.1. Location in the phase diagram

The first question to be answered when studying heavy ion reactions is which temperature and density is reached in the collision. While full dynamical models as described in section 4 provide (model-dependent) answers to this question, the possibility of defining a temperature for the system has more fundamental implications. To be able to assign a single temperature and net baryon chemical potential to the system created in such a highly dynamical process hints at the fact that an equilibrated plasma has formed. Another interpretation involves the idea that hadronization follows statistical behavior and, therefore, the numbers of produced hadrons follow thermal expectations. The second interpretation is supported by the finding

that even in particle production from e^+e^- collisions thermal particle yields are measured [181].

The basic idea to connect final yields at midrapidity or 4π multiplicities to a temperature relies on the formation of a fireball in global equilibrium, that emits all particles at the same instant in time. The thermal model actually has no notion of time evolution or spatial variations, therefore it is limited to a small set of observables, namely the particle abundances. The grand canonical partition function includes all mesons and baryons of the non-interacting hadron resonance gas. The fireball is treated as a grand canonical ensemble with Bose and Fermi statistics depending on the particle spin. The distribution function for particle species i is given by

$$\frac{\mathrm{d}N_i}{\mathrm{d}p^3 \mathrm{d}x^3} = \frac{g_i}{(2\pi)^3} \frac{1}{\exp^{(E-\mu_i)/T} \pm \alpha_i}, \tag{35}$$

where g_i is the degeneracy factor, T is the temperature, $\mu_i = \mu_B B_i + \mu_S S_i$ the particle-specific chemical potential and $\alpha_i = \pm 1$ for fermions and bosons, respectively. Integrating over the entire phase space and assuming Boltzmann statistics—a good approximation for all hadrons except pions when $|\mu_i|/T$ is small—one obtains the following well-known expression for thermal particle production

$$N_i = \frac{g_i V T^3}{2\pi^2} \, \mathrm{e}^{\mu_i/T} \, \frac{m_i^2}{T^2} \, K_2\left(\frac{m_i}{T}\right), \tag{36}$$

where N_i is the number of particles of species i produced in the fireball of volume V at temperature T and chemical potential μ_i. m_i denotes the mass of the particle and K_2 is the modified Bessel function of the second kind. The main quantity determining the particle yields is the mass of the particle species. One expects higher mass particles to be produced with smaller probabilities than lighter particles. To remove sensitivity to the volume, the usual strategy is to look at particle yield ratios that are assumed to be emitted from the same volume. By fitting ratios of several different particle species one can then obtain temperatures and chemical potentials.

It is impressive how well a vast amount of data on particle production in heavy ion collisions can be understood with such simple assumptions. One recent example is shown in figure 18 (left) where particle ratios at midrapidity from Pb+Pb collisions at $\sqrt{s_{\mathrm{NN}}} = 2.76$ TeV are fit within the statistical model. The bars denote the model fit and the symbols show experimental data from the ALICE collaboration. Even the light nuclei follow the trend of thermal production at a unique global temperature and chemical potential. A widely adopted interpretation is that these thermodynamic properties reflect the conditions prevailing at the chemical freeze-out, the moment at which the inelastic reactions cease and the abundances are frozen. (Elastic rescattering is still possible until the kinetic freeze-out which will be discussed in more detail in the next section 5.2.) It is worth pointing out, however, that this interpretation is not logically consistent, because the model also describes the yields of particles that cannot be formed under the chemical freeze-out conditions, such as light nuclei. A physically credible interpretation of the observed yields of such states can be made either on the basis of the dynamical coalescence model [183] or by invoking general properties of strongly coupled, highly excited quantum systems [184].

The right-hand part of figure 18 shows the result of thermal fits to particle yields and multiplicities over a large range of beam energies. As expected, the temperature rises as a function of beam energy, while the net baryon chemical potential decreases. This can be understood by the stopping dynamics of the two nuclei passing through each other. At lower energy the nuclear passing time is longer and more baryons are stopped at midrapidity, while
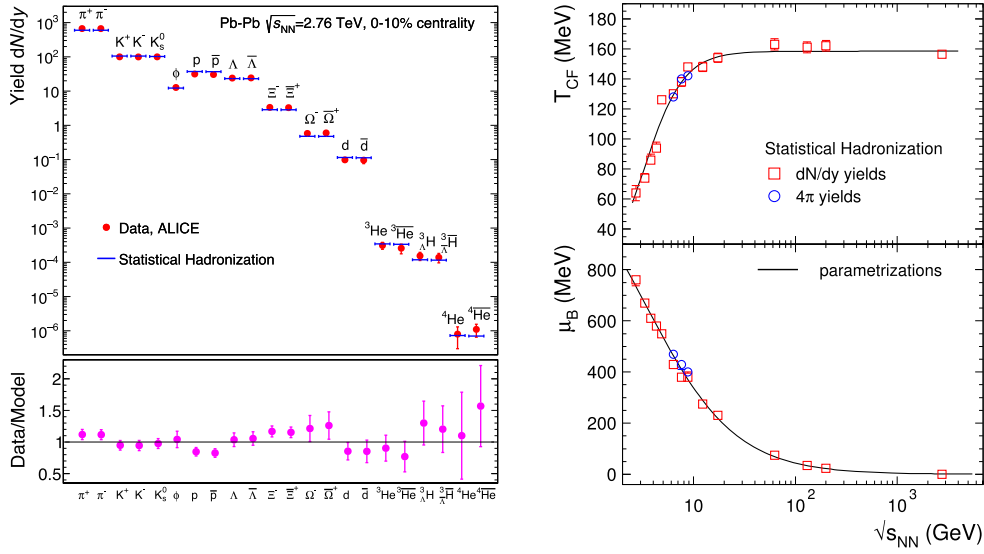
Figure 18. Left: Hadron abundances (dN/dy values at midrapidity) from the ALICE collaboration for central Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV compared to statistical hadronization model analysis. Right: Energy dependence of chemical freeze-out parameters T_{CF} and μ_B. The results are obtained from the statistical hadronization analysis of hadron yields (at midrapidity, dN/dy, and in full phase space, 4π) for central collisions at different energies. Reproduced from [182], with permission from Springer Nature.

at high beam energies the nuclei fly through each other so fast that the baryon number remains localized at forward rapidities (transparency). The parametrizations shown in figure 18 (right) are [182]:

$$T_{CF} = \frac{T_{CF}^{lim}}{1 + \exp(2.60 - (\ln \sqrt{s_{NN}})/0.45)}, \qquad \mu_B = \frac{a}{1 + 0.288\sqrt{s_{NN}}}, \qquad (37)$$

with $\sqrt{s_{NN}}$ in GeV, $T_{CF}^{lim} = 158.4$ MeV and $a = 1307.5$ MeV.

When plotting the values for temperatures and net baryon chemical potential obtained in thermal model fits on the phase diagram one obtains a curve as specified in figure 19. Since those values correspond to a hadron resonance gas, the quark-gluon plasma formation can only happen beyond this line. These values define a lower bound of temperatures and densities that are reached in the corresponding heavy ion collisions, since the chemical freeze-out happens at the end of the evolution. The shape of the line is consistent with the hypothesis that the chemical freeze-out is associated with a constant energy per baryon in the system [185].

At lower beam energies and for smaller system sizes, the assumptions of a grand canonical ensemble may no longer be valid, especially for particles carrying a conserved quantum number such as strangeness. Since s and $\bar{s}$ quarks must be produced in pairs, the conservation of net strangeness in the volume must be taken into consideration. Using a thermal model that includes these ideas, one can even obtain decent fits to particle ratios from low energy heavy ion reactions as shown in the left panel of figure 20 by the HADES collaboration. While it is remarkable how well the extracted temperatures and net baryon chemical potentials fit into the
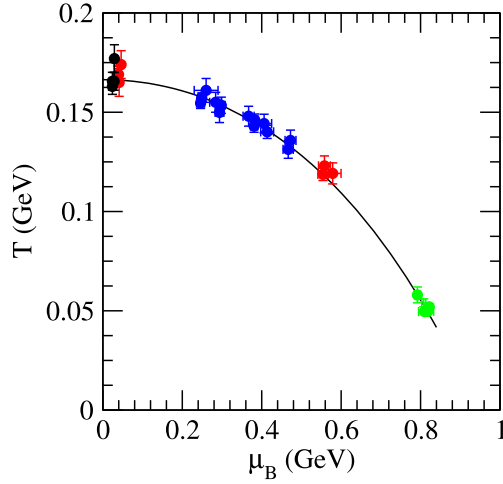
Figure 19. Values of μ_B and T for different energies. The solid line is a parameterization corresponding to $T(\mu_B) \approx 0.17 - 0.13\mu_B^2 - 0.06\mu_B^4$ where T and μ_B are in units of GeV. Reprinted (figure) with permission from [185], Copyright (2006) by the American Physical Society.
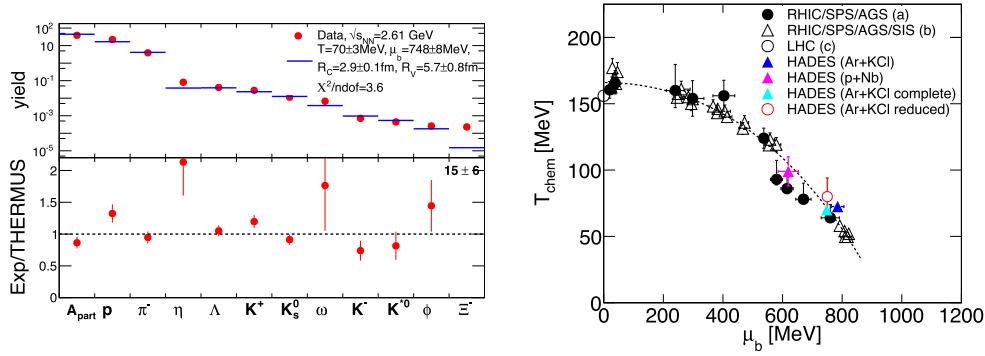
Figure 20. Left: Hadron abundances (dN/dy values at midrapidity) from the HADES collaboration for Ar+KCl collisions at $\sqrt{s_{NN}} = 2.6$ GeV compared to THERMUS calculation. Right: Chemical freeze-out parameters from different systems at different beam energies in the $T - \mu_B$ plane. Reproduced from [186], with permission from Springer Nature.

world data (see right panel of figure 20), some discrepancies are also appearing. For example, the high yield of Ξ baryons needs a different explanation.

At high beam energies, there are some indications that strange particles have a higher chemical freeze-out temperature than hadrons consisting of light quarks only [187]. This could be understood by the smaller cross-sections of strange hadrons with the surrounding matter. Those studies go beyond fitting yields by investigating the relevant susceptibilities through particle number fluctuations of various orders. There are a lot of caveats when comparing these observables directly to grand canonical lattice QCD calculations but, if properly taken into account, allow to extract valuable information on the thermodynamic properties of the system.

If strangeness is found to be in thermal equilibrium with the rest of the produced particle species, it is an indication of quark-gluon plasma formation since the strangeness production is enhanced through partonic production channels compared to reactions at the hadron level [188]. There have been even claims that the peak in the kaon-to-pion ratio as a function of beam energy around $E_{\text{lab}} = 40A$ GeV is a signature for the phase transition to the quark-gluon plasma [189]. Also the 'step' in the mean transverse momentum of kaons as a function of beam energy has been attributed to the latent heat of the phase transition. However, in dynamical model calculations both phenomena can be attributed to more conventional effects, namely the transition from a baryon-dominated to a meson-dominated system, as well as to the softening of the equation of state due to the excitation of more active degrees of freedom associated with hadron resonances [190, 191].

Besides canonical suppression, there are other more advanced variants of the thermal model including, for example, van der Waals interactions [192] or excluded volume effects that can vary for different species [193]. The limitations of the thermal model are that it is a static picture. It is not possible to address observables of more dynamical origin (e.g. anisotropic flow, HBT radii,..) within this approach. Also, resonances cannot be described since the rescattering of daughter particles restricts the ability of experimental collaborations to reconstruct all resonances produced at chemical freeze-out. At LHC energies, baryon annihilation during the late stage of the fireball expansion plays a significant role and affects the proton yields, which also leads to challenges in a fully thermal description [194].

In addition to looking globally at particle production at midrapidity or integrated yields, one can investigate the variation of particle yields as a function of rapidity. For example, the rapidity distribution of net protons gives valuable insight into the baryon stopping dynamics [195]. In the context of the present section, it is important to realize that by studying more forward rapidities higher net baryon densities are accessible also at colliders like RHIC and LHC [196]. Of course, the high rapidity region is typically not equipped with as many detector elements capable of particle identification, but the high energies and high multiplicities might still provide an advantage over low beam energies, especially in fixed target experiments as they have been conducted in the STAR detector at RHIC and at the CERN-SPS.

5.2. Transverse flow

The basic dynamic observables related to transverse flow are transverse momentum distributions of (identified) particles. A more integral, commonly reported measure is provided by the mean transverse momentum $\langle p_T \rangle$. These observables carry information about the radial expansion of the system in addition to thermal information. While the yields of particles are fixed at chemical freeze-out when inelastic reactions cease, the slopes of the transverse momentum spectra contain information about the kinetic freeze-out of the different particle species. Kinetic freeze-out is the moment when the particles no longer suffer elastic collisions and the momenta become fixed. Because elastic cross sections differ from particle to particle the effective temperatures extracted from spectra of different particles do not necessarily agree. Note that the decay of a resonance into the same (final) state from which it was formed, such as the reaction $\pi N \to \Delta \to N\pi$, is understood as a pseudoelastic process, since it does not alter the chemical composition.

Employing the ideas introduced in the previous section 5.1 one can calculate the thermal expectation for particle spectra differentially along transverse momentum under the assumption of a Boltzmann distribution and obtains that
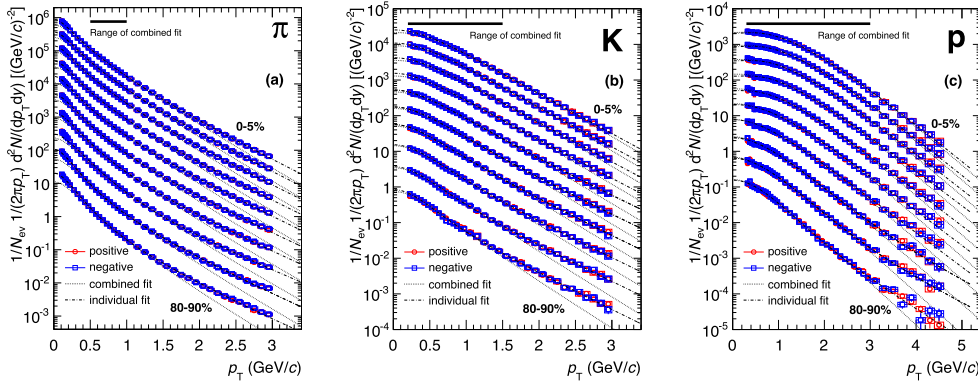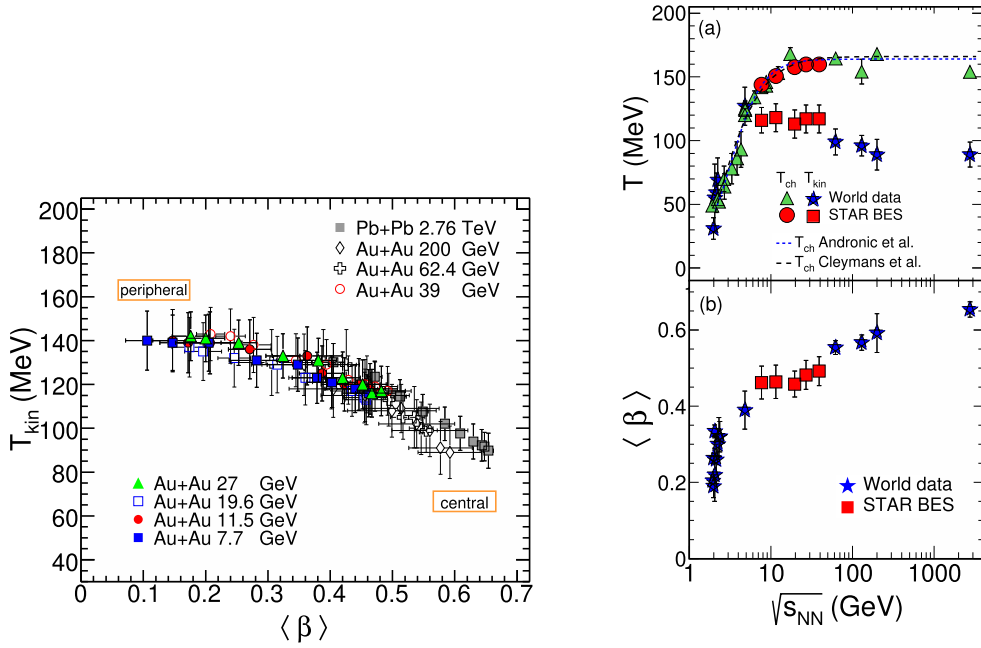
Figure 21. Transverse momentum spectra of pions (a), kaons (b) and protons (c) in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ GeV from the ALICE collaboration for different centralities (central ones are scaled). The dashed curves represent blast-wave fits to individual particles, while dotted curves indicate combined blast-wave fits. Reproduced from [197].

$$\frac{1}{p_T}\frac{\mathrm{d}N}{\mathrm{d}p_T} \propto \mathrm{e}^{-p_T/T_{\mathrm{eff}}}. \tag{38}$$

where p_T is the transverse momentum, and T_{eff} is the slope of the distribution in a semi-logarithmic plot. In figure 21 the transverse momentum spectra for pions, kaons and protons in heavy ion collisions at LHC energies are shown. The shoulder at low p_T indicates a deviation from the purely thermal expectation that becomes more pronounced with increasing hadron mass, in particular, for protons. To obtain exponential distributions one can apply a variable transformation and fit transverse mass spectra $m_T^{-1}(\mathrm{d}N/\mathrm{d}m_T)$ instead of transverse momentum spectra. It is interesting that most of the information on particle production can be summarized by the integrated yield at midrapidity and the mean transverse mass. If a theoretical calculation reproduces both, the full spectra are typically described rather well.

By measuring transverse momentum distributions in different centrality windows and rapidity bins, one can obtain the yields as a function of centrality and rapidity by integrating over these spectra. Since the acceptance of experimental detectors is often limited in the very low momentum region, one needs adequate fit functions. The particles with very low momenta are bent strongly in the magnetic fields and therefore do not reach the tracking detectors. To estimate systematic uncertainties, different functional forms are employed, e.g. pure exponential functions in p_T or m_T, functions including quantum statistics or Tsallis-Levy distributions.

To extract kinetic freeze-out temperatures and average flow velocities from transverse momentum spectra of several particle species the blast-wave model is employed. The assumption here is that all particles are emitted from a boosted fireball that follows a common radial expansion profile. The different behavior of the different species is then mainly attributed to the different masses that result in different mean transverse momenta even for the same underlying velocity profile. To obtain integrated yields and mean transverse momenta, the spectra shown in figure 21 have been fitted individually for each species with a blast-wave function [199]:

Figure 22. Left: results for kinetic freeze-out temperature and radial velocity from fits to transverse momentum spectra from the RHIC beam energy scan and LHC. Right: chemical and kinetic freeze-out temperatures and radial flow velocities at kinetic freeze-out over a collision energy range spanning three orders of magnitude. Reprinted (figure) with permission from [198], Copyright (2017) by the American Physical Society.

$$\frac{1}{p_T}\frac{\mathrm{d}N}{\mathrm{d}p_T} \propto \int_0^R r\mathrm{d}r\, m_T\, I_0\!\left(\frac{p_T \sinh\rho}{T_{kin}}\right) K_1\!\left(\frac{m_T \cosh\rho}{T_{kin}}\right), \tag{39}$$

where the velocity profile ρ is described by

$$\rho = \tanh^{-1}\beta_T(r) = \tanh^{-1}(\beta_s(r/R)^n). \tag{40}$$

Here, $m_T = \sqrt{p_T^2 + m^2}$ is the transverse mass, I_0 and K_1 are modified Bessel functions, r is the radial distance in the transverse plane, R is the transverse radius of the fireball, $\beta_T(r)$ is the transverse expansion velocity and β_s is the transverse expansion velocity at the surface.

Combining the spectra of different particle species in a joint fit makes it possible to extract kinetic freeze-out temperatures and average flow velocities for different beam energies and centralities as shown in figure 22. The results are rather similar at high RHIC and LHC energies. As a function of centrality, one can see a trend that more central collisions result in lower temperatures and higher transverse velocities, while more peripheral collisions end up at higher temperatures and with less explosive radial expansion. The apparently counter-intuitive ordering of the deduced temperatures indicates that the fireball in central collisions produces more particles and is 'cooking' longer, such that the particles decouple at lower temperatures when the system had more time to expand. One caveat is that the results are rather sensitive to the fit ranges, which have to be carefully chosen (see the black bars in figure 21).

The world data shown in the right panel of figure 22 reveal the expected behavior of rising radial expansion velocities with increasing collision energy. The extracted temperatures show

a more interesting pattern. At low beam energies the chemical freeze-out overlaps with the kinetic freeze-out, while above $\sqrt{s_{NN}} \sim 8$ GeV the kinetic freeze-out temperatures are lower than the chemical ones and stay constant or even decrease at higher beam energies. One way to understand this invokes the same argument as above for central collisions, that the fireball has a longer lifetime at high beam energies and, therefore, the particles decouple later in the evolution.

During the final expansion of the hadronic fireball both, the momenta of the hadrons and the chemical composition of the fireball, change. The latter can occur in two ways: by decay of unstable, excited hadrons, and by inelastic collisions among hadrons. The latter are generally thought to diminish rapidly as the fireball cools, since most low-energy interactions among ground-state hadrons are elastic. A notable exception to this rule are baryon-anti-baryon annihilation reactions, which have large cross-sections at low energies. Hadron momenta, of course, can be changed by elastic scattering. This applies especially to the most abundant hadrons at high beam energies, pions and nucleons, which scatter with a very large cross section due to the Δ-resonance.

One important consideration when comparing model calculations to experimental data is to 'match apples to apples'. For transverse mass spectra this seems straightforward, but there can be complications like feed-down corrections that have to be taken into account. In many models the Λ hyperon is regarded as stable particle, because it is stable with respect to the strong interaction. In practice, some of the produced Λ hyperons undergo a weak decay before they reach the detector, which is at a macroscopic distance from the collision vertex. Therefore, the reported number of protons differs by about 40 % depending on whether the feed-down from Λ hyperons is being considered or not. Similarly, the Σ^0 decays to almost 100% into Λ and needs to be added in model calculations before comparing to experimental data.

Transverse momentum spectra and particle yields are also the basic observables that a hydrodynamic evolution needs to describe. In figure 23 calculations within the VISHNU hybrid approach based on the VISH2 + 1 dimensional viscous hydrodynamic code and hadronic rescattering by UrQMD are compared to pion and proton spectra in Au+Au collisions at the highest RHIC energy. One can see that different levels of initial state fluctuations described by Monte-Carlo Glauber or a saturation-based MC-KLN approach do not affect the spectra much. Also, varying the constant shear viscosity-over-entropy density ratio η/s employed in the hydrodynamic evolution from 0 to 0.24 does not have a large effect. (See below for the dependence on bulk viscosity.) Of course, smaller variations might become visible when one compares results on a linear scale. In general, transverse momentum spectra can be used to gauge the basic hydrodynamic evolution parameters. The anisotropic flow discussed in the next section 5.4 has a higher sensitivity to the initial state fluctuations and to the properties of the quark-gluon plasma and its transport coefficients.

Figure 24 (left) shows the mean transverse momentum of pions, kaons and protons in heavy ion collisions at LHC energies, measured by the ALICE collaboration, in comparison with a hybrid calculation relying on IP-Glasma fluctuating initial conditions, the 3 + 1 dimensional viscous hydrodynamic implementation MUSIC and hadronic rescattering by UrQMD. As a function of centrality, the mean transverse momenta are found to be very sensitive to the inclusion of bulk viscosity in the evolution. A temperature-dependent bulk viscosity with a peak around the transition temperature was employed in this study. While there is the caveat that this calculation still used an equation of state with a transition temperature around $T_c = 190$ MeV, the qualitative differences are expected to remain also with an equation of state fitted to state-of-the-art lattice calculations (see e.g. the recent study of deuterons and their sensitivity to bulk viscosity in [203]). The right panel of figure 24
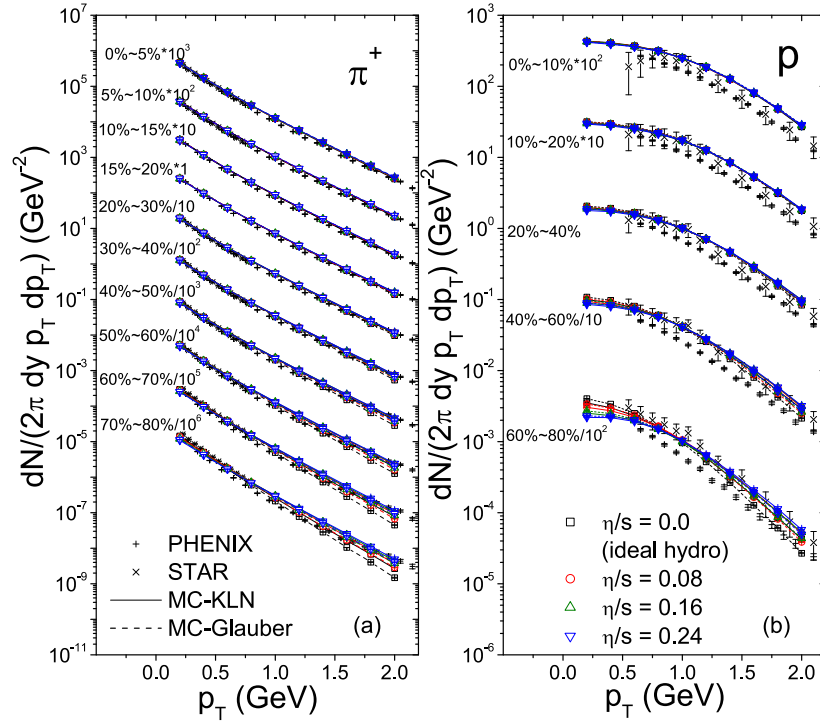
Figure 23. Transverse momentum spectra of pions and protons for different centrality windows in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV compared to the VISHNU hybrid approach with different initial conditions and different values of shear viscosity over entropy density. Reprinted (figure) with permission from [200], Copyright (2011) by the American Physical Society.

shows the effect of hadronic rescattering. While pions and kaons are only mildly affected, the transverse momentum of protons increases by about 30% during the late stage due to rescattering. This can be attributed to the large cross sections of nucleons with fast-moving pions; the phenomenon is sometimes referred to as 'pion wind'.

5.3. Geometry of the fireball

Measurements of correlation functions between particle pairs provide information about the geometry of the emission region in heavy-ion collisions. The inferred size of the so-called 'region of homogeneity' reflects the kinetic freeze-out, since these measurements are performed on final state hadrons. This technique is originally applied in the size estimation of stars by Hanbury-Brown and Twiss (HBT). The shape of the correlation function is typically assumed to be Gaussian with the following parametrization:

$$
C(\vec{q}) = (1 - \lambda) + \lambda \, K_{\text{Coul}}(q_{\text{inv}})
$$
$$
\times \exp\left(-q_{\text{o}}^2 R_{\text{o}}^2 - q_{\text{s}}^2 R_{\text{s}}^2 - q_{\text{l}}^2 R_{\text{l}}^2 - 2q_{\text{o}}q_{\text{s}} R_{\text{os}}^2 - 2q_{\text{o}}q_{\text{l}} R_{\text{ol}}^2\right), \tag{41}
$$

where a value of the chaoticity parameter $\lambda < 1$ accounts for dilution of the HBT signal due to particle correlations from resonance decays, and the different radii (R_{o}, R_{s} and R_{l}) refer to different directions. The out-direction (R_{o}) points along the direction of the average
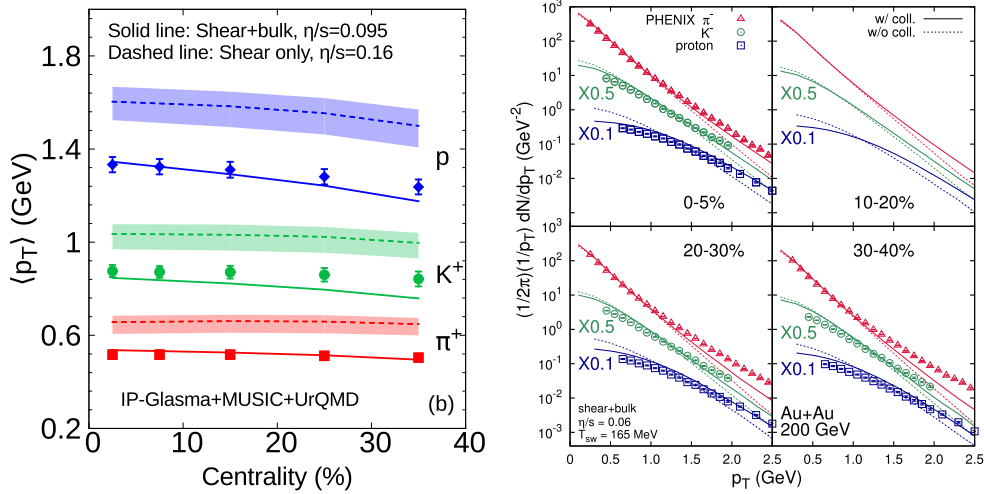
Figure 24. Left: mean transverse momenta as a function of centrality in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ GeV, the bands around the dashed lines indicate the effect of changing T_{switch} within the MUSIC+UrQMD hybrid approach compared to experimental data. Reprinted (figure) with permission from [201], Copyright (2015) by the American Physical Society.Right: transverse momentum spectra for pions, kaons and protons in different centrality bins calculated within a hybrid approach compared to PHENIX measurements in AuAu collisions at $\sqrt{s_{NN}} = 200$ GeV and the difference between the dashed and full lines indicates the effect of hadronic rescattering. Reprinted (figure) with permission from [202], Copyright (2018) by the American Physical Society.

momentum of the particle pair, the long-direction (R_l) along the beam direction and the side-direction R_s denotes the third orthogonal direction. q_{inv} denotes the relative momentum of the two particles measured in the rest frame of the particle pair; modifications of the relative wave function of the pair is most conveniently expressed in terms of this quantity. The contribution proportional to λ accounts for quantum correlations. When the data are insufficient for a three-dimensional analysis, the correlation function is often assumed to be spherically symmetric,

$$C(\vec{q}) = (1 - \lambda) + \lambda \, K_{Coul}(q_{inv})\exp(-|\vec{q}|^2 R_{inv}^2) \qquad (42)$$

with a single Gaussian source radius R_{inv}. A review on the topic can be found in [31]. As a function of mean transverse mass of the particle pairs, there is a global decreasing trend [204, 205] This can be easily understood due to the fact that faster particles are emitted earlier in the evolution and therefore from smaller source sizes than slower ones.

Identical particle correlations have also been measured for neutral and charged kaons and for protons [206]. The results are generally consistent with those obtained for pions, but slightly smaller source sizes have been measured for kaons perhaps pointing to their earlier kinetic freeze-out. Information about source sizes can also be gleaned from unlike charged particle correlations, where the information resides in the factor $K_{Coul}(q)$ that accounts for final-state Coulomb and other interactions between the emitted particles (see e.g. [207]). In the future, looking at HBT correlations for dileptons and photons can provide more information about the geometry of the system during the evolution since electromagnetic probes leave the fireball from all stages without disturbance from rescattering.
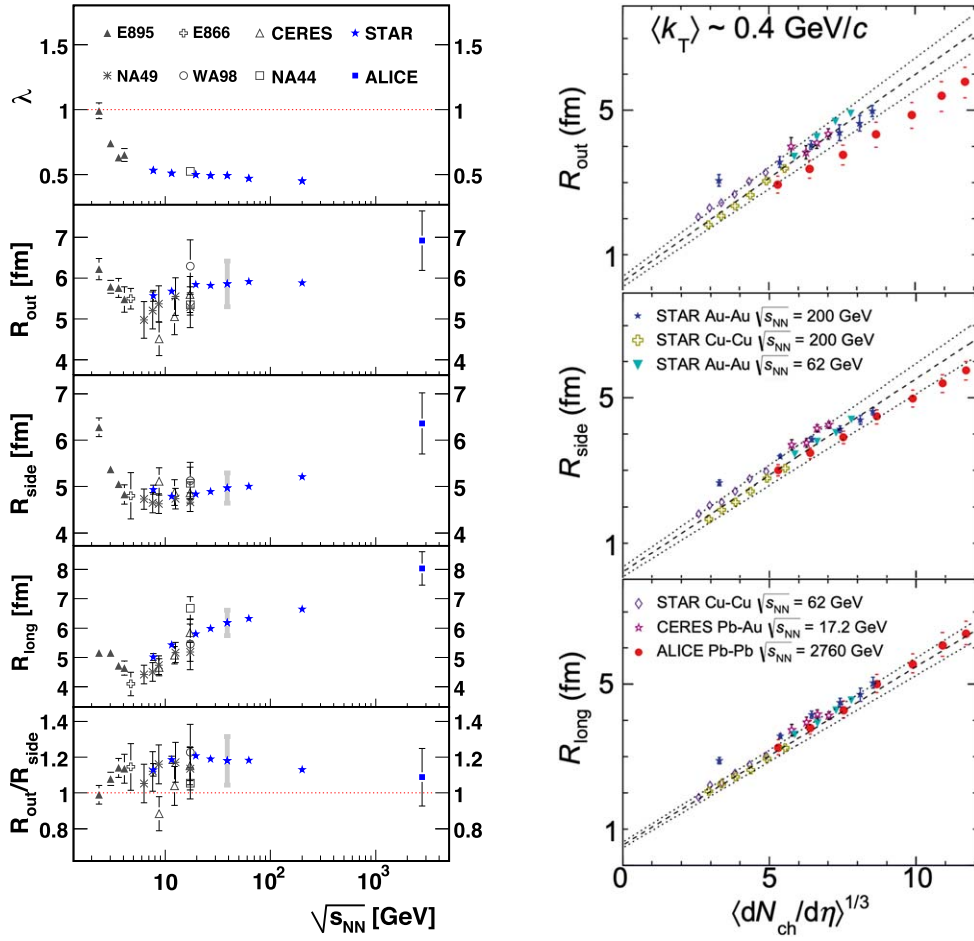
Figure 25. Left panel: collision energy dependence of the λ parameter, the HBT radii, and the ratio R_o/R_s derived from identical pion pairs. Reprinted (figure) with permission from [204], Copyright (2015) by the American Physical Society. Right panel: dependence of the HBT radii on the charged particle pseudorapidity density $dN_{ch}/d\eta$ for different collision system and collision energies. The dashed straight lines are linear fits to the data excluding ALICE. Reproduced from [205]. CC BY 3.0.

The excitation function of the radius parameters as well as the λ parameter and the R_o/R_s ratio is summarized in the left panel of figure 25. The ratio of the two transverse radii is predicted to be sensitive to the lifetime of the system. Interestingly, this ratio stays rather constant over a large range of beam energies. The fraction of correlated particles λ is higher at lower beam energies and saturates at higher energies to a lower constant value, which points to an increasing fraction of particles originating from longer-lived resonances in the sample. R_{out} and R_{side} stay more or less constant as a function of beam energy while R_{long} is rising significantly indicating a longer lifetime of the overall system evolution.

The right panel of figure 25 shows the dependence of the HBT radii on the charged particle rapidity density [205]. If the particle freeze-out always occurs at the same density, the radii should be linear functions of $(dN_{ch}/d\eta)^{1/3}$. The data, especially those for R_l, are roughly
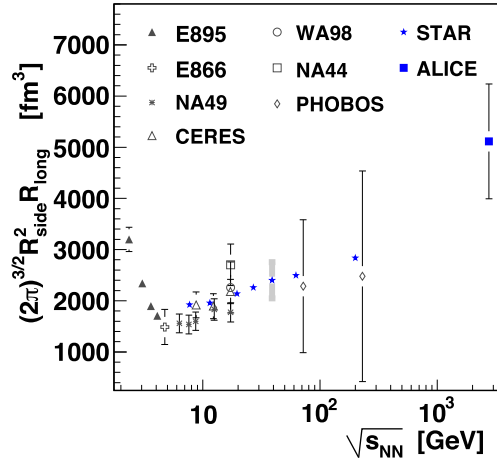
Figure 26. Volume of the emission region extracted from HBT radii of identical pions in Au+Au (Pb+Pb) collisions as a function of collision energy. Reprinted (figure) with permission from [204], Copyright (2015) by the American Physical Society.

consistent with this hypothesis. By combining the transverse HBT radii one can infer a measure of the fireball volume at kinetic freeze-out as shown in figure 26. These volumes can be compared to the one required by the thermodynamic description in section 5.1 and the qualitative behavior as a function of beam energy is consistent. In general, the analyzes of HBT radii and their comparison to hydrodynamic calculations allow one to connect coordinate and momentum space descriptions that are otherwise hard to achieve, see e.g. [153]. The complexity of this connection was analyzed in [208] where it was shown that multiple aspects of the dynamical evolution influence the size of the HBT radii and that it is non-trivial to get the bulk evolution in agreement with the correlation measurements.

A large value of the R_o/R_s ratio was proposed as one of the signatures of a first-order phase transition as a long duration of particle emission increases the outward-directed correlation length R_o. When the system undergoes a phase transition and the pressure is reduced significantly, the lifetime is expected to be significantly extended [210, 211]. In such a scenario one expects an increase in the R_o/R_s ratio as a function of beam energy. Figure 27 shows a calculation within a hybrid approach based on UrQMD initial conditions, (3+1)-dimensional ideal hydrodynamic evolution, and final hadronic transport within UrQMD. Different switching energy density criteria as well as the equation of state influence the lifetime. The noticeable result is that the only curve that qualitatively follows the slight peak as a function of beam energy that the experimental data from the NA49 collaboration suggests is the one with a first-order phase transition. Detailed correlation observables have the potential to hint at a first-order phase transition between hadron gas and quark-gluon plasma although, when the dynamical evolution is modeled more realistically, the effect is much smaller than originally estimated.

5.4. Anisotropic flow and the 'Perfect' fluid

The anisotropic flow coefficients as defined in section 2.4 contain interesting information about the properties of the quark-gluon plasma as well as the initial conditions. The angular modulations of collective flow in the plane transverse to the beam axis are analyzed in a Fourier decomposition. Non-vanishing values of v_n have been measured to high precision at
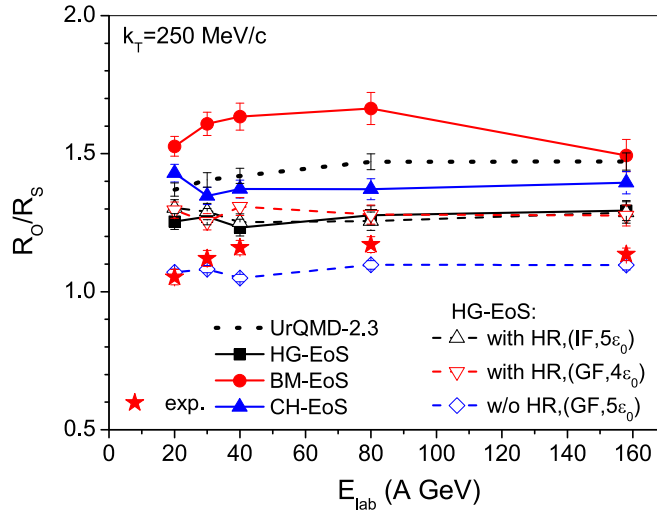
Figure 27. R_o/R_s ratio in central Pb+Pb collisions measured by NA49 as a function of beam energy compared to several calculations within a transport+hydrodynamics hybrid approach. Reproduced from [209]. CC BY 3.0.

RHIC and LHC as a function of transverse momentum, particle species and centrality. When hydrodynamic calculations are tuned to describe the yield and transverse momentum spectra for a system, the anisotropic flow can be very well predicted [120, 121]. This central finding that lies at the foundation of our 'standard model' of heavy ion collisions is based on a $\leqslant 10\%$ effect (for the average values) on the background radial expansion discussed in the previous section 5.2. The left panel of figure 28 shows one such example of a (3+1)-dimensional viscous hydrodynamic calculation for Au+Au collisions at $\sqrt{s_{\mathrm{NN}}} = 200$ GeV that fits the charged particle anisotropic flow as a function of transverse momentum.

The agreement of hydrodynamic calculations with anisotropic flow measurements forms the basis for the paradigm that the quark-gluon plasma exhibits an extremely low ratio of shear viscosity to entropy density and is therefore one of the most 'perfect' fluids in nature. Interestingly, this feature is shared with ultra-cold matter that can be probed in the laboratory, where atoms are confined in a trap in an almond-shaped configuration. Once released the particle stream along the pressure lines and the coordinate space anisotropy is transformed into a momentum space anisotropy [212]. The complex response function of the nonlinear hydrodynamic evolution can be dissected into eccentricities of different order and combinations of eccentricities that lead to the final-state flow coefficients (see e.g. [213]).

Higher order flow coefficients probe smaller-scale structures and are more sensitive to the properties of the plasma. Figure 28 (right) demonstrates the effect of increasing the shear viscosity-over-entropy density ratio from 0 to 0.08 and 0.16. Higher viscosities lead to smaller flow coefficients since the initial state structures are diluted more quickly. Without initial state fluctuations the odd flow coefficients would be zero by symmetry. But since 2010 it has been recognized that triangular flow (and higher coefficients) are actually non-zero resulting from small scale structures in the initial state [106]. Figure 29 shows the first measurement of triangular flow as a function of the number of participants in Au+Au collisions at $\sqrt{s_{\mathrm{NN}}} = 200$ GeV. While all calculations agree for elliptic flow in central collisions, the v_3 measurement is highly sensitive to the initial state structures (see [215] for a review). On the one hand, this poses the challenge that collective flow is not only sensitive to
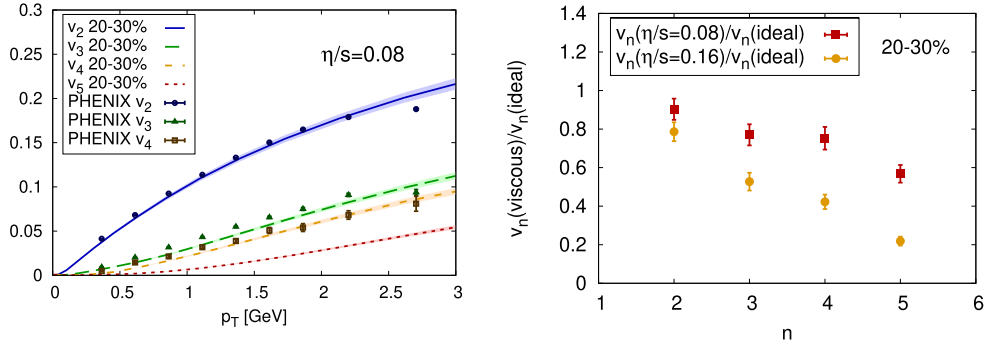
Figure 28. Left: v_n coefficients in semi-central Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV calculated in viscous hydrodynamics with Monte-Carlo Glauber initial conditions and Cooper–Frye freeze-out on a hypersurface of constant temperature. Right: ratio of v_n coefficients for different values of effective shear viscosity over entropy ratio. Reprinted (figure) with permission from [214], Copyright (2012) by the American Physical Society.
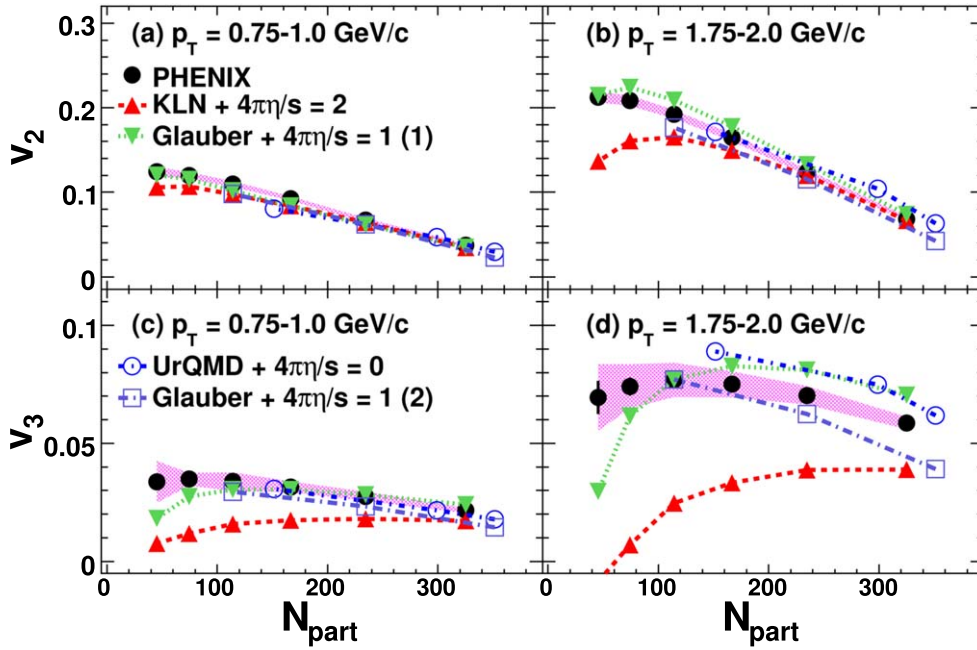
Figure 29. Comparison of $v_2\{\Psi_2\}$ (panels (a) and (b)) and $v_3\{\Psi_3\}$ (panels (c) and (d)) in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV as a function of centrality compared to different fluid dynamic calculations based on various initial conditions ('MC-KLN + $4\pi\eta/s = 2$' and 'Glauber + $4\pi\eta/s = 1$ (1)' [216]; 'Glauber + $4\pi\eta/s = 1$ (2)' [217]; and 'UrQMD' [218]). Reprinted (figure) with permission from [219], Copyright (2011) by the American Physical Society.

the medium properties but also to the details of the initial state. On the other hand, this also offers the opportunity to learn something about the initial state created by two nuclei colliding close to the speed of light.

By systematic improvements of theory and experiment it has been confirmed that all observed collective flow data are consistent with a quark-gluon plasma that exhibits a very low specific shear viscosity that is close to the lower bound predicted by AdS/CFT for a strongly coupled gauge plasma of $\eta/s = (4\pi)^{-1} = 0.08$. The averaged flow coefficients as a function of centrality are the most sensitive observables for this purpose. The full event-by-event distributions of the flow coefficients are mainly sensitive to the initial state and its fluctuations, while the viscosity mainly influences the average. Nowadays, this complex inverse problem involving many parameters and many observables is attacked with Bayesian analysis techniques as discussed in section 5.6.

The flow data for identified particles are sensitive probes of the hadronization mechanism and the dynamics of hadronic rescattering in the late stage that can be described by transitioning to a transport treatment for the fireball evolution (see [146] for a review). Additional information on the details of the initial conditions and the hydrodynamic response can be inferred from a vast amount of data on correlations of flow coefficients of different order, the correlation of the flow coefficients with transverse momentum and similar more advanced observables.

Even at the highest currently available collision energies the hot matter produced in the collision is not entirely boost invariant. One way by which this shows up is the gradual decorrelation of the event plane angles Ψ_n (9) measured in different pseudorapidity windows $\Delta\eta$. The correlation coefficient $r_n(\eta)$ is defined in terms of the complex flow vectors $\mathbf{q}_n = \sum_i w_i e^{in\phi_i}/(\sum_i w_i)$ measured in each event, where the sum runs over all particles in a selected pseudorapidity window and the w_i account for detector efficiencies [220], as

$$r_n(\eta) = \frac{\langle \mathbf{q}_n(-\eta)\mathbf{q}_n^*(\eta_0) + \mathbf{q}_n(\eta)\mathbf{q}_n^*(-\eta_0)\rangle}{\langle \mathbf{q}_n(\eta)\mathbf{q}_n^*(\eta_0) + \mathbf{q}_n(-\eta)\mathbf{q}_n^*(-\eta_0)\rangle}. \tag{43}$$

Here brackets indicate an event average, the asterisk indicates the complex conjugate, and η_0 denotes a reference pseudorapidity window usually chosen at a far-forward or far-backward pseudorapidity. Figure 30 shows the longitudinal decorrelation of the event planes Ψ_n for $n = 2, 3, 4$ in Xe+Xe collisions at the highest LHC energy. The $n = 2$ event plane is seen to decorrelate less rapidly than the $n = 3, 4$ planes. This phenomenon can be attributed to the predominantly geometric origin of elliptic flow, whereas the higher anisotropic flow coefficients are more sensitive to initial density fluctuations, for odd n exclusively so. Results for different energies at RHIC, shown in figure 31, show that the decorrelation occurs much more rapidly at lower energies, in agreement with the expectation of larger deviations from boost invariance.

The current theoretical understanding is that the decorrelation can be traced back to rapidity-dependent fluctuations in the gluon densities of the colliding nuclei and to the gluon-gluon interactions that seed the initial state of the fireball. In addition, hydrodynamic fluctuations during the expansion of the quark-gluon plasma probably also contribute to the decorrelation. Theorists have successfully modeled these fluctuations and found rough agreement with the measured decorrelation (see e.g. [223]).

In general, collective flow is among the main observables that are sensitive to the phase transition between the hadron gas and the quark-gluon plasma. This idea can be easily understood: since the equation of state is encoded in the pressure as a function of the energy density (and possibly net baryon density), the anisotropic flow will react differently
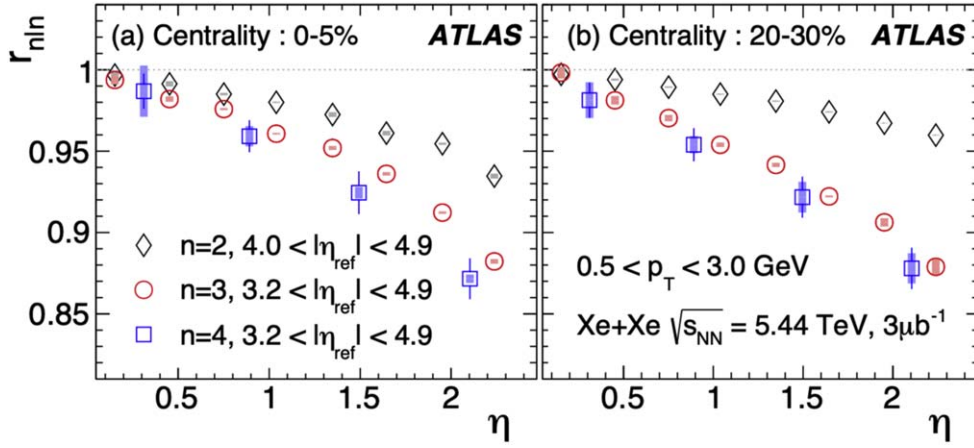
Figure 30. Longitudinal decorrelation of the event plane Ψ_n for $n = 2, 3, 4$ in Xe+Xe collisions in two different centrality windows at the highest LHC energy measured by ATLAS. Reproduced from [221]. CC BY 4.0.
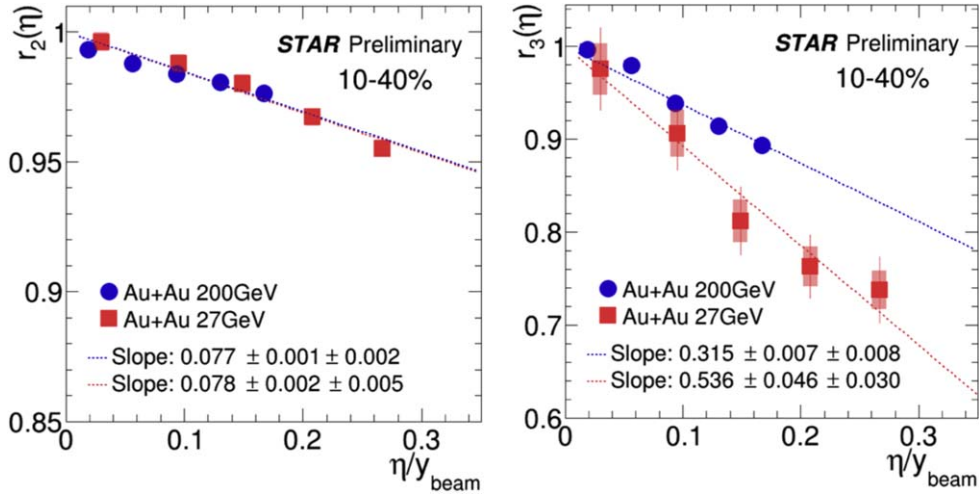
Figure 31. Longitudinal decorrelation of the event plane Ψ_n for $n = 2, 3$ in Au+Au collisions at RHIC at two different collision energies measured by STAR. Reprinted from [222], Copyright (2021), with permission from Elsevier.

depending on the pressure profiles that are probed by heavy ion collisions as a function of time. In other words, the flow coefficients that are primarily built up in a quark-gluon plasma with its very low specific viscosity should eventually vanish when the beam energy is reduced and the plasma is no longer formed. Figure 32 shows the integrated elliptic flow of charged particles as a function of the beam energy in Au+Au collisions. Interestingly, at low collision energies the elliptic flow is generated entirely by hadronic transport while at high energies the fraction of the flow built up during the hydrodynamic stage is more than 60%. At lower collision energies, the so-called directed flow v_1 is an important observable as well. $v_1 = \langle p_x/p_\perp \rangle$ quantifies the 'bounce-off' of particles in the reaction plane. This rapidity-odd
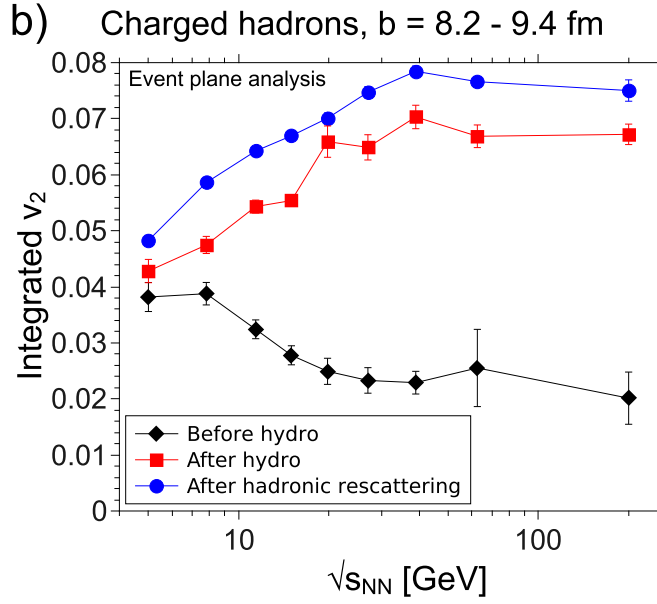
Figure 32. Development of elliptic flow in the different reaction stages of Au+Au collisions over the energy range of the RHIC beam energy scan within the UrQMD hybrid approach From [224]. Reprinted (figure) with permission from [224], Copyright (2013) by the American Physical Society.

flow coefficient can be summarized by fitting the slope dv_1/dy around midrapidity. The beam energy dependence of dv_1/dy for protons is expected to show non-monotonic behavior in case of a first-order phase transition [225]. In fact, understanding the collective flow and, in particular, the directed flow poses a challenge, since it is also very sensitive to the treatment of the interactions with the spectators and the interface between hydrodynamics and transport in hybrid calculations [226].

5.5. Small droplets of quark-gluon plasma

Long-range correlations among emitted particles in rapidity, akin to the phenomenon of anisotropic collective flow, are also found in high-multiplicity p+p and p+Pb collisions at the LHC [227–230]. Figure 33 shows two-body correlations of charged particles with transverse momenta in the range $1\,\text{GeV}/c < p_T < 3\,\text{GeV}/c$ versus azimuthal angle difference $\Delta\phi$ and pseudorapidity difference $\Delta\eta$ in high-multiplicity ($N_{ch} > 110$) events. The left panel is for p+p collisions, the right panel is for p+Pb collisions measured by CMS. The feature of interest in the near-side 'ridge' visible at $\Delta\phi \approx 0$ and extending over the full pseudorapidity acceptance $|\Delta\eta| \leqslant 4$. The ridge is clearly more pronounced in p+Pb collisions than in p+p collisions, which may indicate a larger degree of collectivity.

A crucial test whether the ridge phenomenon observed in these small collision systems is caused by geometry-driven collective flow was carried out by the PHENIX collaboration, which compared p+Au, d+Au, and ^{3}He+Au collisions at RHIC. Because the geometric structure of the light ions (p, d, ^{3}He) is distinctively different (see left panel of figure 34), one expects significant, geometry-driven differences in the elliptic and triangular anisotropic flow coefficients v_2 and v_3. This is, indeed, what was observed by the experiment, as seen in the right panel of figure 34. In v_2 the p+Au system stands out as less elliptically deformed; in v_3
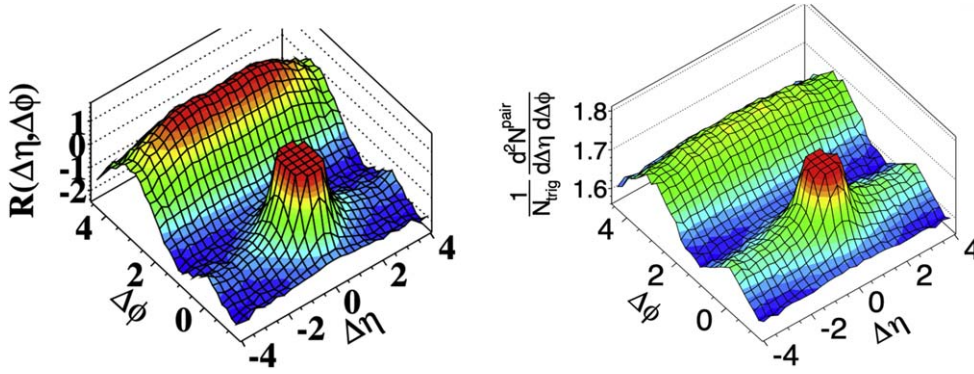
Figure 33. Two-particle correlations of charged particles with transverse momenta in the range $1\,\text{GeV}/c < p_T < 3\,\text{GeV}/c$ versus azimuthal angle difference $\Delta\phi$ and pseudorapidity difference $\Delta\eta$. Only high-multiplicity events with more than 110 charged tracks were selected for this analysis. Left panel: p+p collisions at $\sqrt{s} = 7$ TeV; Reproduced with permission from [227]. CC BY-NC 2.0 right panel: p+Pb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV. Reproduced with permission from [228]. CC BY-NC-ND 4.0.

the ^{3}He+Au system stands out as having a significantly larger triangular deformation. Signatures of collective flow are also seen in the particle-specific elliptic flow $v_2(p_T)$ in high-multiplicity p+Pb events at $\sqrt{s_{\text{NN}}} = 8.16$ TeV (see figure 35). The v_2 of K_s^0 and D^0 show a pronounced mass splitting at low p_T, but become approximately equal for $p_T > 5\,\text{GeV}/c$. The $v_2(p_T)$ for identified baryons and mesons exhibit the usual valence-quark number splitting.

Additional evidence for medium-like behavior in small systems comes from the gradual approach to full thermal (grand canonical) equilibrium of multi-strange baryon production [233]. Figure 36 shows the ratio of Ω and $\overline{\Omega}$ baryons to charged pions in p+p, p+Pb, and Pb+Pb collisions at various LHC energies as a function of $dN_{\text{ch}}/d\eta$. The steady increase with multiplicity is interpreted as evidence that the volume and lifetime of the hot quark-gluon plasma are growing with multiplicity. It is difficult to imagine that this steady growth is not accompanied by collective flow in response to the thermal pressure inside the fireball.

One often presented counter-indication against formation of a collectively behaving QGP in p+A collisions is that jet quenching has not been observed. Recent data from ALICE for $15\,\text{GeV}/c < p_T^{\text{ch}} < 50\,\text{GeV}/c$ jets in high-multiplicity p+Pb collisions limit the energy loss for a $R = 0.4$ jet cone to less than $0.4\,\text{GeV}/c$, more than an order of magnitude less than the energy loss measured for $R = 0.4$ jets in central Pb+Pb collisions [235]. This can likely be understood as a consequence of the fact that hard scattered partons need time to evolve down to the virtuality scale Q^2 at which the surrounding medium interacts with them [236]. For a parton with energy p_T, the kinematic virtuality within a transverse distance L from the scattering vertex is $Q^2 > p_T/L$. On the other hand, the virtuality scale associated with rescattering in the medium is

$$Q_{\text{med}}^2 \leqslant \int_{\tau_0}^{L} \mathrm{d}\tau \hat{q}(\tau) \sim (\hat{q}/T^3)T_0^3 \tau_0 \ln(L/\tau_0)), \tag{44}$$

where the subscript '0' indicates values at the initial moment of hydrodynamical evolution. For $L \leqslant 3$ fm in p+Pb one estimates $Q^2 > 2Q_{\text{med}}^2$, which means that the parton has left the medium before it reaches the virtuality scale at which the medium can begin to influence its QCD evolution.
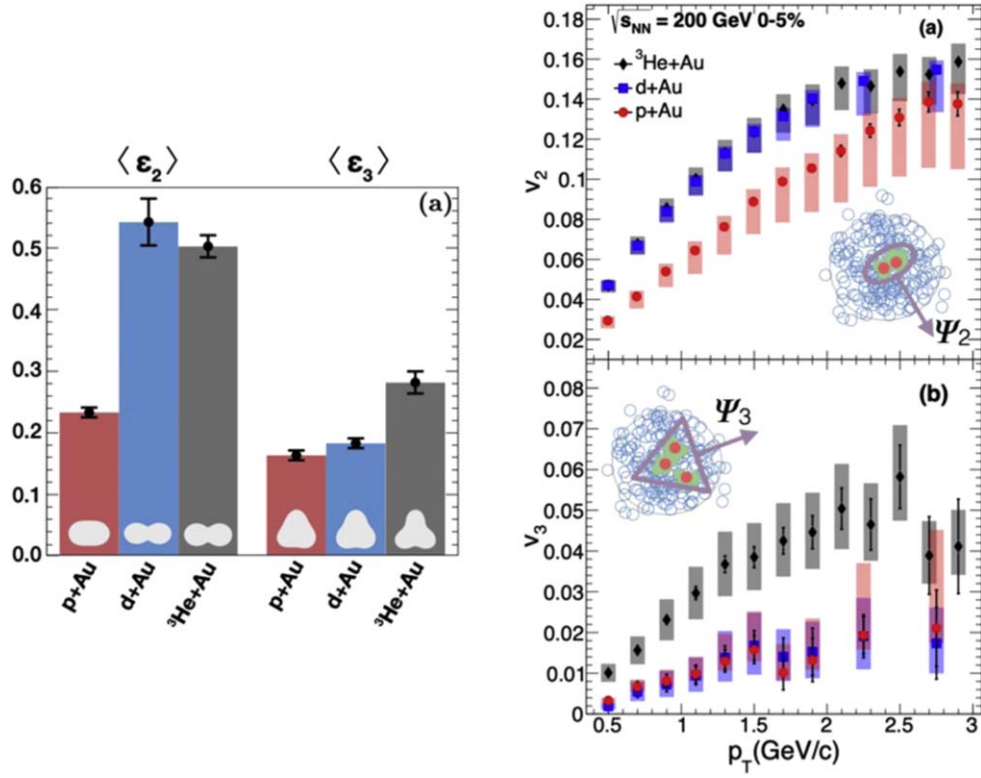
Figure 34. Left panel: Elliptic and triangular eccentricities of the proton (p), deuteron (d), and ^{3}He. Right panel: Elliptic and triangular flow coefficients $v_2(p_T)$ and $v_3(p_T)$ for the p+Au, d+Au, and ^{3}He+Au collision systems at the same energy $\sqrt{s_{NN}} = 200$ GeV. Reproduced from [231], with permission from Springer Nature.
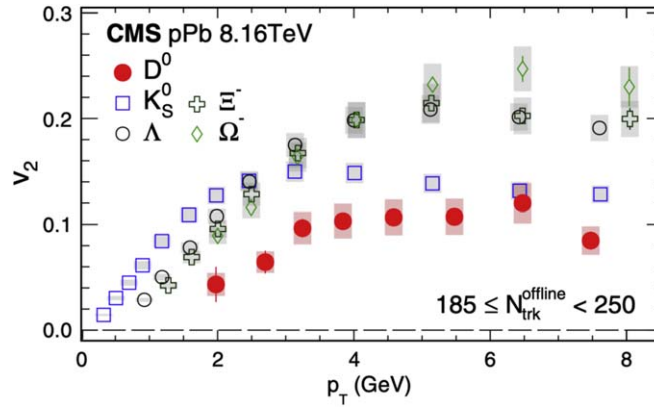
Figure 35. Elliptic flow coefficients $v_2(p_T)$ for identified hadrons in high-multiplicity p+Pb collisions at $\sqrt{s_{NN}} = 8.16$ TeV. The mass splitting between K_s^0 and D^0 mesons, as well as the usual baryon-meson splitting are clearly visible From [232]. Reproduced with permission from [232]. CC BY-NC-ND 4.0.
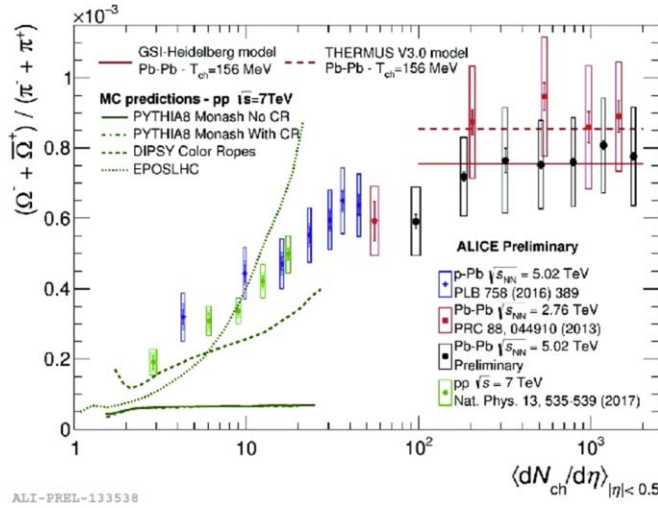
Figure 36. Yield ratio of Ω and $\overline{\Omega}$ baryons relative to charged pions ($\pi^{\pm}$) in p+p, p+Pb, and Pb+Pb collisions at various energies versus charged particle multiplicity $dN_{ch}/d\eta$. Reproduced from [234], with permission from Springer Nature.

The large flow gradients in the initial phase of the heavy ion reaction, especially for smaller collision system, severely strain the applicability of viscous hydrodynamics, which relies on a gradient expansion of the collective flow field. A measure of the applicability of hydrodynamics is the Knudsen number Kn, which describes the ratio of the mean free path λ to the characteristic system size L:

$$Kn = \lambda/L. \tag{45}$$

Here L could either be the size of the system or the typical length scale on which the fluid properties vary. A small Knudsen number, $Kn \ll 1$ implies that viscous hydrodynamics should provide a good description of the system. At early times during the collision and for small collision systems, such as p+A, this condition is violated.

Numerical simulations have shown, however, that viscous hydrodynamics provides a remarkably robust description even under conditions where the Knudsen number is not small. This behavior has been traced to the existence of hydrodynamic 'attractors', collective flow patterns that emerge as asymptotic solutions of the transport equations for many different initial conditions and even different microscopic transport dynamics (see [237] for a review). The existence of these attractors ensures remarkable insensitivity of the expansion dynamics against many uncertainties about the initial state and against the not well understood dynamics of the fluid before equilibration.

Figure 37 shows the hydrodynamic attractors for three different exactly solvable microscopic transport dynamics in the boost invariant (Bjorken) scenario. Flows for widely different initial conditions are seen to rapidly converge on the attractor within $\tau < 1/T$, well before the hydrodynamic gradient expansion becomes reliable.

Attractors not only apply to hydrodynamical flow patterns but also emerge in the framework of kinetic theory. In this setting, attractors have been found to smoothly interpolate between the very early stage of the collision where free streaming provides for a good approximation and the late stage when viscous hydrodynamics applies. This behavior is illustrated in figure 38, where the evolution of the energy density is shown for two different
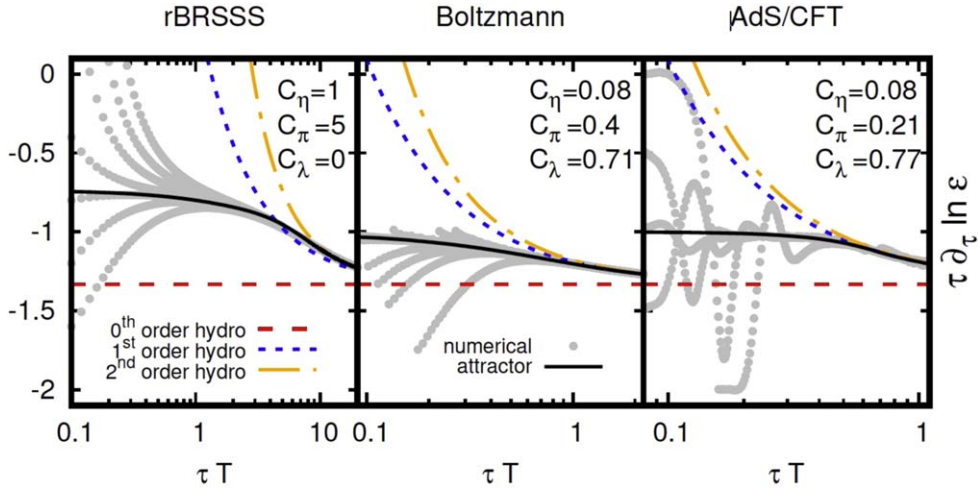
Figure 37. Hydrodynamic attractors, shown as solid black curves, for three different microscopic dynamics in the Bjorken model, depicted together with exact solutions (gray dots) for different initial conditions. For comparison the figure shows the results from ideal (here called '0th-order') hydrodynamics (red dashed line) and first-, and second-order viscous hydrodynamics (blue dashed and orange dash-dotted lines), which all approach the attractor asymptotically. The left panel shows results for the Borel resummed Baier–Romatschke–Son–Starinets–Stephanov (rBRSSS) formulation of conformal relativistic viscous hydrodynamics [238]; the central panel shows various solutions of the Boltzmann equation in the collision time approximation; the right panel shows results from numerical solution of the gravitational equations in the holographically dual Anti-de Sitter space (see section 3.3). The parameters C_η, C_π, C_λ denote dimensionless combinations of the coefficients of the dissipative terms in the BRSSS theory. For more detailed explanations, we refer to the original publication [239] and supplemental material cited therein.

values of the 't Hooft coupling constant $\lambda = g^2 N_c$ in the kinetic theory, which merge smoothly into the early (free streaming) and late (hydrodynamic) evolution. Altogether, these discoveries have strengthened our confidence in the robustness of the existing framework for modeling relativistic heavy ion collisions over the entire duration of the collision.

5.6. Quantitative constraints of equation of state and transport coefficients

Heavy ion reactions follow a rather complex dynamical evolution that involves switching between non-equilibrium descriptions for the initial (parton gas) and final (hadron gas) stage and viscous hydrodynamic evolution during the intermediate stage (QGP fluid). All the pieces of the theoretical framework come with their own parameters controlling, for example, the amount of initial state fluctuations, the transition criteria between descriptions, as well as the relevant physics properties, such as the equation of state and the transport coefficients. As discussed in section 5.4 the observables are often sensitive to more than one parameter in a non-trivial interplay. Note that there is a difference between the model parameters and the parameters of the system under consideration. The beam energy, centrality and other choices are the controllable settings that have to be matched between theory and experiment for meaningful comparisons.

One option to cope with such a multi-parameter, multi-observable problem is to apply the methodology of Bayesian multi-parameter analysis. The main advantage of such an approach
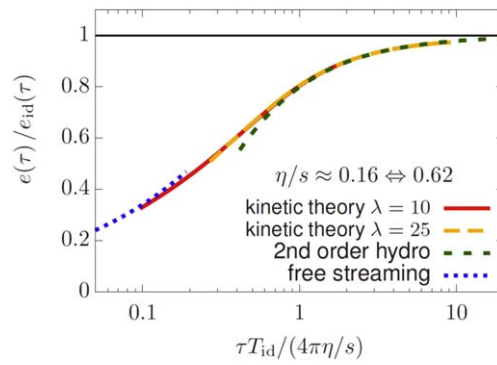
Figure 38. Attractor for kinetic transport in the Bjorken model, shown together with early stage expansion by free streaming (blue dots) and late stage hydrodynamic flow (green dashes). The results for the kinetic theory are shown for two values of the 't Hooft coupling constant $\lambda = g^2 N_c = 10$, 25 (solid red, dashed orange lines). Reproduced from [240]. CC BY 4.0.

is that one obtains a multi-dimensional sensitivity analysis as a by-product and statistically meaningful parameter extractions with quantified uncertainty. Important inputs are the statistical and systematic uncertainties of the measurements and of the theoretical models, which are sometimes difficult to assess. The pioneering studies that have applied Bayesian analysis for the first time in the field of heavy ion physics were carried out by the MADAI[12] collaboration [241]. By now, many groups have adopted the procedures and there are further collaborative efforts between statisticians and heavy ion groups (e.g. JETSCAPE[13], MUSES[14] and BAND[15] collaborations). The application of such advanced statistical techniques only makes sense, when the underlying theoretical description is established, otherwise no meaningful results can be obtained. If major physics components are missing ('unknown unknowns'), a Bayesian analysis will not produce valid results.

The following steps are usually taken:

 (i) Select the theoretical models and define the parameters to be varied.
 (ii) Define the prior ranges for the parameters as large as practically feasible.
 (iii) Run the full evolution model at enough points in the multi-parameter space, typically determined by Latin hypercube sampling.
 (iv) Select the observables to be employed in the study.
 (v) Employ principal component analysis (PCA) to determine orthogonal components of the observables.
 (vi) Construct an emulator for each observable trained on the full model runs.
 (vii) Use Markov Chain Monte-Carlo (MCMC) techniques to determine the posterior distribution of the parameters.
 (viii) Perform a closure test for verification that the posterior distribution covers results of full model runs that have not been used in the training of the emulator.

While explaining all the details of such a multi-parameter Bayesian analysis goes beyond the purpose of this review, the most important concepts will be covered here. Bayes' theorem is formulated as follows

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}. \tag{46}$$

The posterior distribution $P(A|B)$ containing the information of how likely parameter set A is under the set of observations B is given by the inverse probability of observables B given parameters A ($P(B|A)$) with a prior distribution $P(A)$ and normalized by the integral $P(B) = \int dA \, P(B|A)P(A)$. In this manner, Bayes' relation permits to invert the conditional probabilities to one that is easier to determine. In other words, we can formulate the question which parameters fit the observations best based on the outcome of simulations for a certain set of parameters. Since the integral in the denominator is usually hard to calculate, one generally ignores the denominator and restricts the analysis to relative probabilities instead of absolute ones.

Since evaluating the theoretical model for a multitude of parameter combinations becomes quickly too expensive even for current high-performance computing facilities, surrogate models are important to enable a wide parameter scan. Gaussian process emulators are constructed in the following way. Any arbitrary functional form describing the dependence of the observables on the parameters can be approximated with a Gaussian process GP that encodes the mean and the covariance

$$f(x) \sim GP(\text{mean}(x), \text{cov}(x, x')). \tag{47}$$

In general, one can envision such a Gaussian process emulator as a way to fit the training data points with associated quantified uncertainties for the values between training points. Therefore, a crucial input is the chosen kernel function

$$k(\vec{x}_p, \vec{x}_q) = C^2 \exp\left(-\frac{1}{2}\sum_{i=1}^{s}\frac{|x_{p,i} - x_{q,i}|^2}{l_i^2}\right), \tag{48}$$

for parameter values $\vec{x}_p$ and PCA values of observables $\vec{x}_q$. The parameters C and l_i have to be chosen with care not to overfit the training data, but also to not leave too much freedom in the construction of the emulator. Figure 39 shows the behavior of unconditioned draws (left) as well as the result after proper training (right). The 95% confidence interval indicated by the gray bands is fully constrained at the training points and increases further away from them as expected. Uncertainties at the training points can be incorporated by an additional white noise kernel into the Gaussian process.

In figure 40 the first application of Bayesian methods to heavy ion collisions is shown. A (2+1)-dimensional hybrid approach was applied to Au+Au collisions at the highest RHIC energy and to Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. The results were compared to data for bulk observables including particle spectra, anisotropic flow and HBT correlations. When drawing parameters weighted with their posterior distribution, it is apparent that the experimental data prefer an equation of state that is consistent with lattice QCD calculations. This is a very nice confirmation that our dynamical description actually prefers an equation of state similar to lattice QCD results based on observables at two different collider energies. This is an important finding that confirms the predicted properties of hot and dense QCD matter from experimental data.
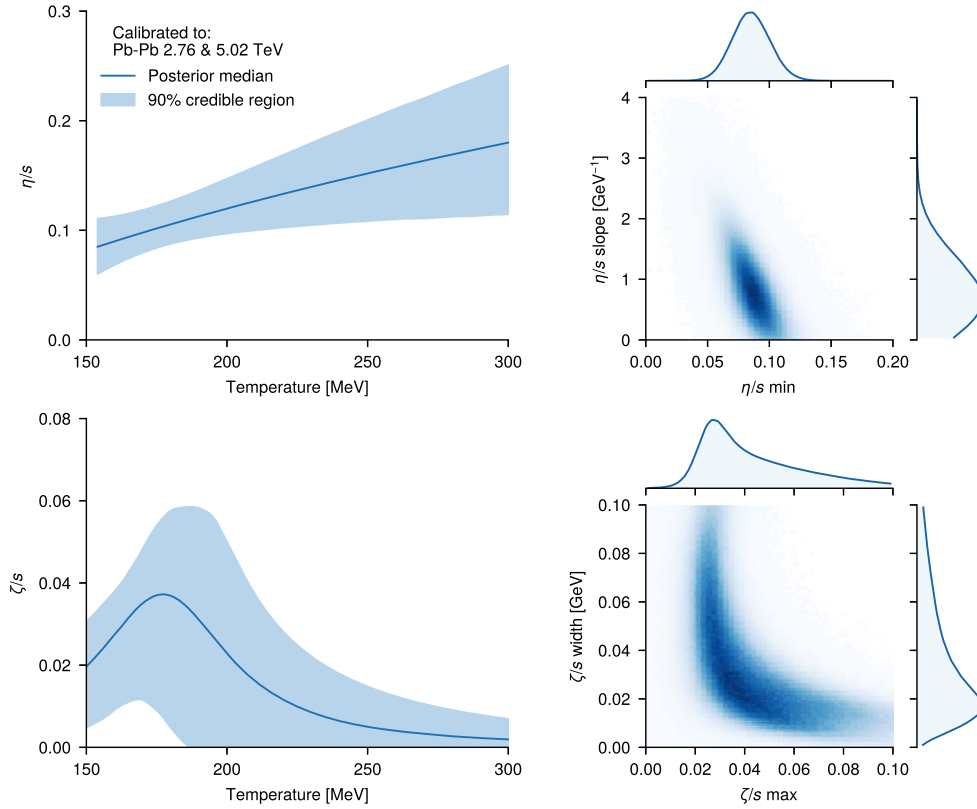
Figure 41 shows the posterior distributions for shear and bulk viscosity obtained from a hybrid calculation compared to experimental data. Within this work, 14 parameters were varied within the Trento+VISH2 + 1+UrQMD hybrid approach and the bulk properties of
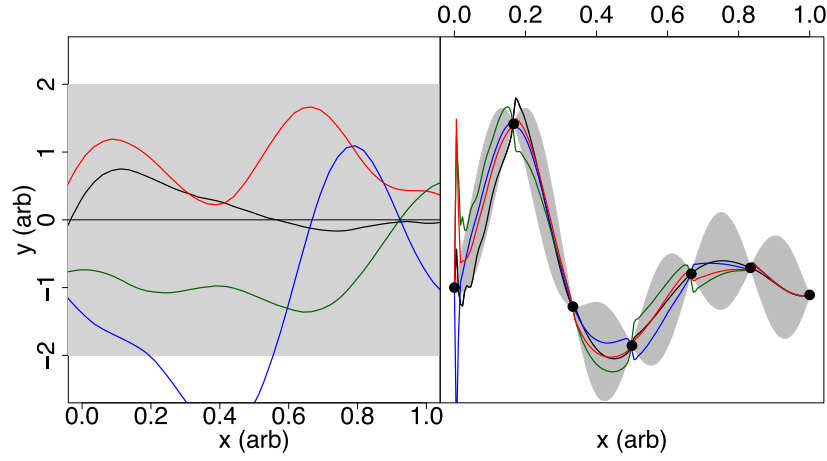
Figure 39. Left: Unconditioned draws from a Gaussian Process with a mean of zero and constant unit variance. Right: Draws from the same process after conditioning on 7 training points (black circles). The gray band in both panels is a pointwise 95% confidence interval. Reprinted (figure) with permission from [242], Copyright (2014) by the American Physical Society.
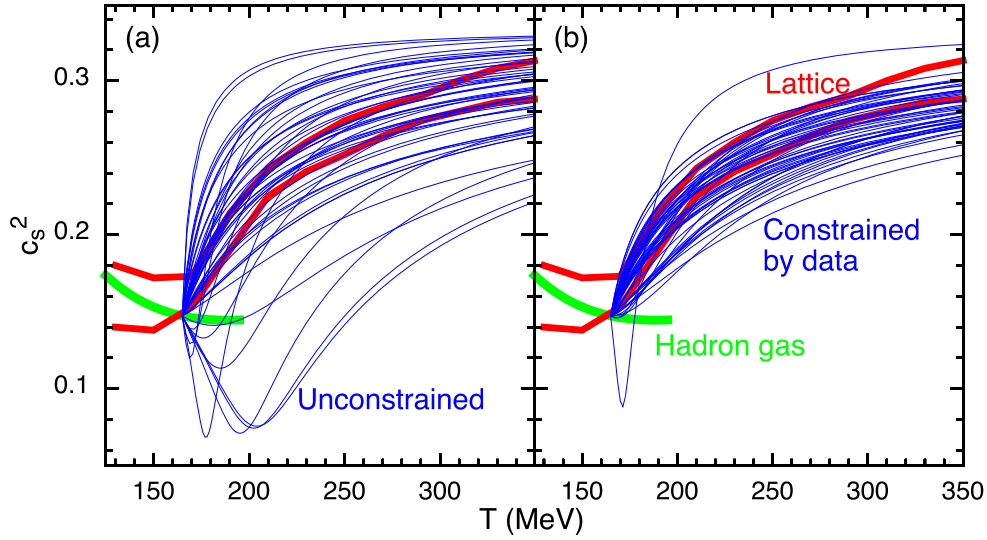
Figure 40. Fifty equations of state were generated by randomly choosing equation of state parameters from the prior distribution and weighted by the posterior likelihood. The thin blue lines in the left panel (a) are unconstrained samples from the prior; those in the right panel (b) were constrained by data from $\sqrt{s_{NN}} = 200\,\text{GeV}$ Au+Au collisions and $\sqrt{s_{NN}} = 2.76\,\text{TeV}$ Pb+Pb collisions. The two thick red lines in each figure represent the range of lattice equations of state shown in [43], and the short thick green line shows the equation of state of a non-interacting hadron gas. Reprinted (figure) with permission from [243], Copyright (2015) by the American Physical Society.

Figure 41. Posterior distributions for the parameters determining the temperature dependence of shear and bulk viscosity as well as the 90% confidence intervals shown for the transport properties of the quark-gluon plasma extracted with a Bayesian analysis of bulk properties at LHC energies within a hybrid approach. Reproduced from [244], with permission from Springer Nature.

hadrons measured in Pb+Pb collisions at 2 LHC energies were taken into account. The Bayesian analysis quantifies constraints on the temperature dependence of transport coefficients and properties of the initial state at the same time. Other analyzes have focused on small systems [245] or employed different implementations for the hydrodynamic and hadronic transport evolution [149].

Extending the list of observables to transverse momentum fluctuations and anisotropic flow correlations has been the subject of recent studies (see e.g. [246, 247]). Obtaining comprehensive quantitative conclusions on the QGP properties from all these efforts will be a major task in the future. While the different analyzes agree on major features, details like the preferred size of hot spots in the initial state or the maximum value of the bulk viscosity differ substantially between different analyzes. A very interesting further application of Bayesian analysis is demonstrated in [248], where the sensitivity of the parameters to potential future measurements in oxygen–oxygen collisions at LHC has been assessed based on 'best fit' parameters from prior studies.

While most Bayesian analysis have concentrated on the soft sector at RHIC and LHC energies, where the well established 'standard model' for the dynamical evolution of heavy ion reactions is applicable, there are a few studies targeting hard probes and lower beam

energies. In [249] the jet quenching transport coefficient $\hat{q}$ governing the transverse momentum transport of a hard parton traveling through a medium has been quantified within the JETSCAPE framework [160]. Employing the MATTER and LBT energy loss modules for high and low virtuality partons, respectively, the temperature and energy dependence of $\hat{q}$ can be constrained by a comparison to data from RHIC and LHC. The results are compatible with prior constraints from the JET collaboration [166].

In [250] a Bayesian analysis for the charm diffusion coefficient was reported, and in [251] results for the beam energy dependence of the shear viscosity-over-entropy density were presented. Both results are consistent with prior works and expectations for the qualitative behavior of the transport coefficients. As the range of Bayesian model-data comparison efforts grows, one must be careful to avoid applying Bayesian analyzes in regimes where the theoretical model is under insufficient control and account properly for the variability in modeling choices (see e.g. [149]). In the future, it will be rewarding to see more and more observables from the hard and soft sector confronted with a unified theoretical description. A straightforward extension of previous work will be to do a Bayesian analysis for hard probes on a well-calibrated (by Bayesian methods) soft background.

5.7. Event-by-event fluctuations

Modern experiments with heavy ion collisions typically record many millions of events under identical conditions. Due to quantum fluctuations there will always be event-by-event fluctuations, even if the species, beam energy, and impact parameter selection is restricted. In heavy ion reactions there are many sources of fluctuations, some of them trivial (like statistical fluctuations), some of them far from trivial (like the dynamical fluctuations associated with a critical endpoint). Here, we will concentrate on fluctuations of conserved quantities at high beam energies as well as fluctuations associated with the quark-gluon plasma phase transition at lower beam energies. The event-by-event fluctuations in the initial state resulting in higher order flow coefficients have been addressed in section 5.4.

As shown in section 5.1 the system formed in heavy ion collisions can be regarded to first approximation as being in thermal equilibrium. In that case, the fluctuations of conserved quantum numbers follow the expectations from the grand canonical ensemble as known from statistical mechanics (see [252] for a review). There is one caveat: when the charge under investigation is only produced in small quantities or the volume under consideration is close to the entire system, then exact conservation laws have to be taken into account in a canonical, instead of grand canonical, approach.

The idea to measure fluctuation observables originates from their association with properties of the system created in heavy ion collisions. For example, the fluctuations of mean transverse momentum are expected to reflect temperature fluctuations. If the system is hotter the particles are emitted with larger transverse momenta, while they obtain less transverse momentum in a colder fireball. For thermal fluctuations of conserved quantum charges, such as net baryon number B, strangeness S, and electric charge Q, the expected size is sensitive to the degrees of freedom that are active in the system. Figure 42 depicts the expected differences in fluctuations of net baryon number and net charge as a function of beam energy for a quark-gluon plasma and a hadron gas. Since the quarks carry fractions of baryon number and electric charge the corresponding fluctuations are smaller. Experimental data are found to be mainly consistent with the fluctuations expected from a hadron resonance gas. One possible explanation is that the hadronization process washes the partonic fluctuations out, and finally only hadronic fluctuations are observed. This has been demonstrated in a dynamic coalescence approach, where the hadronization process was modeled microscopically in an
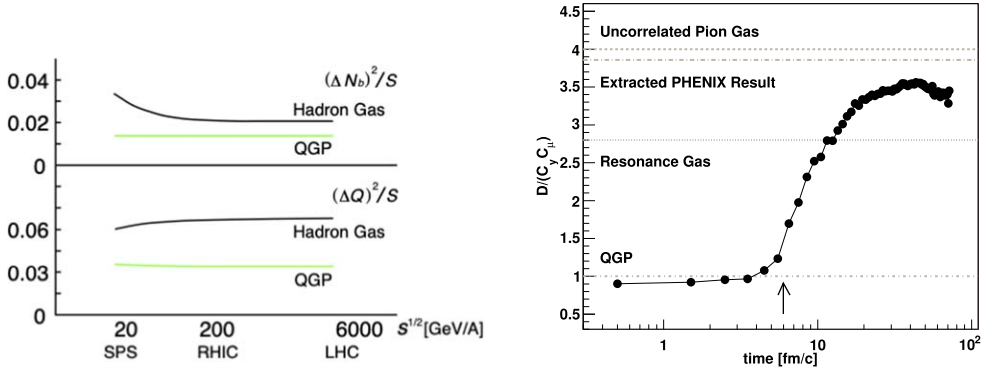
Figure 42. Left: Schematic drawing of the beam energy dependence of the net baryon number and charge fluctuations per unit entropy for a hadronic gas and a quark-gluon plasma. Reprinted (figure) with permission from [253], Copyright (2000) by the American Physical Society. Right: corrected charge fluctuations $\tilde{D}$ as a function of time within a hadronization model (arrow depicts time of hadronization) for Au+Au reactions at $\sqrt{s_{NN}} = 200$ GeV (full symbols). Also shown are the values for an uncorrelated pion gas, a resonance gas and a quark-gluon plasma. Reproduced from [254]. CC BY 4.0.

expanding system (see the right panel of figure 42). Nowadays the interest in the mean number of pairwise produced conserved charges has shifted to correlation observables, such as balance functions for charged particles [255, 256].

Another very intriguing application of fluctuation measurements is their direct comparison to lattice QCD calculations (see [257, 258] for a review and lecture notes on this topic). Fluctuations of conserved charges cannot be created by rapidly occurring local processes, only by much slower diffusion of particles between different regions of the fireball [253, 259]. Because diffusion proceeds much more rapidly in the QGP stage than in the hadronic gas phase, these fluctuations are expected to reflect the conditions just before hadronization. The thermal δQ_i are related to susceptibilities $\chi_n^{Q_i}$ by

$$\langle (\delta Q_i)^n \rangle = T^n \frac{\partial^n}{\partial \mu_i^n} \ln Z(T, \mu_i) = VT^3 \chi_n^{Q_i}, \tag{49}$$

where Q_i is the conserved charge of interest, T the temperature, μ_i the corresponding chemical potential and Z the partition function. Experimentally, the moments of the conserved charge distribution can be measured by standard statistical quantities:

$$M \text{ (mean)}, \quad \sigma^2 \text{ (variance)}, \quad S \text{ (skewness)}, \quad \kappa \text{ (kurtosis)}. \tag{50}$$

By associating each beam energy $\sqrt{s}$ with pairs of temperature and baryon chemical potential (T, μ_B), these moments correspond to certain susceptibilities as stated in equation (49). To remove the volume dependence, one usually considers ratios

$$\begin{aligned} S\sigma &= \chi_3/\chi_2 \quad ; \quad \kappa\sigma^2 = \chi_4/\chi_2 \\ M/\sigma^2 &= \chi_1/\chi_2 \quad ; \quad S\sigma^3/M = \chi_3/\chi_1, \end{aligned} \tag{51}$$

where the symbols refer to quantities defined in (49) and (50). Studies of the influence of volume fluctuations on such comparisons between experiment and lattice QCD calculations can be found in [260, 261].
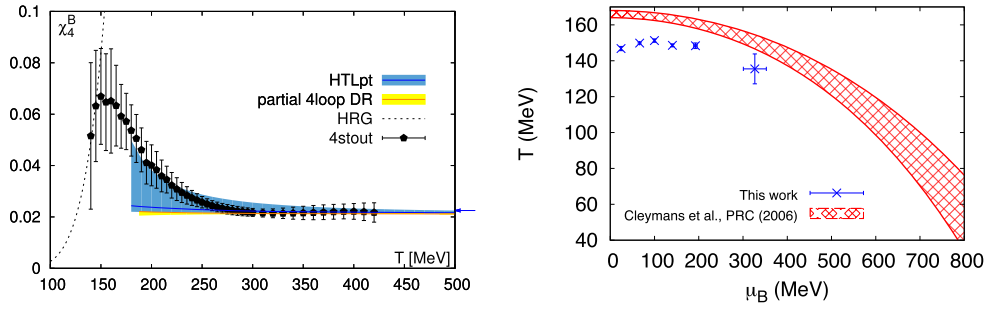
Figure 43. Left panel: quartic baryon number susceptibility χ_4^B defined in (49), calculated via lattice QCD (black dots with error bars, labeled '4stout'), in comparison with thermal perturbation theory (HTLpt) [60] and dimensional reduction (DR) [262] and the hadron resonance gas model (HRG). The arrow on the right edge marks the Stefan–Boltzmann limit for a free parton gas. Reprinted (figure) with permission from [263], Copyright (2015) by the American Physical Society. Right: Freeze-out parameters in the $(T - \mu_B)$ plane. The figure compares the chemical freeze-out curve [185] (red band) with the values obtained from the analysis of σ^2/M for net electric charge and net protons (blue symbols) comparing a hadron resonance gas model to STAR data. Reproduced from [264]. CC BY 3.0.

The left panel of figure 43 shows the lattice-QCD results for the quartic baryon number susceptibility χ_4^B as a function of temperature and compares them with the predictions from several effective approaches. At low temperature the lattice results merge smoothly to the hadron resonance gas model (HRG), while they approach perturbative calculations (HTLpt, partial 4-loop DR) at high temperature, eventually approaching the free parton gas limit.

The comparison of fluctuation measurements to susceptibilities of conserved quantum numbers calculated from lattice QCD is an alternative method to determine the chemical freeze-out conditions in terms of temperature and chemical potential. The results for net baryon number and net charge fluctuations are in rough agreement with the findings according to the thermal model (see 5.1). This agreement, which is illustrated in the right panel of figure 43, is expected since the lattice results in that regime correspond to a hadron-resonance gas.

In [187] it was suggested that comparisons between lattice calculations and strangeness fluctuations indicate a somewhat higher decoupling temperature than the one for light hadrons. This is in line with the expectations from microscopic models, where strange particles have a smaller cross section with other hadrons than the non-strange particles, most prominently protons and pions.

While it is exciting to directly compare *ab initio* lattice calculations to experimental data, one has to be aware of the limitations: depending on the kinematic cuts of the measurement, the comparison to a grand canonical ensemble calculation may be appropriate or not: typically only net proton fluctuations are measured and the mapping to net baryon number fluctuations carries uncertainties; also final state interactions and non-equilibrium effects may affect fluctuations.

The second important application of fluctuation observables is related to their expected sensitivity to the QCD critical endpoint. Finding signatures of the critical endpoint of the first-order phase transition between quark-gluon plasma and hadron gas is one of the main motivations for the heavy ion program at finite densities. The theory of phase transition dynamics predicts that the correlation length increases when the system passes through a
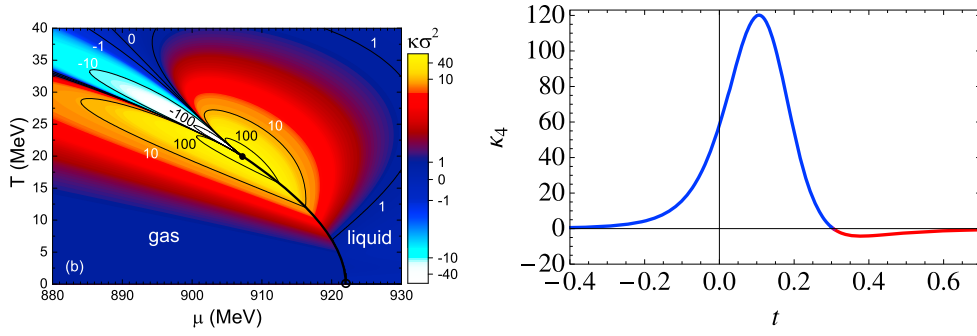
Figure 44. Left: The scaled kurtosis $\kappa\sigma^2$ calculated for symmetric nuclear matter in (T, μ_B) coordinates within Van der Waals equation of state for fermions. Reprinted (figure) with permission from [266], Copyright (2015) by the American Physical Society. Right: dependence of the baryon number kurtosis κ_4 on the reduce temperature parameter t the freeze-out curve (in arbitrary units). $t = 0$ corresponds to the location of the critical endpoint; $t < 0$ is the region where the phase transition is of first order. Reprinted (figure) with permission from [267], Copyright (2011) by the American Physical Society.

critical region. In particular, higher moments of the distributions of conserved quantum numbers are related to a higher power of the increased correlation length. In the idealized equilibrium scenario, the correlation length as well as the higher moments diverge. In heavy ion collisions this divergence is prevented by finite-size and finite-lifetime effects [265].

The fluctuation moments (51) can be determined experimentally as the cumulants C_n of the measured particle distributions [266]. For example, the kurtosis κ is obtained by evaluating the event average

$$\kappa\sigma^4 = C_4 = \langle(\Delta N)^4\rangle - 3(\langle(\Delta N)^2\rangle)^2 \tag{52}$$

with the variance $\sigma^2 = C_2 = \langle(\Delta N)^2\rangle$. Figure 44 (left) shows the scaled kurtosis $\kappa\sigma^2 = C_4/C_2$ in the phase diagram of nuclear matter. There are interesting structures and sign changes expected in the region around the critical endpoint of the QCD phase transition, which are analogous to those known in standard liquid–gas phase transitions. If one follows a typical freeze-out line in the phase diagram, figure 44 (right) depicts the expected beam energy dependence of the kurtosis involving a peak and then a dip structure when going from high to low beam energies.

Experimentally, measurements of the excitation function of the kurtosis of the net proton distribution have been carried out within the (first) RHIC beam energy scan program by the STAR collaboration and by the HADES collaboration at GSI (see figure 45 for a compilation of results). These are extremely challenging measurements since the impact of efficiencies and kinematic cuts for a multi-particle correlation measurement has to be controlled. As can be seen in figure 45 the data indicate a non-trivial behavior as a function of collision energy that is not compatible with the prediction of a transport (UrQMD) calculation which does not contain information about a critical point. The global net baryon number conservation influences the results, but not enough to bring them into agreement with experimental data. NA61 has looked at lower-order fluctuation observables as a function of beam energy and system size but could, so far, find no sign of critical behavior. In the future, higher precision data from the Beam Energy Scan II at RHIC, CBM at FAIR, MPD at NICA and other heavy ion physics programs at lower beam energies will complement the existing measurements.
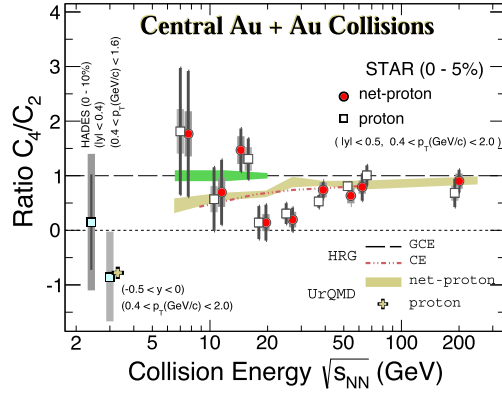
Figure 45. Collision energy dependence of the ratio of cumulants, C_4/C_2, for protons (open squares) and net protons (red circles) from central Au+Au collisions [268–270] (the symbols for protons are slightly shifted horizontally). Reprinted (figure) with permission from [271], Copyright (2021) by the American Physical Society.

On the theory side, there have been many developments targeting a dynamic non-equilibrium evolution through a critical endpoint (see [272] for a recent summary). All of these efforts are based on extending the fluid dynamics description to include thermal fluctuations. One approach involves adding thermal noise to the relativistic hydrodynamic evolution, which is numerically challenging. A complementary approach propagates the two-particle correlations on top of a hydrodynamic background ('Hydro+' formalism [273]). Both calculations agree in their finding that the non-equilibrium evolution has significant effects on the magnitude and behavior of the kurtosis, even though they are carried out in simplified settings. A full (3+1)-dimensional non-equilibrium evolution including the critical dynamics still remains a future challenge for a quantitative understanding of the experimental measurements.

Event-by-event fluctuations in heavy ion physics can also be used to select events of interest. The 'event shape engineering' technique groups, for example, events according to certain features like the magnitude of radial flow or certain anisotropic flow coefficients to provide further handles beyond centrality and beam energy. Machine learning techniques may make it possible to access interesting information by feeding information from single events into an artificial intelligence (AI) system. Of course this has to be done with great care, since the machine is not smarter than the best theoretical model on the market that was used to train the neural network. Systematic uncertainties inherent in such approaches are difficult to assess.

5.8. Hadronization and quark collectivity

Hadron production from heavy ion collision in the transverse momentum region below a few GeV/c exhibits two striking features: (1) The baryon-to-meson ratio, for both protons and antiprotons, in central Au+Au grows steadily with p_T for $p_T \lesssim 3$ GeV/c reaching a value three times as large as in peripheral collisions as shown in figure 46 [274]. (2) The elliptic flow coefficient $v_2(p_T)$ for mesons and baryons shows a distinctly different behavior, with the baryon v_2 saturating at larger values than the meson v_2 [275], as shown in figure 47 for several species of identified hadrons [276]. Especially noteworthy is the observation that ϕ-mesons (green symbols), which have approximately the same mass as nucleons, follow the proton
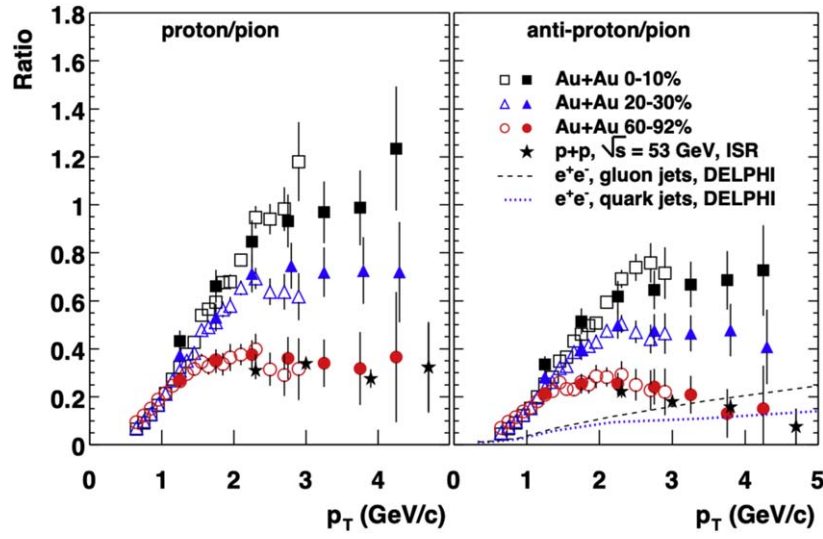
Figure 46. Proton-to-pion ratio (left) and antiproton-to-pion ratio (right) for Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV collisions for several centrality windows as a function of transverse momentum p_T. Reprinted (figure) with permission from [274], Copyright (2003) by the American Physical Society.
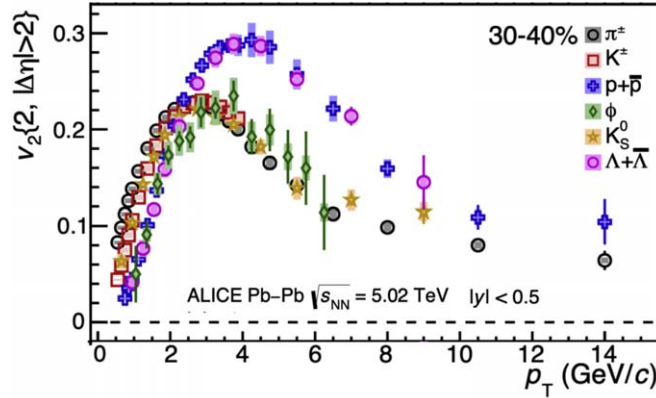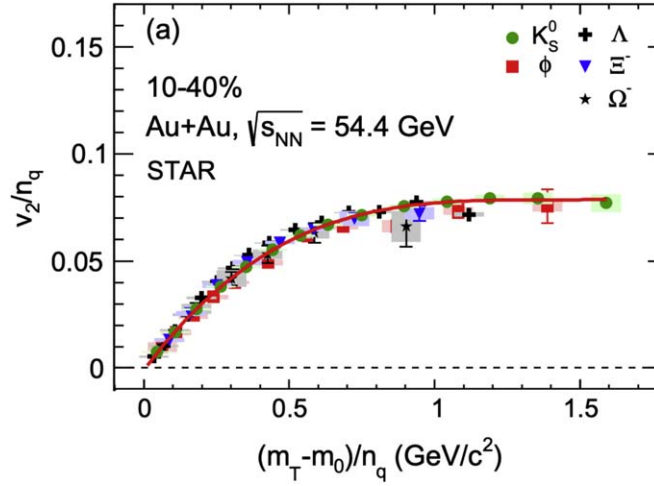

Figure 47. Elliptic flow coefficient $v_2(p_T)$ for various identified hadrons in $30 - 40\%$ central $\sqrt{s_{NN}} = 5.02$ GeV Pb+Pb collisions. (Adapted from [276]). Reproduced from [276]. CC BY 4.0.

data points at low p_T indicating hadronic collective flow, but follow the data points for other (much lighter) mesons at $p_T > 2$ GeV/c indicating that the collective flow is established at the quark level. This finding is consistent with the prediction that the formation of low-p_T hadrons occurs mainly during the bulk QGP hadronization, whereas higher-p_T hadrons are preferentially emitted from the QGP surface where quark recombination dominates (see figure 13).

Both phenomena can be explained by the mechanism of hadron formation by quark recombination (or coalescence) from the quark-gluon plasma, in which each valence quark

Figure 48. Scaled elliptic flow coefficient v_2/n_q for five different hadron species in $\sqrt{s_{NN}} = 54.4$ GeV Au+Au collisions as a function of the scaling variable $(m_T - m_0)/n_q$. The solid red line indicates a fit to the K_s^0 data From [282]. Reprinted (figure) with permission from [282], Copyright (2023) by the American Physical Society.

inherits the collective flow properties of the QGP fluid. Since baryons contain three valence quarks whereas mesons contain only two, baryons experience a stronger push from collective flow towards higher-pT than mesons [277–279]. For the elliptic flow $v_2(p_T)$ this mechanism implies a scaling law with valence quark number n_q [280]:

$$v_2(p_T) = n_q v_2^{(q)}(p_T/n_q), \tag{53}$$

where $v_2^{(q)}(p_T)$ denotes the elliptic flow coefficient for (anti-)quarks in the QGP.

Strictly speaking, the scaling with p_T/n_q can only be justified in the kinematic domain where hadron masses can either be neglected or described additively by constituent quark masses. In order to apply the scaling law heuristically over a wider momentum range, especially down to small momenta p_T, it is customary to compare the elliptic flow of different hadrons as a function of the scaled transverse mass $(m_T - m_0)/n_q$, where m_0 is the rest mass of the hadron and $m_T(p_T) = \sqrt{p_T^2 + m_0^2}$. This version of the valence quark scaling law has been found to be remarkably well obeyed by a large number of hadron species over a wide collision energy range. Two examples from the recent RHIC beam energy scan are shown in figure 48. A review of theoretical and experimental aspects of the quark recombination mechanism can be found in [281].

The valence quark scaling (53) has also been observed in identified particle emission patterns at the LHC, where the scaling is observed to hold even for the higher flow anisotropy coefficients v_3 and v_4 [276]. Figure 49 shows valence quark-number scaled flow coefficients $v_n(p_T/n_q)/n_q$ for $n = 2, 3, 4$ for several identified hadrons in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV.

The scaled flow coefficients in figure 49 exhibit broad peaks around $p_T/n_q \approx 1.5$ GeV/c. In the fragmentation-recombination scenario, this peak corresponds to a gradual transition to the fragmentation dominated regime. In the transition region it is plausible that quarks from a parton shower recombine with thermal, collectively flowing partons [283]. An implementation of this idea in a dynamical model (EPOS) describes hadron formation at intermediate
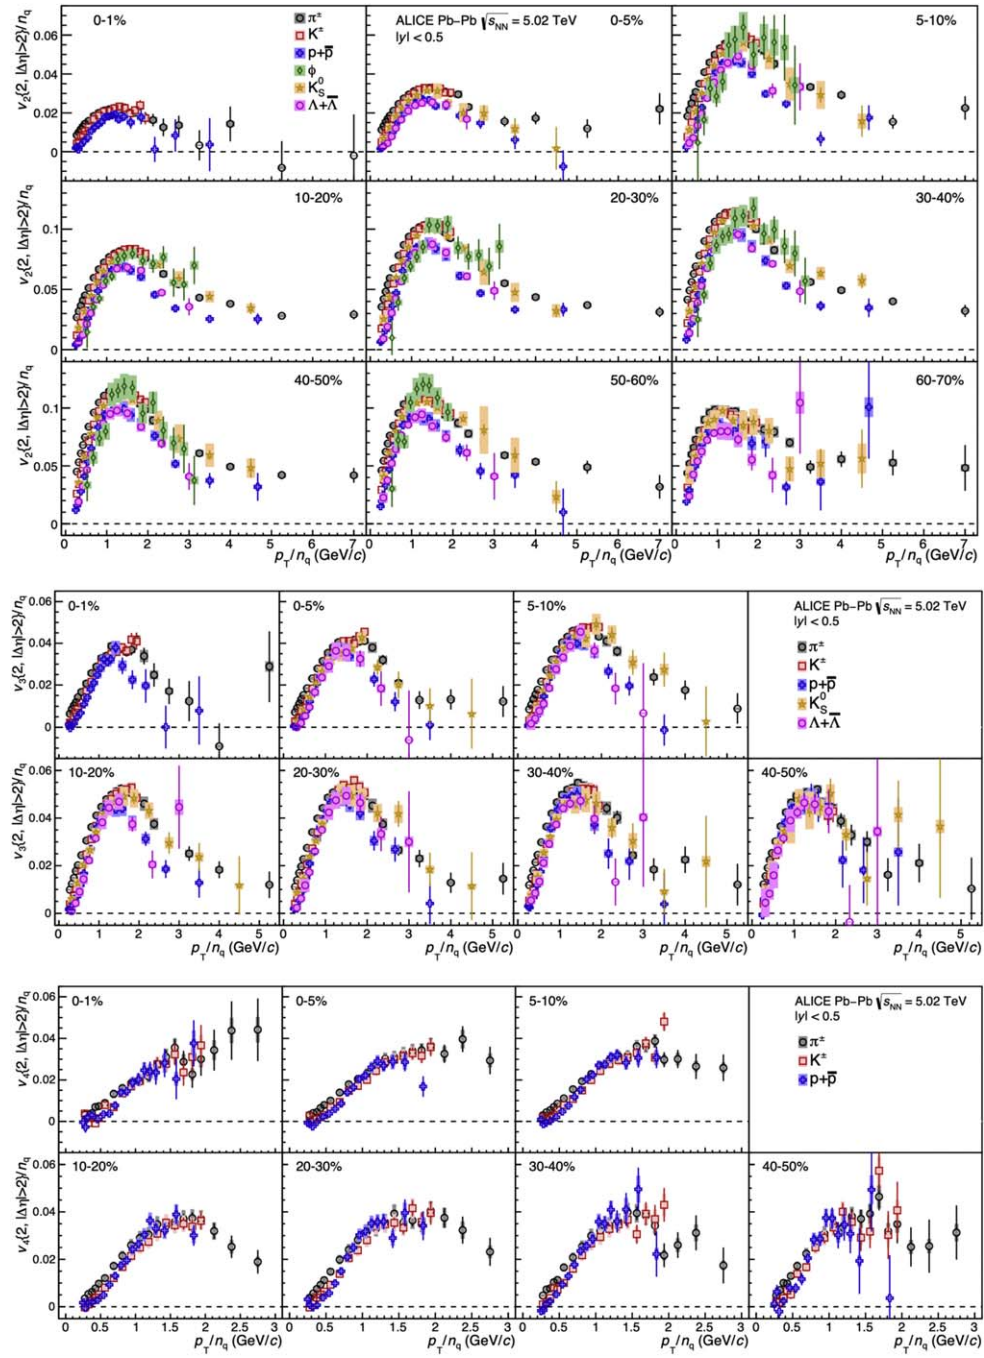
Figure 49. Valence quark-number scaled anisotropic flow coefficients v_n/n_q for six identified hadron species in $\sqrt{s_{NN}} = 5.02$ TeV Pb+Pb collisions as a function of the scaling variable p_T/n_q. Top panels: $n = 2$; central panels: $n = 3$; Bottom panels: $n = 4$. Note that the momentum scale in the bottom set of panels is different. Reproduced from [276]. CC BY 4.0.
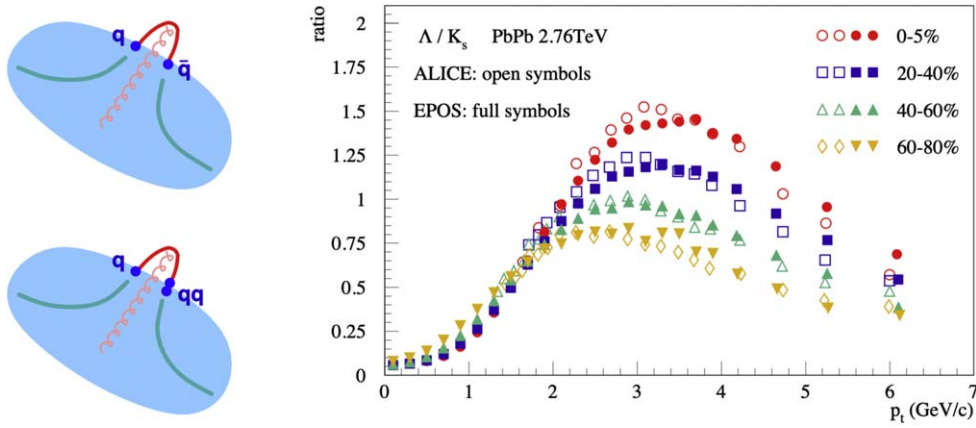
Figure 50. Left panel: Recombination of thermal quarks from the QGP fluid with energetic string fragments. Right panel: p_T-dependence of the hyperon-to-kaon ratio N_Λ/N_K for different centrality windows in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. ALICE data (open symbol) are shown together with results from the EPOS model (solid symbols). Reprinted (figure) with permission from [284], Copyright (2012) by the American Physical Society.

values of p_T as string (color flux-tube) fragmentation in the presence of a thermal parton fluid [284, 285]. Figure 50 compares ALICE data [286] for the p_T-dependence of the hyperon-to-kaon ratio N_Λ/N_K for different centrality windows in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV with results from the EPOS model that accounts for recombination of the leading parton with thermal partons as well as fragmentation by quark-pair production from the vacuum [285].

At low momenta ($p_T < 1.5$ GeV/c) hadronic rescattering affects baryon and meson flow differently and amplifies the mass splitting observed in the unscaled elliptic flow $v_2(p_T)$ [287]. Pions move much faster than baryons and push them out to larger p_T ('pion wind', see also the end of section 5.2), while heavy baryons have the opposite effect on pions and other mesons. Note that this mechanism does not work for ϕ-mesons, as these do not have large cross sections with pions or nucleons [288]. The increase of the mass splitting in $v_2(p_T)$ is clearly visible in figure 51, which shows $v_2(p_T)$ for pions, kaons, and protons calculated with and without hadronic rescattering in comparison with PHENIX data from Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV [289].

5.9. Jet quenching and parton energy loss

The emission of hadrons at high transverse momentum ($p_T > 6$ GeV/c) in relativistic heavy ion collisions is suppressed because high-momentum quarks and gluons lose energy when they propagate through the quark-gluon plasma [290, 291]. For light quarks or gluons the dominant energy mechanism is gluon radiation in association with scattering off a virtual gluon in the QGP. In the BDMPS-Z formalism of multiple scattering the scattering power of the QGP is encoded in the parameter $\hat{q}$, which describes the average squared momentum exchange with the medium per unit path length, $\hat{q} = d\langle q_T^2 \rangle/dx$ [292–294]. The BDMPS-Z approach is well suited to describe energy loss in a thick QGP. The GLV formalism, which is based on an opacity expansion, is better suited for a thin QGP. The higher-twist formalism [157] aims at the description of the full virtuality evolution of a jet created by an energetic quark or gluon and is expected to apply to both, short and long path lengths (see also [295]).
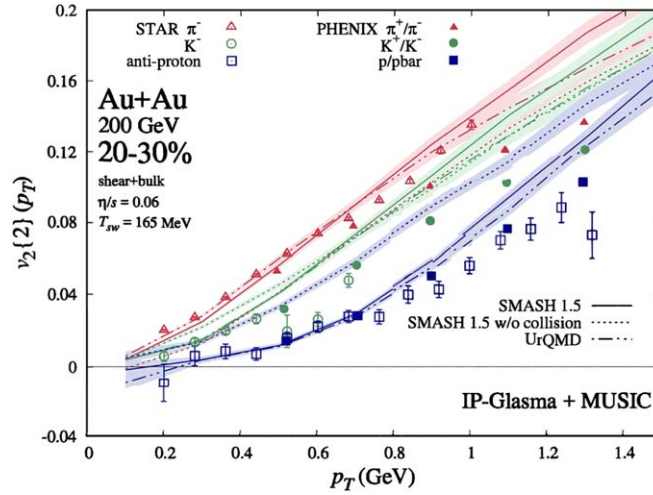
Figure 51. Elliptic flow coefficient $v_2(p_T)$ for pions, kaons, and protons calculated within a hybrid transport model [289] with and without hadronic rescattering in comparison with PHENIX data for Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV. The solid and dash-dotted lines show calculations including hadronic rescattering; the dotted lines are calculated without rescattering. Reproduced from [289]. CC BY 4.0.

Experimental evidence that high-p_T hadron suppression is a final-state effect comes most directly from the comparison of the nuclear modification factors R_{AA} (7) for hadrons with that for photons [296]. Figure 52 shows that $R_{AA} \approx 1$ for photons in Au+Au collisions at the top RHIC energy, whereas $R_{AA} \approx 0.25$ for pions and η-mesons. Similar results have been obtained for Z^0 or $W^{\pm}$ bosons in p+Pb and Pb+Pb collisions at LHC [297–301], where the data are consistent with nuclear modification of the initial-state parton distributions but show no final-state suppression of the boson yield when identified by their leptonic decay modes. However, as Z^0 bosons decay within 0.1 fm/c, the hadron showers created in the hadronic decay mode ($q\bar{q}$) may provide for a calibrated probe of the jet-medium interactions [302].

Schematically, the hadron spectrum can be expressed as a convolution of the parton distribution functions $f_i^{(A)}(x)$ in the colliding nuclei with the hard QCD scattering cross section and a fragmentation function $D_{i \to h}^{(\text{med})}(z)$ that is modified by the medium

$$\frac{\mathrm{d}N_h}{\mathrm{d}p_T}(p_T) = f_i^{(A)}(x_1) \otimes f_j^{(A)}(x_2) \otimes \frac{\mathrm{d}\sigma_{ij}}{\mathrm{d}p_T}(p_T/z) \otimes D_{i,j \to h}^{(\text{med})}(z). \tag{54}$$

It is usually assumed that the hadronization of the quarks and gluons contained in the evolving parton cascade occurs in vacuum outside the medium. For hadrons deriving from the leading parton in this cascade, the medium modification can then be expressed as a shift in the momentum fraction in the vacuum fragmentation function:

$$D_{i \to h}^{(\text{med})}(z) \approx \left\langle D_{i \to h}^{(0)}\left(\frac{z}{1 - z\Delta p_T/p_T}\right) \right\rangle, \tag{55}$$

where Δp_T is the momentum loss of the leading parton and $\langle \cdots \rangle$ indicates an average over the position and orientation of the hard scattering event. The magnitude of the suppression thus depends not only on the amount of energy loss but also on the steepness of the unmodified
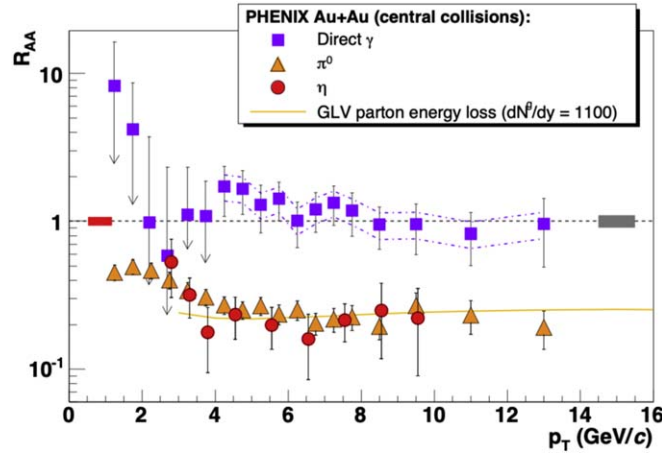
Figure 52. Nuclear suppression factors $R_{AA}(p_T)$ for direct photons, neutral pions and η-mesons in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV. Both meson species exhibit the same level of suppression, while photons are not suppressed. Reprinted (figure) with permission from [296], Copyright (2006) by the American Physical Society.
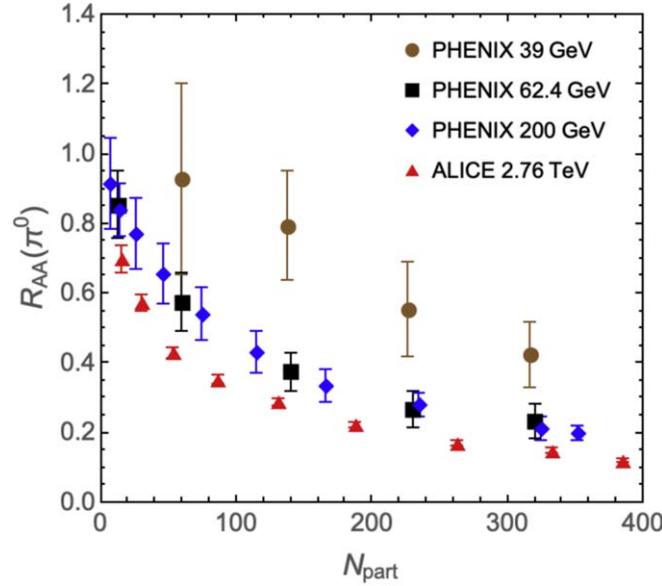
Figure 53. Nuclear suppression factors R_{AA} for neutral pions in Au+Au collisions as a function of the number of participant nucleons N_{part} at $\sqrt{s_{NN}} = 39$, 62.4, 200 GeV and Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. Reprinted (figure) with permission from [303], Copyright (2012) by the American Physical Society, Reproduced from [304]. CC BY 4.0.

hadron spectrum. As the spectrum becomes flatter at high collision energy, the same energy loss causes less suppression. In spite of this effect, experimental data from RHIC and LHC confirms that the suppression increases with collision energy, as shown in figure 53 for neutral pions with $p_T > 6$ GeV/c, implying a strong increase in the energy loss. This
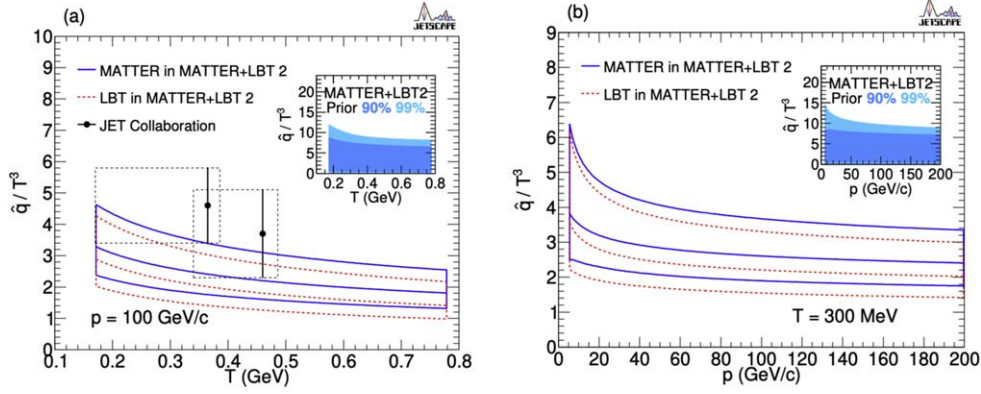
Figure 54. Reduced jet quenching parameter $\hat{q}/T^3$ for quark-initiated jets in a quark-gluon plasma determined by a Bayesian analysis. $\hat{q}/T^3$ is shown as a function of temperature T (left panel) and quark-momentum p (right panel). The solid blue and dotted red lines indicate uncertainty domains deduced by comparison of two different jet quenching algorithms with data. The black dots with error bars in the left panel show the values reported in the original JET Collaboration analysis [166]. Reprinted (figure) with permission from [249], Copyright (2021) by the American Physical Society.

observation agrees with expectations, as the energy loss parameter $\hat{q}$ grows rapidly with temperature: $\hat{q}/T^3 \approx 2 - 5$ (see figure 54).

The increase with N_{part} reflects the strong path-length dependence of the radiative energy loss which, for a medium of constant density, is approximately given by [294]:

$$\Delta E = -\frac{1}{2} C_2 \alpha_s \hat{q} L^2, \tag{56}$$

where C_2 denotes the SU(3) Casimir operator for the energetic parton. Figure 55 shows an attempt to relate the fractional momentum loss S_{loss} of energetic partons to the average squared path-length L^2 of the parton inside the medium. S_{loss} is obtained by comparing the spectrum of hadrons measured in $\sqrt{s_{NN}} = 200\,\text{GeV}$ Au+Au collisions with the binary collision-scaled hadron spectrum in p+p collisions in different centrality windows. The data are consistent with a linear dependence on L^2 as predicted by (56).

An independent assessment of the path-length dependence of parton energy loss can be obtained by measuring the azimuthal anisotropy of the nuclear suppression with respect to the collision plane. This anisotropy can be expressed by the Fourier coefficient $v_2(p_T)$, defined as

$$\frac{\mathrm{d}N(p_T)}{\mathrm{d}\phi} \propto 1 + 2v_2(p_T)\cos(2\phi) \tag{57}$$

as the 'elliptic flow' coefficient. The average difference in path length dL for partons emitted perpendicular to the collision plane compared with those emitted along the plane is quite large as shown in the left panel of figure 56, implying that a significant dependence on the emission angle relative to the collision plane is to be expected. This expectation is confirmed by data from Au+Au collisions at RHIC, see the right panel of figure 56 which shows a sizable value of v_2 for hadrons up to 10 GeV/c momentum. Data from Pb+Pb collisions at LHC for much higher-pT, shown in figure 57, reveal a strong correlation with the elliptic flow coefficient v_2
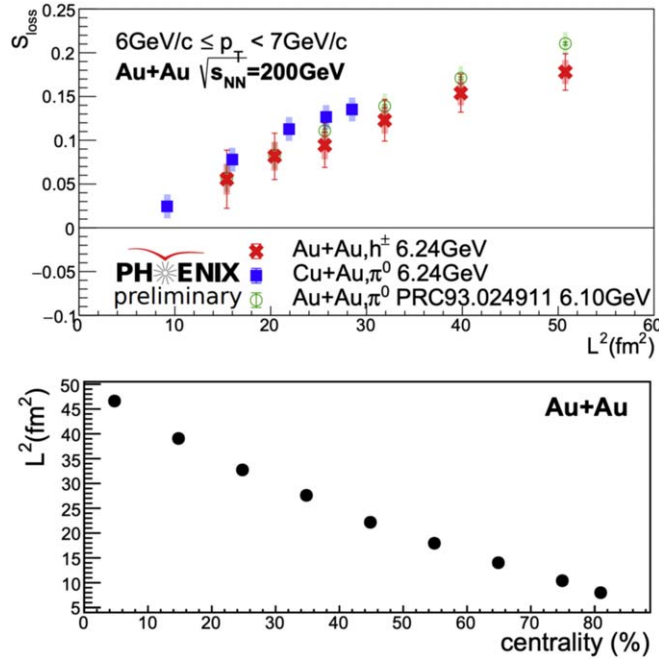
Figure 55. Upper panel: fractional momentum loss S_{loss} versus average squared path-length L^2 for several species of high-p_T hadrons emitted in Au+Au collision at $\sqrt{s_{\mathrm{N}}} = 200$ GeV collisions. When the parton hadronizes after leaving the medium, the fractional momentum loss of the hadrons should be equal to that of the primary parton. Lower panel: average squared path-length L^2 for hard-scattered partons in Au+Au collisions as a function of centrality. The linear dependence of S_{loss} on L^2 is clearly visible in the upper panel. Reproduced with permission from [305]. CC BY-NC-ND 4.0.

measured in the low-p_T region and indicate that the anisotropy at high p_T has the same geometric origin.

Unless they are absorbed by the QGP, the gluons that are radiated by a fast parton in the medium, remain part of the full jet. What fraction of the radiated energy is recovered in measurements of the full jet energy depends on the opening angle R (the 'jet radius') that is used to define the jet. Typical values used for such studies are $R = 0.2 - 0.6$. Smaller jet radii imply a larger energy loss. This trend is readily apparent in figure 58, which shows the relative magnitude of $R_{\mathrm{CP}}(p_T)$ for jets with total transverse momentum p_T and radius $0.2 \leqslant R \leqslant 0.5$. (See the text below equation (7) for an explanation of R_{CP}.)

Full jet quenching exhibits many of the same phenomena as the suppression of single hadrons at high p_T, except that all results depend quantitatively on the jet radius R. This means that full jet quenching measurements are not only sensitive to longitudinal energy loss, but also to the angular redistribution of the energy within the jet [309]. Measurements of full jet suppression thus enable more differential measurements, e.g. the study of how quenching modifies the jet shape in terms of the longitudinal momentum fraction of a hadron within the jet, $z_h = p_T^h / p_T^{\mathrm{jet}}$, and the relative angle $r < R$ of the momentum of a hadron with respect to the jet axis (see [310] for a review on jet measurements).

Often the jet shape $\rho(\xi, r)$ is expressed in terms of the variables $\xi = \ln(1/z)$ and r. Examples of the modification of the jet shape in central Pb+Pb collisions compared with p+p
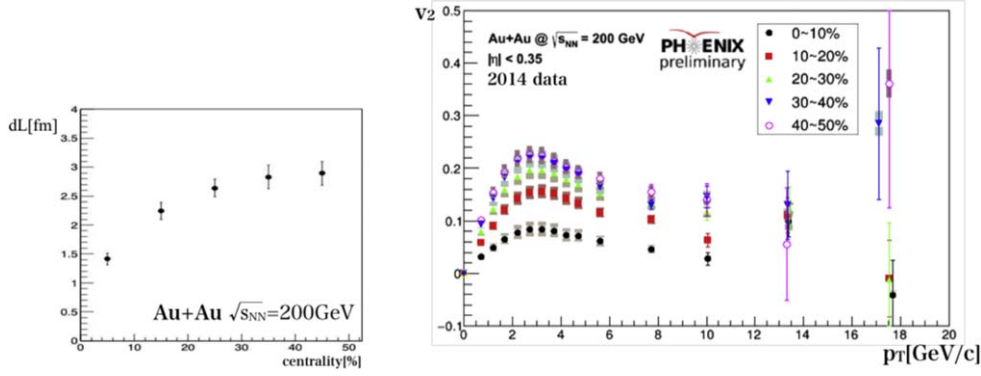
Figure 56. Left panel: Average geometric difference $dL = L_{in} - L_{out}$ between the path-length for a parton emitted perpendicular to the reaction plane (L_{out}) compared with a parton emitted along the reaction plane (L_{in}) in a Au+Au collision as a function of centrality. Right panel: Azimuthal anisotropy coefficient $v_2(p_T)$ for charged hadrons in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV for different centralities. Reproduced from [306]. CC BY 4.0.
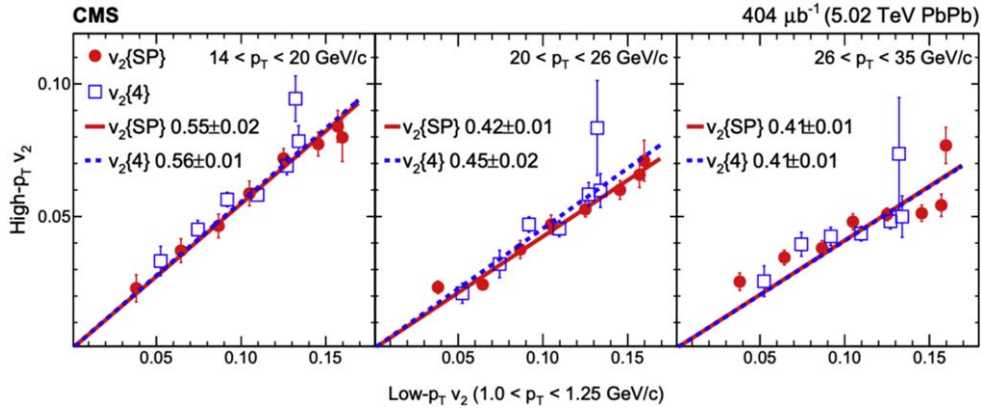
Figure 57. Correlation between the azimuthal anisotropy coefficient $v_2(p_T)$ of charged hadrons in the high-p_T region with the same coefficient measured in the low-p_T region where it is considered a measure of the collective ('elliptic') flow. The data are for Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV. The linear correlation indicates that both phenomena have the same underlying geometric origin. Reproduced from [307]. CC BY 4.0.

collisions are shown in figure 59. The data show 'softening' of the shape of the jet in terms of a redistribution of the energy in the jet to smaller z and larger angles r. This is precisely the pattern expected from in-medium gluon radiation, which involves lower parton virtuality than vacuum radiation and does not exhibit the same angular ordering that suppresses low-z, large-angle radiation.

Jets produced back-to-back with an isolated high-energy photon are preferentially initiated by a quark that has been scattered to large angles, e.g. by the Compton-like process $gq \rightarrow \gamma q$. Simulations with an event generator, such as PYTHIA, HERWIG or SHERPA, show that photon
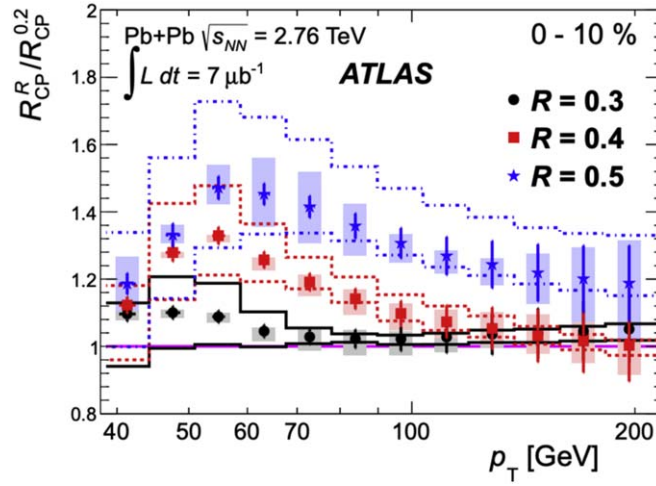
Figure 58. Relative suppression factor $R_{CP}^R(p_T)/R_{CP}^{0.2}(p_T)$ for jets in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV derived from a comparison of central and peripheral collisions. The superscript R denotes the cone angle that is used in the jet finding algorithm. As R is increased, more of the energy contained in the jet is captured. A larger value of R_{CP}^R means less suppression. The figure shows that more in-medium radiated energy is recovered for larger jet cone angles R. Reproduced with permission from [308]. CC BY-NC-ND 4.0.

tagging enhances the fraction of jets initiated by hard scattered quarks over those initiated by gluons, from 35%–50% for inclusive jets to 70%–80% for photon-tagged jets at LHC energies over the jet energy range 60 GeV $< E_T <$ 200 GeV [313]. A comparison between photon-tagged jets and inclusive jets thus allows to probe the color charge dependence of parton energy loss expressed by the dependence of the energy loss (56) on the color-SU(3) Casimir operator ($C_2 = 4/3$ for quarks and $C_2 = 3$ for gluons). Data from ATLAS shown in figure 60 confirm the expectation that jets initiated by quarks are less suppressed than those initiated by gluons, manifested in the stronger suppression of inclusive jets.

Jets generally occur in pairs (di-jets) where one jet balances the transverse momentum of the other. This means that the relative di-jet distribution is strongly peaked at 180° in azimuthal angle. While highly correlated in azimuthal angle ϕ, di-jets are not strongly correlated in pseudorapidity η but separated by a variable gap $\Delta\eta \sim \ln(x_2/x_1)$, where x_i are the momentum fractions of the colliding partons that produce the di-jet. In order to localize both partners of the di-jet, one employs two trigger particles, one for each jet, or two calorimeter-based triggers. The particle distribution in each jet is then measured relative to the (η,ϕ) coordinates of the respective trigger and denoted as [314]

$$\frac{1}{N_{\text{trig}}} \frac{\mathrm{d}^2N}{\mathrm{d}(\Delta\eta)\mathrm{d}(\Delta\phi)}. \tag{58}$$

In order to isolate di-jets in heavy ion collisions one typically also imposes a lower p_T-cutoff on the included particles and subtracts the randomized background from mixed minimum-bias events [314].

There are two main observables that have been studied for di-jets in A+A collisions. One is the additional nuclear suppression of high-p_T hadron pairs (di-hadrons) relative to the suppression of single inclusive hadrons. Such a suppression is to be expected, because both
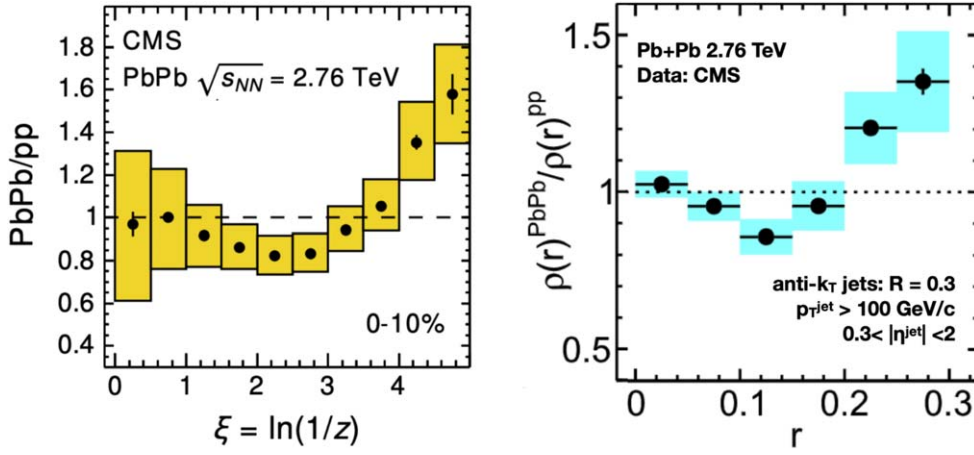
Figure 59. Modification of the jet shape in $0\%-10\%$ central Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV relative to p+p collisions for $100\,\text{GeV}/c < p_T < 300\,\text{GeV}/c$ jets with $R = 0.3$. Left panel: modification of the longitudinal jet shape $\rho(\xi)$. The hadron density in the jet is suppressed at moderate values of $z = e^{-\xi}$ and enhanced at small values, reflecting the increase in the soft components of the jet caused by additional gluon radiation in the QGP. Reproduced from [311]. CC BY 3.0. Right panel: modification of the transverse jet shape $\rho(r)$. The hadron density in the jet is shifted to larger angles $r > 0.2$, reflecting the redistribution of energy within the jet cone by gluon radiation in the QGP. Reprinted from [312], Copyright (2014), with permission from Elsevier.
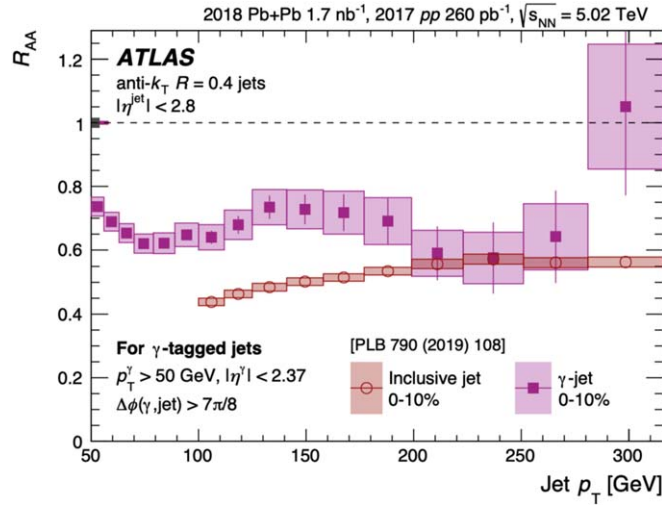
Figure 60. Comparison of the suppression factor $R_{\text{AA}}(p_T)$ of $R = 0.4$ jets in central Pb +Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV for inclusive jets and photon-tagged jets. The observation that R_{AA} is closer to unity for photon-tagged jets than for inclusive jet in the range $p_T^{\text{jet}} < 200\,\text{GeV}$ can be attributed to the fact that these contain a much smaller fraction (20%–30%) of gluon-initiated jets than inclusive jets (50%–65%). Reproduced from [313]. CC BY 4.0.
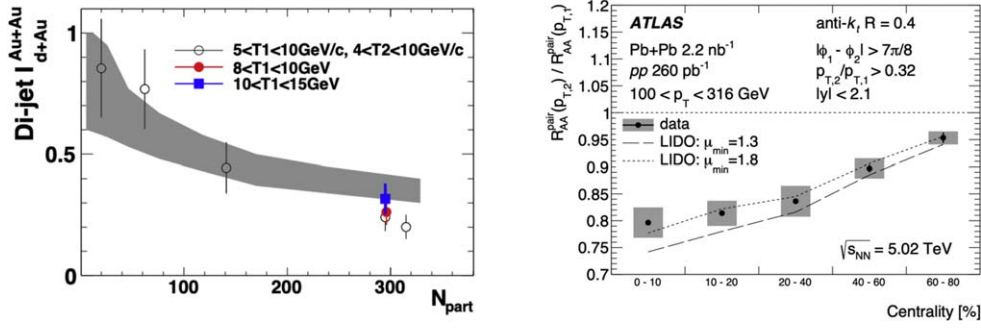
Figure 61. Left panel: Di-hadron suppression factor I_{AA} versus centrality (participant number N_{part}) for Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV measured by STAR. The different symbols indicate different trigger selections indicated by the ranges of $T1$ and $T2$. The gray band shows the expectation for di-jet surface emission. See [314] for details. Right panel: Di-jet suppression factor I_{AA}^{dijet} measured by ATLAS in $\sqrt{s_{NN}} = 5.02$ TeV Pb+Pb collisions in comparison with theoretical results obtained with the LIDO model for jet transport in a QGP medium [315]. For details see [316].

di-jet precursor partons propagate through the quark-gluon plasma and lose energy. The additional suppression for inclusive hadron pairs is expressed in terms of the quantity

$$I_{AA} = \frac{R_{AA}^{\text{di-jet triggers}}}{R_{AA}^{\text{single triggers}}}. \tag{59}$$

The left panel of figure 61 shows the I_{AA} in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV, relative to the baseline from d+Au collisions at the same energy. As one can see, the additional suppression in central collisions is comparable to the single suppression factor R_{AA} shown in figure 53. This is to be expected as the average path lengths of both scattered partons in the medium are comparable, and thus both partons suffer similar energy loss. The right panel of figure 61 shows the analogous di-jet suppression factor I_{AA}^{dijet} for the complete jets measured in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV. This quantity is to be compared with the single jet R_{AA}^{jet} shown in figure 60. The additional suppression grows with centrality, but is generally less severe than the suppression observed for single jets.

The other observable is the di-jet asymmetry A_J or, equivalently, the di-jet imbalance ratio x_J, defined as

$$a_J = \frac{p_{T,1} - p_{T,2}}{p_{T,1} + p_{T,2}}, \qquad x_J = \frac{p_{T,2}}{p_{T,1}}, \tag{60}$$

where $p_{T,1}$ and $p_{T,2}$ denote the transverse momenta of the leading and sub-leading jet, respectively. The di-jet balance ratio x_J in central Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV is shown in figure 62 in comparison with the same ratio in p+p collisions. The trigger conditions were either $p_{T,1} > 100$ GeV (left panel) or $p_{T,1} > 200$ GeV (right panel) and $p_{T,2} > 25$ GeV for jets within $|\eta| < 2.1$. For the lower trigger energy (left panel) the di-jet balance ratio distribution in Pb+Pb collisions is peaked around $x_J \approx 0.5$ and differs strongly from the distribution observed in p+p collisions.

This behavior can be interpreted as follows: when the jet is produced well outside the center of the fireball, one of the jets traverses a substantially shorter distance through the medium than the other. This causes a larger energy loss, which is reflected in a ratio x_J
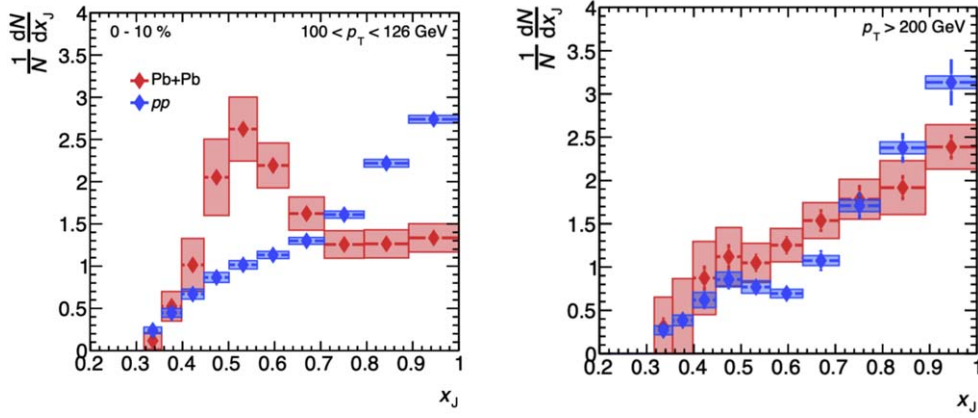
Figure 62. Di-jet balance ratio x_J in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV measured by ATLAS. Left panel: $p_{T,1} > 100$ GeV; right panel: $p_{T,1} > 200$ GeV. Reprinted from [317], Copyright (2016), with permission from Elsevier.

significantly smaller than unity. Interestingly, the difference between Pb+Pb and p+p collisions shrinks with increasing trigger threshold until the distributions are statistically indistinguishable for $p_{T,1} > 200$ GeV. The same trend is found when one goes from central to peripheral collisions [316, 317].

5.10. Heavy quark probes

Hadrons containing heavy quarks (c or b-quarks) are of interest both theoretically and experimentally. One distinguishes hadrons with open heavy flavor, such as D- and B-mesons or Λ_c baryons, and those with hidden heavy flavor, such as charmonium ($c\bar{c}$) and bottomonium ($b\bar{b}$). Open heavy flavor hadrons in their ground states decay weakly and live long enough so that their decay can be identified by micro-vertex detectors. Hadrons with hidden heavy flavor can decay either electromagnetically or by gluon-mediated strong interaction. For states with quantum number $J^{PC} = 1^{--}$, e.g. J/ψ and Υ, the strong decay channel ($J/\psi(\Upsilon) \to ggg$) is suppressed by a factor $(\pi^2 - 9)\alpha_s^3$ involving an algebraic near cancellation and a higher power of *alpha_s*. Accordingly, the decay into lepton pairs has a rather large branching ratio, which makes these states readily detectable. An extensive survey of heavy quark physics in relativistic heavy ion collisions (as of 2015) can be found in [318, 319].

On the theoretical side, heavy quarks are interesting because they are almost exclusively produced during the initial stage of the reaction by hard QCD processes, mainly $g + g \to Q + \bar{Q}$. They may or may not subsequently thermalize, which is an interesting question by itself, but their number remains essentially conserved from then until hadronic freeze-out. Inside the QGP, when light quarks are deconfined, hadrons containing both heavy and light quarks cannot exist. During hadronization, such hadrons are created by recombination of deconfined heavy quarks with light quarks or antiquarks. However, hadrons containing solely heavy quarks may survive under conditions not too far above the deconfinement threshold because their binding radii are small and their binding energies are large compared with the temperature.

The question, above which temperature $T_m^{(H)}$ a specific heavy heavy quark bound state H 'melts', has been studied in great detail using lattice QCD. Initial investigations focused on static color screening studies [320] but more recently the focus has shifted to the investigation

of dynamic properties encoded in the spectral functions [321, 322], which include non-static effects, such as ionization by thermal gluons. More generally, the heavy quark mass provides for a large momentum scale $M_Q \gg \Lambda_{\mathrm{QCD}}$ that enables various effective field theory approaches to QCD, known as heavy quark effective theory (HQET) or nonrelativistic QCD (NRQCD). In combination with HTL perturbation theory techniques, these approaches form a rigorous theoretical framework for the calculation of transport properties of heavy quarks in the QGP, including the formation and destruction of quarkonia [323].

Originally, the inclusive study of the transport properties of heavy quarks in the QGP mainly relied on the measurement of leptons (e, μ) emitted in their semi-leptonic weak decays [324]. The discrimination between leptons from b-decays and those of c-decays further requires the identification of the decay vertex where one uses, on a statistical basis, the property that hadrons containing single b-quarks (open beauty hadrons) have a longer average lifetime than those containing single c-quarks (open charm hadrons). The comparison of the inclusive lepton spectrum measured in A+A collision with the binary collision-scaled spectrum measured in p + p collisions yields information about the transport of heavy quarks in the QGP. This information is usually presented as nuclear modification factor R_{AA} plotted as a function of p_T. In order to relate to nuclear modification of the heavy quark spectrum, the lepton spectrum requires unfolding with the decay spectrum of the parent hadrons in their rest frame, which is an ill-defined procedure. One therefore usually compares the data with calculations that include the weak decays of open heavy flavor hadrons.

Figure 63 shows that lepton spectra from heavy flavor decays exhibit similar nuclear modification features as those of light hadrons. The left panel, which shows flavor separated R_{AA} for leptons from c versus b decays in Au+Au collisions at RHIC [325, 327], provides evidence that c-quarks experience strong rescattering in the QGP, resulting in a suppression which is comparable to that of light quarks. The suppression effect for leptons from b-quark decays is significantly smaller. This is expected, as the energy loss of a b-quark in collisions with thermal partons is reduced by a factor m_c/m_b when compared with that of c-quarks. In addition, radiative energy loss by heavy quarks exhibits a dead-cone effect [161, 328] that increases with quark mass. Although medium-induced gluon radiation is predicted to partially fill the radiation dead-cone, a mass dependent reduction of the radiative energy loss is predicted [329].

The right panel of figure 63 shows the muon R_{AA} for unseparated b- and c-decays in central $\sqrt{s_{\mathrm{NN}}} = 2.76$, 5.02 TeV Pb+Pb collisions at the LHC [326]. Again, the nuclear modification exhibits similar features as that measured for light charged hadrons with a minimum around $p_T \sim 10\,\mathrm{GeV}/c$. A similar plot for $\sqrt{s_{\mathrm{NN}}} = 5.02$ TeV [330] shows that the observed trend is generally well reproduced by theoretical calculations that include collisional and radiative energy loss.

With the advent of high precision tracking and vertexing of charged particles in heavy ion collisions, high statistics spectra of identified open charm hadrons have become available that do not require unfolding. Open charm hadrons that exhibit a reconstructable hadronic weak decay mode include D^0, $D^\pm$, and Λ_c. The best data are available from ALICE for identified D-mesons for Pb+Pb collisions at LHC [331]. The left panel of figure 64, showing $R_{\mathrm{AA}}(p_T)$ in central collisions, confirms that all species of D-mesons are equally suppressed, in agreement with the hypothesis that the suppression mainly reflects the energy loss of c-quarks in the QGP. The right panel of figure 64 shows that the suppression strongly depends on centrality as expected from the path-length and energy density dependence of the energy loss.

The strong suppression of D-mesons, similar to the pattern observed for light hadrons, raises the question whether c-quarks thermalize in the QGP and participate in the collective flow and if so, to what degree. A partial answer to this question is afforded by the
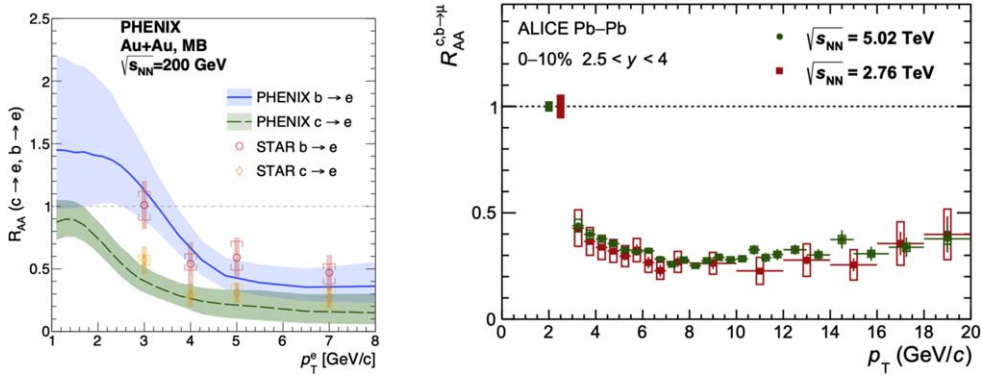
Figure 63. Nuclear suppression factor $R_{AA}(p_T)$ for single electrons from semi-leptonic heavy quark decays. Left panel: PHENIX and STAR results for vertex separated electrons from b- and c-quark decays measured in central Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV. Reproduced from [325]. CC BY 4.0 Right panel: results from ALICE for muons from semileptonic b, c-decays in $\sqrt{s_{NN}} = 2.76$, 5.02 TeV central Pb+Pb collisions. Reproduced from [326]. CC BY 4.0.
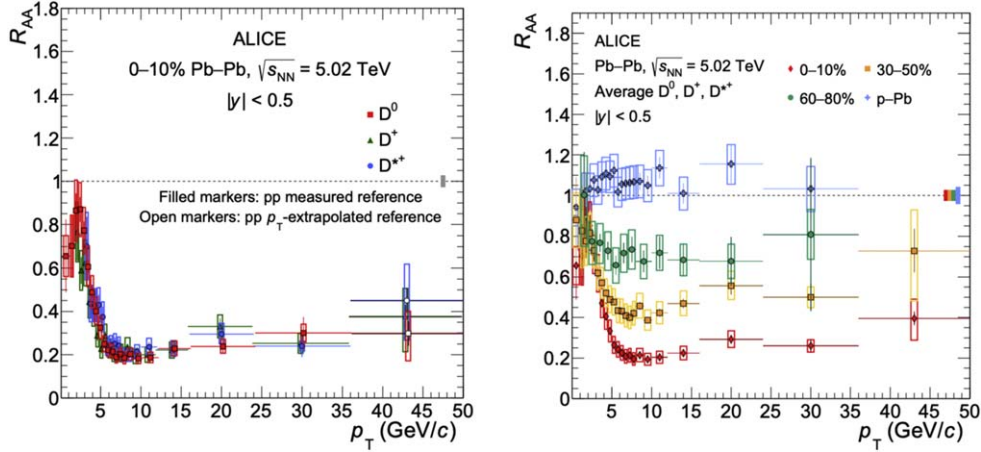
Figure 64. Nuclear suppression factor $R_{AA}(p_T)$ for identified D-mesons in Pb+Pb collisions at the LHC measured by ALICE. Left panel: $R_{AA}(p_T)$ for D^0, D^+, and D^{*+} mesons in central collisions at $\sqrt{s_{NN}} = 5.02$ TeV. Right panel: $R_{AA}(p_T)$ for all identified D-mesons in three centrality windows at $\sqrt{s_{NN}} = 5.02$ TeV Pb+Pb collisions and in p+Pb collisions at the same energy. Reproduced from [331]. CC BY 4.0.

measurement of the elliptic flow coefficient v_2 for D-mesons. Results for $v_2(p_T)$ of identified D-mesons measured by ALICE are shown in figure 65 in comparison with $v_2(p_T)$ for charged pions in two centrality windows. Except at the lowest measured p_T the D-mesons exhibit almost the same amount of elliptic flow as pions, which indicates that they participate in the overall collective flow of the QGP. The reduced v_2 at low p_T is expected because D-mesons have almost twice the mass of a proton and thus should show an even stronger kinematic reduction of v_2 at low p_T than protons.
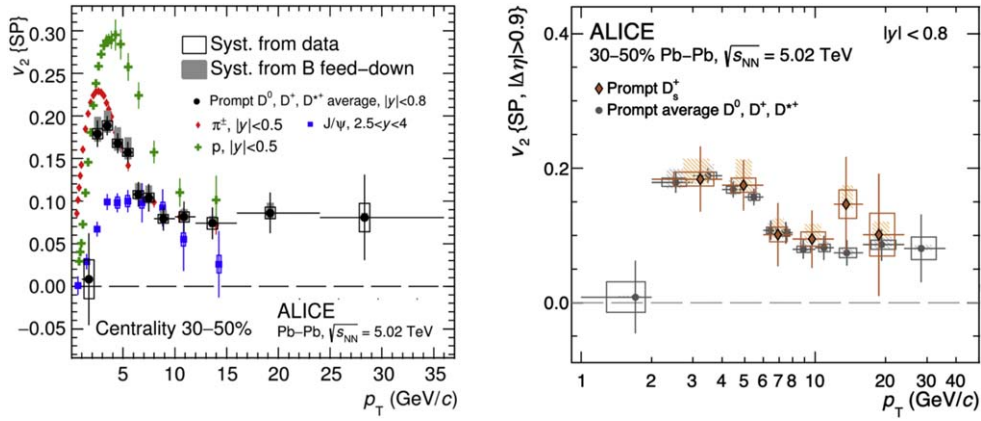
Figure 65. Elliptic flow $v_2(p_T)$ of identified D-mesons in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV measured by ALICE. The left panel shows results for a mid-centrality window (30%–50%), in comparison with the v_2 for charged pions, protons, and J/ψ From [332]. Reproduced from [332]. CC BY 4.0. The right panel shows the v_2 for D and D_s mesons in the same centrality window From [333]. Reproduced from [333]. CC BY 4.0.

Bound states of heavy quarks (charmonium, Upsilon, and B_c) are sensitive to the color screening length in the QGP [334]. When the color screening length r_D, also called the color Debye length, is shorter than the radius of the heavy quark bound state, the bound state dissolves ('melts') in the QGP and becomes a broad resonance. For each bound state there is a characteristic threshold temperature [320], also called the Mott temperature T_M [335]. In addition to color screening, the other important contribution to this process is ionization by thermal gluons in the QGP.

The two determinants of the Mott temperature are the radius of the bound state and its binding energy. In the non-relativistic limit and using natural units, these are given by $R_{Q\bar{Q}} = N^2(\alpha_s m_{Q\bar{Q}})^{-1}$ and $B_{Q\bar{Q}} = N^2 \alpha_s^2 m_{Q\bar{Q}}$, where $m_{Q\bar{Q}}$ is the reduced mass and $N \geqslant 1$ is the principal quantum number of the Coulombic bound state. For charmonium the 1s and 2s states are bound in the vacuum (J/ψ and ψ'); for bottomonium the 3s state is also bound (Υ, Υ', and Υ''.) One thus expects that when the temperature is raised above the critical temperature T_c, first the ψ' and Υ'' states melt (close to T_c), then the J/ψ and Υ' states (around $1.5 T_x$) and eventually the Υ ground-state (slightly above $2T_c$). This predicted phenomenon is known as 'sequential melting'.

Over the past three decades the suppression of quarkonium production in heavy ion collisions, compared with scaled proton–proton collisions, has been measured in great detail. The suppression is commonly expressed in terms of the ratio R_{AA}, similar to the suppression of jet production. Results for charmonium (J/ψ) suppression in Au+Au collisions at RHIC and in Pb+Pb collisions at LHC are shown in figure 66. The suppression of bottomonium ($\Upsilon(ns)$) production in Pb+Pb collisions at LHC is shown in figure 67.

The interpretation of the R_{AA} data is complicated by several effects. The two most important ones are:

- The primary production process for heavy quark pairs, $gg \rightarrow Q\bar{Q}$, is suppressed in nuclear collisions because the nuclear gluon distribution at small Bjorken-x is screened
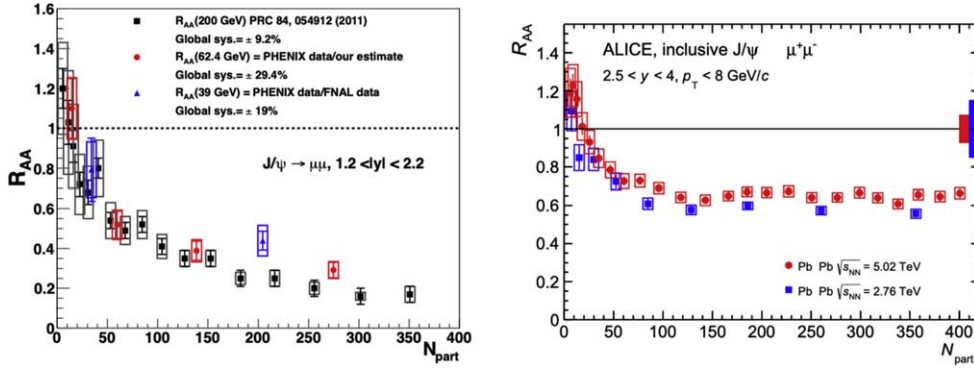
Figure 66. Nuclear suppression factor R_{AA} for J/ψ production as a function of centrality, measured by the number of participant nucleons N_{part}. Left panel: Au+Au collisions at $\sqrt{s_{NN}} = 200, 62.4, 39$ GeV from [336]. Reprinted (figure) with permission from [336], Copyright (2012) by the American Physical Society. Right panel: Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV from [337]. Reproduced from [337]. CC BY 4.0.
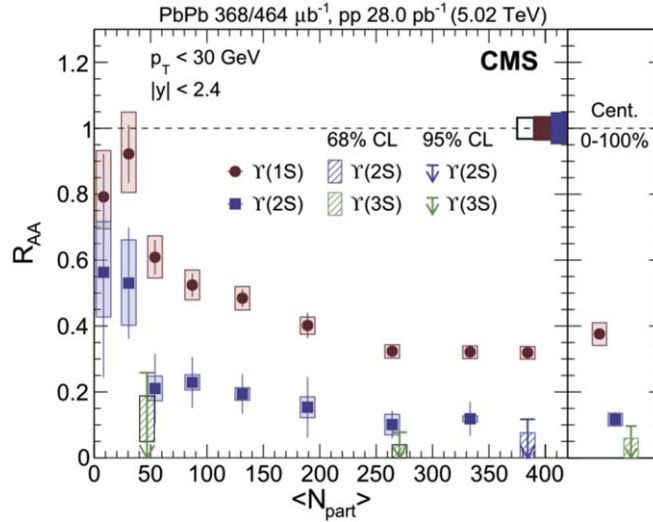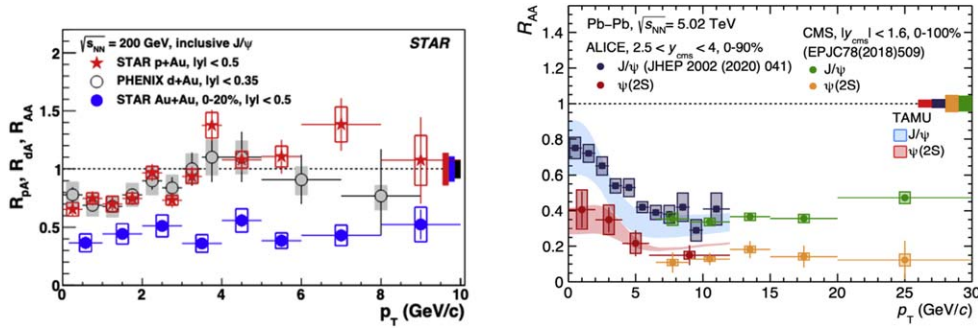

Figure 67. Nuclear suppression factor R_{AA} for bottomonium production in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV as a function of centrality, measured by the number of participant nucleons N_{part}. Reproduced from [338]. CC BY 4.0.

('shadowed'). This effect, which can be studied in $p(d) + A$ collisions, is mainly observed at low transverse momentum p_T, as is visible in figure 68 (left panel).

- In Pb+Pb collisions at LHC energies, $c\bar{c}$ pairs are copiously produced and are thought to be thermalized in the QGP. These pairs can coalesce into J/ψ and ψ' mesons when the QGP hadronizes. This regeneration mechanism leads to a striking difference in the centrality dependence of J/ψ suppression at LHC (right panel of figure 66), which is constant over a wide centrality range, compared with the centrality dependence at RHIC (left panel of figure 66), which shows a continued drop toward central collisions. The

Figure 68. Left panel: Nuclear suppression factor $R_{(p/d)A}(p_T)$ for J/ψ production in p
+Au and d+Au collisions at $\sqrt{s_{NN}} = 200$ GeV from [339]. Reproduced from [339].
CC BY 4.0. The approximately 30% suppression at small p_T is mainly due to gluon
shadowing in the Au nucleus. Right panel: nuclear suppression factor $R_{AA}(p_T)$ for J/ψ
and ψ' production in Pb+Pb $\sqrt{s_{NN}} = 5.02$ TeV from [340]. Reproduced from [340].
CC BY 4.0. The enhancement at small p_T is attributed to recombination of c and $\bar{c}$
quarks during hadronization (regeneration).

different behavior is caused by a relative enhancement of charmonium production at low
p_T at LHC, which is shown in the right panel of figure 68. The same mechanism, albeit
less pronounced, may already be at work in Au+Au collisions at RHIC at central rapidity
($|y| < 0.35$) where J/ψ production is found to be less suppressed than at forward rapidity
($|y| > 1.2$). The stronger suppression of ψ' compared with J/ψ seen in this figure is also
evidence for the sequential melting concept.

5.11. Electromagnetic probes

Electromagnetic probes are theoretically interesting, because they do not interact strongly and
therefore are emitted without final-state effects from all stages of the reaction. Many hadrons
and other particles are also produced during the whole evolution of the heavy ion reaction, but
the main advantage of electromagnetic probes is that they reach the detector undisturbed. A
dilepton pair or a photon escapes even the hot and dense reaction zone without further
interaction. However, their relatively weak interaction makes them experimentally challen-
ging, because their yields are small, and the experimental signal often has large backgrounds
from weak or electromagnetic decays of abundant hadrons that create photons or leptons in
the final state (e.g. $\pi^0 \to \gamma\gamma$, $D^0 \to K^-\mu^+\nu_\mu$). Any interpretation of experimental results relies
therefore on theoretical input on the origin of the contributions from the different sources and
stages of the reaction.

There are two main physics questions that can be addressed with dilepton and photon
measurements. By investigating the invariant mass spectrum of dileptons emitted from vector
mesons one can get insight about the properties of the spectral functions of resonances in the
medium (see [346] for a recent review). The ρ-meson is of special interest, since the idea is to
study chiral symmetry restoration by observing the spectral functions of the ρ and its chiral
partner a_1 become degenerate. More recently, the focus has shifted to signatures of chiral
mixing, since the measurement of the a_1 spectral function is out of reach. The second main
topic is the idea of extracting a temperature of the quark-gluon plasma from the thermal
radiation. For this purpose, photons might seem more straightforward, but one has to account
for a blue shift in the photon spectrum due to radial flow. The slope of dilepton spectra in the
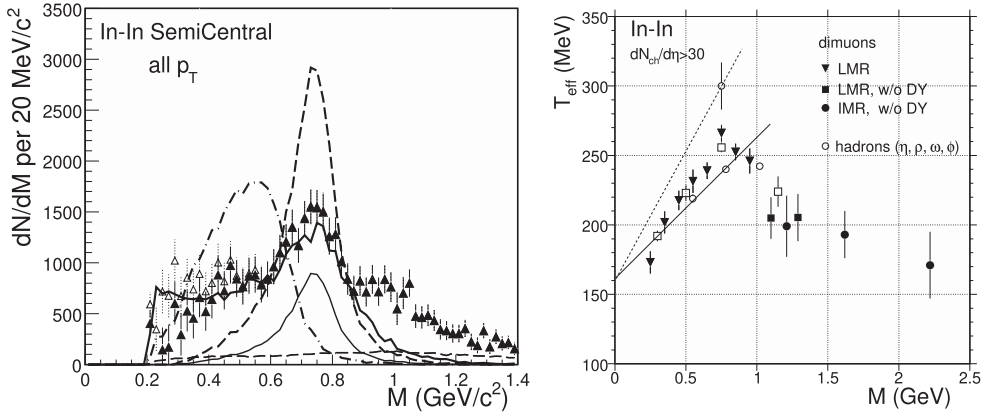
Figure 69. Left: Comparison of the excess e^+e^- mass spectrum for 158 A GeV/c In+In collisions at $dN_{ch}/d\eta$=140 to the cocktail ρ (thin solid), unmodified ρ (dashed), in-medium broadening ρ [341, 342] (thick solid), in-medium moving ρ related to [343] (dashed-dotted). Reprinted (figure) with permission from [344], Copyright (2006) by the American Physical Society. Right: inverse slope parameter $T_{\rm eff}$ versus dimuon mass M for $dN_{ch}/d\eta > 30$ with open charm subtraction. Reprinted (figure) with permission from [345], Copyright (2008) by the American Physical Society.

invariant mass region between the ϕ-meson and the J/ψ, $1.1 \text{GeV} < m_{\rm inv} < 3 \text{ GeV}$ provides a more direct measure of the thermal radiation from the plasma, because $m_{\rm inv}$ is not affected by flow.

Figure 69 depicts the most precise dilepton measurement in a heavy ion environment to date. The invariant mass spectrum of dimuons was recorded for In–In collisions at $E_{\rm lab} = 158A$ GeV at the CERN-SPS. After contribution from the 'cocktail' of known hadronic sources has been subtracted the remaining excess yield (see the left panel of figure 69) provides a measure of the spectral function of the ρ-meson inside the hot and dense medium. From these results, it can be concluded that the ρ-meson is strongly modified in the medium and mainly broadened, while a mass shift of the pole position has not been observed. The differential dimoun measurements make it possible to fit transverse momentum spectra in different invariant mass bins and extract the effective temperature as the inverse slope shown in the right panel of figure 69. In the low mass region (LMR) the extracted values agree with the ones from the hadronic spectra, consistent with a hadronic origin, while above 1 GeV in invariant mass the thermal slope saturates and suggests emission from an equilibrated quark-gluon plasma. In this mass region, there is a correlated background from heavy quark decays that needs to be carefully subtracted.

Moving to higher beam energies the STAR collaboration has measured dielectron invariant mass spectra as shown in figure 70. The excess yield above the hadronic cocktail emission indicates that the spectral function of the ρ-meson is also broadened at these higher beam energies. Theoretical calculations include thermal dilepton rates from effective field theory folded with a fireball model [348, 349], coarse-grained UrQMD transport calculations with the same thermal dilepton rates [350, 351] and fully microscopic non-equilibrium calculations by within the PHSD approach [145, 352]. In the future, ALICE is expected to be also able to measure precise, background subtracted dilepton spectra in Pb+Pb collisions at LHC energies.
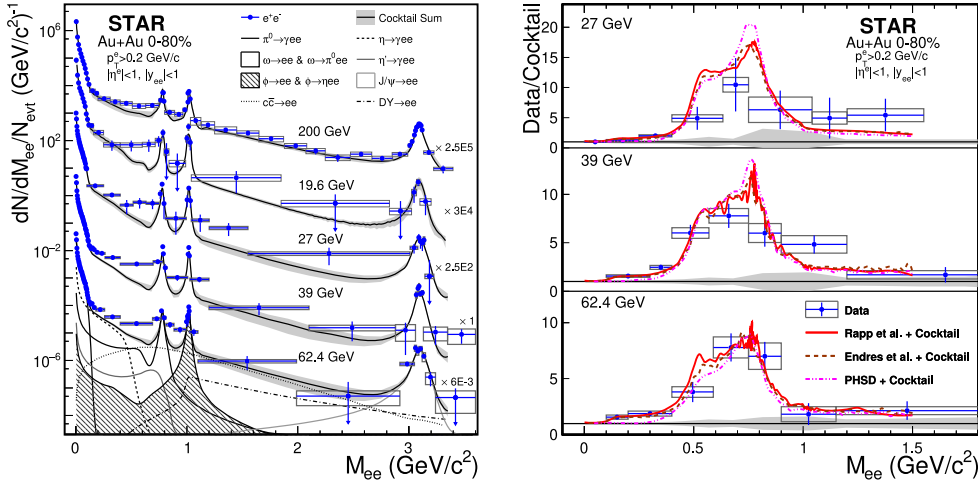
Figure 70. Left: Dielectron invariant mass spectra within the STAR acceptance from $\sqrt{s_{NN}}$=19.6, 27, 39, 62.4, and 200 GeV 0%−80% most-central Au+Au collisions (scaled for visibility). Right: the ratio of the invariant mass spectra to the cocktail with the ω and ϕ yields removed compared to model calculations. Reprinted (figure) with permission from [347], Copyright (2023) by the American Physical Society.

At GSI, the HADES experiment is dedicated to investigating dilepton emission from elementary and heavy ion reactions. In this baryon-dominated environment the ρ-meson is mainly modified due to its interactions with the baryonic resonances. Figure 71 shows the extracted thermal emission from Au+Au collisions at $E_{lab} = 1.23A$ GeV. There are clear indications of medium modifications in comparison with the vacuum environment indicated by the agreement with coarse-grained (CG) transport calculations including medium-modified spectral functions for vector mesons. The calculation (see right panel of figure 71) shows this difference explicitly for the radiation from the ρ- and ω-mesons, the main vector mesons contributing in this mass range. In the future, HADES and the CBM experiment will measure the excitation function of thermal dilepton emission with the goal to identify signatures of a first-order phase transition.

Photon production from heavy ion collisions is dominated by sources from hadronic decays, most prominently the π^0 that decays to photons with an almost 100% branching ratio. Therefore, it is very challenging to experimentally extract the primary photons, usually called 'direct photons'. The first measurement of a direct photon spectrum was accomplished by the WA98 collaboration at SPS [364]. More recently, PHENIX and ALICE have published transverse momentum spectra and flow measurements for direct photons. A compilation of the spectra is shown in figure 72. In order to compare spectra measured in different centrality windows they have been divided by a normalization factor $(dN_{ch}/d\eta)^{1.25}$. This factor is motivated by the fact that the charge particle multiplicity $(dN_{ch}/d\eta)$ scales with the number of participant nucleons, N_{part}, whereas the direct photon yield is expected to scale with the number of binary nucleon–nucleon collisions $N_{coll} \sim N_{part}^{4/3}$. At higher transverse momenta the spectra match the expectations from perturbative QCD calculations [365], while at lower momenta an exponential behavior is observed. By fitting the slope of the transverse momentum spectra of direct photons, one can infer an effective temperature of the quark-gluon plasma as depicted in figure 73.
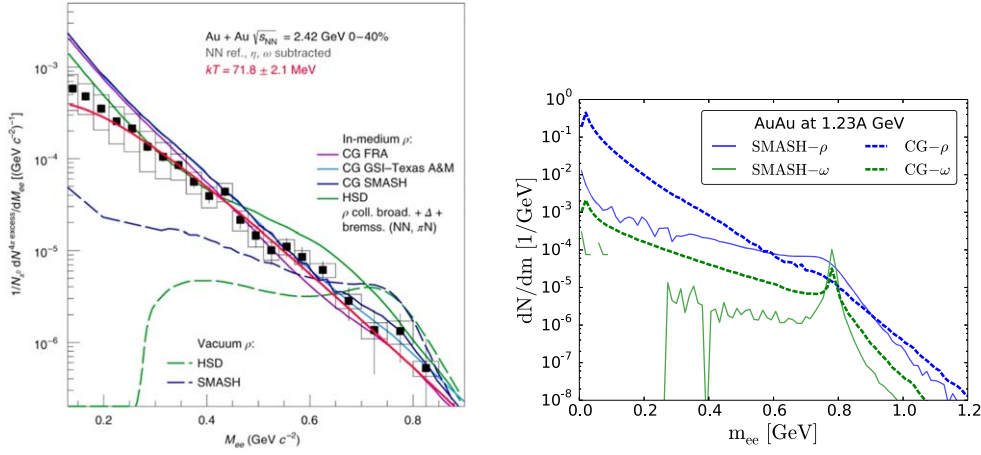
Figure 71. Left: Excess yield of dileptons extracted by subtracting η, ω contributions as well as the *NN* reference normalized to the number of neutral pions, the red curve shows a thermal $dN/dM_{ee} \propto (M_{ee})^{3/2} \exp(-M_{ee}/T)$ fit. Reproduced from [353], with permission from Springer Nature. Right: comparison of invariant mass spectra of dielectrons produced by ρ and ω in Au+Au collisions at $E_{kin} = 1.23\,A$ GeV within the coarse-graining approach versus the default SMASH dilepton production. Reprinted (figure) with permission from [354], Copyright (2018) by the American Physical Society.

In figure 74 state-of-the-art calculations considering direct photon emission from all stages of the reaction are compared with spectra and elliptic flow data from the ALICE collaboration in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. Again the prompt photons from the early hard collisions dominate the high p_T region. At lower transverse photon momenta the thermal emission from the hydrodynamic medium dominates. The photons emitted during the pre-equilibrium stage are depicted by the full line. The magnitude of photon elliptic flow it depends on the time at which full chemical equilibrium between quarks and gluons is achieved (indicated by the different values of τ_{chem}).

In general, it is still hard to explain the photon production yield and elliptic flow at the same time consistently in one theoretical calculation. The intuitive reason is as follows: processes that increase the yield need to occur early in the evolution when the plasma is hotter, while larger elliptic flow is reached in the later stages of the evolution. Therefore, elliptic flow is increased, if later sources are enhanced, for example in [369] the contribution of the non-equilibrium hadronic stage has been shown to be significant for low transverse momenta. In the future, it is going to be crucial to investigate emission of electromagnetic probes from well-calibrated models for the bulk evolution [368].

5.12. Production of light nuclei and exotic hadrons

There are two different mechanism that lead to the production of light nuclei in heavy ion reactions. In the fragmentation regions, at very forward and backward rapidity, the spectator remnants can fragment and reach the detector as a multitude of smaller and larger nuclei. This can happen because the spectator remnants emerge from the collision in a highly excited state, and collisions among nucleons within these spectator remnants often result in their disintegration. We will not discuss this phenomenon further as the main interest here is on the production of light (anti-)nuclei at midrapidity within the hot and dense collision region. The
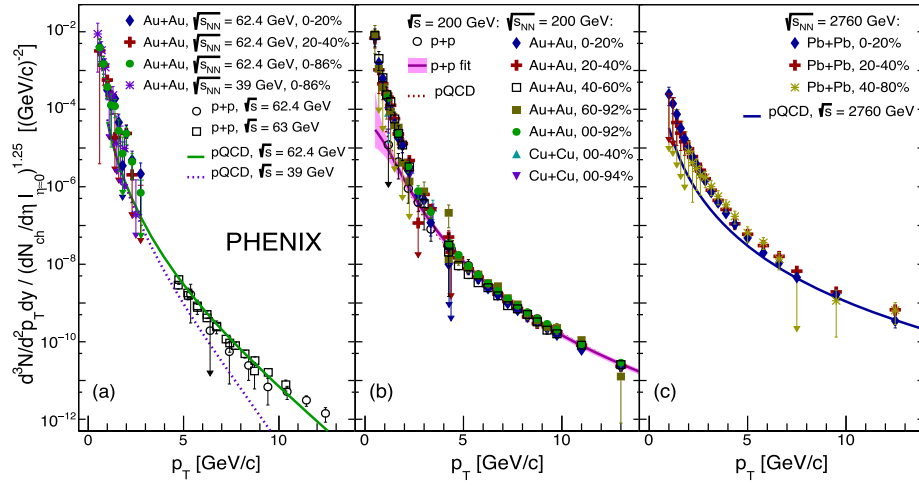
Figure 72. Direct photon p_T-spectra normalized by $(\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}\eta)^{1.25}$ for (a) the minimum bias Au+Au 39 and 62.4 GeV data sets, (b) various centrality selected 200 GeV Au +Au [355–357] and Cu+Cu [358] data sets, and (c) various centrality selected Pb+Pb 2.76 TeV data sets [359]. Also shown are (a) p+p data from the ISR [360, 361] and (b) p+p 200 GeV data [362]. Reprinted (figure) with permission from [363], Copyright (2022) by the American Physical Society.
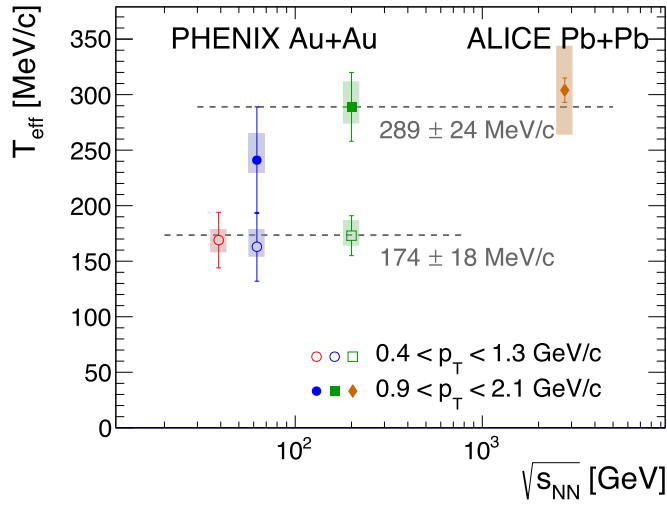
Figure 73. Inverse slopes, T_{eff}, obtained from fitting the combined data from central collisions is compared to the fit results of the individual data sets at 62.4, 200, and 2760 GeV. Also included is the value for $\sqrt{s_{\mathrm{NN}}} = 39$ GeV obtained from fitting the minimum bias data set in the lower-$p_\perp$ range. Reprinted (figure) with permission from [363], Copyright (2022) by the American Physical Society.
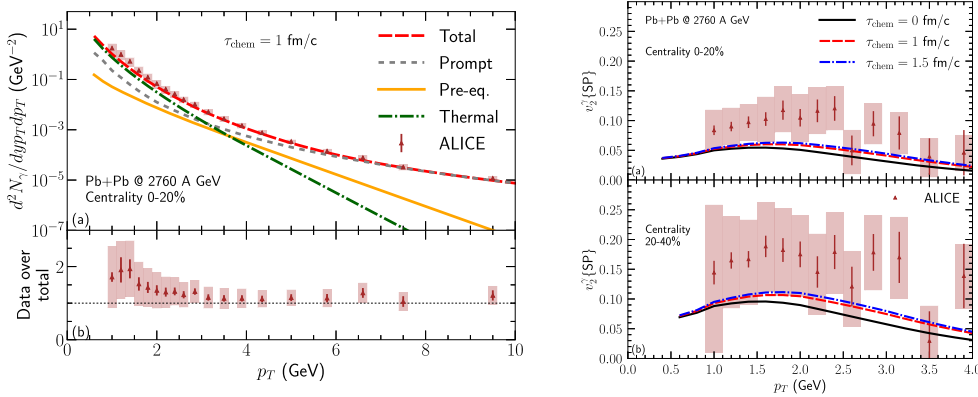
**Figure 74.** Left: (a) Direct photon yield in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, in the 0%−20% centrality class with depicting the different components compared to data from the ALICE Collaboration [366]. Right: direct photon elliptic flow $v_2^{\gamma}(p_T)$ in $\sqrt{s_{NN}} = 2.76$ TeV Pb+Pb collisions for different chemical equilibration times compared to experimental data from the ALICE collaboration [367]. Reprinted (figure) with permission from [368], Copyright (2022) by the American Physical Society.

production of nuclei is interesting because of its dependence on the nucleon–nucleon interaction, but also as a probe of possible differences in the properties of matter and antimatter.

Figure 18 shows the production yields of many particle species including light nuclei in Pb +Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV [182]. Due to the vanishing chemical potential at this high energy the yields of deuterons and anti-deuterons, helium-3, hypertriton and helium-4 and their anti-nuclei are pairwise identical. Generally, the predictions within the statistical hadronization model agree very well with the measurements for the same temperature as for all other hadron species. This poses immediately a question of current debate: how can particles that have small binding energies of a few MeV freeze-out chemically from a fireball of hot and dense strongly-interacting matter at a temperature that is many times higher ($T_{ch} \sim 150$ MeV)?

Besides the thermal production of light nuclei, another production mechanism is proposed. The idea is that only individual hadrons are produced from the fireball, but later on the nucleons or antinucleons combine to form light nuclei. This coalescence picture [371, 372] involves calculating the overlap in phase-space of all nucleons and drawing conclusions about the abundance of light nuclei from there. In such a picture it is expected that the transverse momentum spectra of light nuclei and their anisotropic flow coefficients v_n scale according to the number of nucleons contained in a nucleus, in analogy to the recombination approach and the partonic quark number scaling discussed above (see section 5.8). Figure 75 depicts the expected yields at midrapidity in a thermal and a coalescence production approach. The rapid fall-off of the yields with beam energy is caused by reduced baryon stopping at higher energies, which means that fewer valence quarks that can form complex nuclei are present in the fireball. The highest yields are therefore expected in the beam energy range $E_{lab} = 10 - 20A$ GeV.

Due to recent increased interest in the topic, several models have been developed that aim to explain the mechanism behind this behavior. One model attempts to describe the early chemical freeze-out via the Saha equation in analogy to cosmology [373] or rate equations [374]. Another approach involves microscopic calculations of the non-equilibrium dynamics
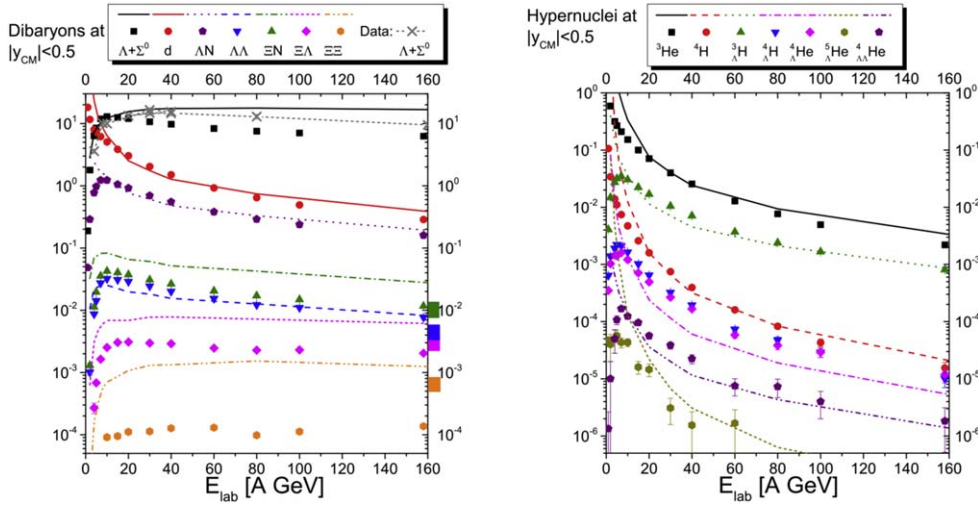
Figure 75. Excitation function of dibaryons (left) and hypernuclei (right) produced in central Au+Au collisions calculated within a thermal production model from the UrQMD hybrid approach (full lines) compared to a coalescence approach (symbols). Reprinted from [370], Copyright (2012), with permission from Elsevier.

of light nuclei in the hadronic stage of the reaction [375]. At low beam energies in a baryon-rich environment, the main reactions are the nucleon catalysis reactions, while at high beam energies the pion catalysis is more important. Calculations within a hadronic transport approach suggest that the chemical equilibrium is maintained due to the high reaction cross sections during the rescattering phase in nuclear collisions [376, 377]. The kinetic and chemical decoupling almost coincide in such an approach.

Hypernuclei and their properties are of particular interest since one might be able to infer knowledge about the $\Lambda - N$ interaction. For a while, the hypertriton ($^3_\Lambda$H 'lifetime puzzle' attracted much interest, since the first measurements of the lifetime of the hypertriton seemed to deviate from the lifetime of a free Λ hyperon. However, newer measurements by ALICE (see figure 76), STAR, and HADES, and an even more precise result recently reported by ALICE [379], point to an excellent agreement with the world data for the Λ lifetime.

A rather recent idea that has already been applied with considerable success, is to measure the femtoscopic correlations of exotic hadrons in elementary and heavy ion reactions. These final-state correlation measurements can be connected to the hadronic interactions of those particles. In this manner it is possible to extract quantitative information about the interaction of Ξ, Ω and other exotic hadrons containing one or more strange quarks with other hadrons (see the recent reviews [380, 381]).

5.13. Vorticity and polarization

The initial state of a non-central heavy ion collision is characterized by a very large angular momentum in the center-of-mass frame. For example, a collision between two ^{208}Pb nuclei at $\sqrt{s_{\mathrm{NN}}} = 5.02$ TeV with impact parameter $b = 7.5$ fm commands an angular momentum $L = A\sqrt{s_{\mathrm{NN}}}b/2 \approx 2 \times 10^7 \hbar$. Only a tiny fraction of this angular momentum ends up in the central rapidity region where most observed particles are produced. However, even this small fraction endows the QGP in non-central collisions with a sizable vorticity.
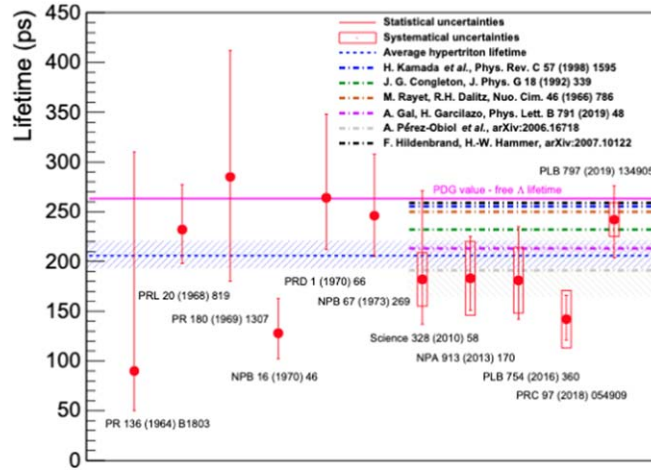
Figure 76. Comparison of measurements of the hypertriton lifetime in heavy ion reactions to theoretical calculations and the lifetime of the free Λ hyperon. From [378].

Assuming thermal equilibrium, the degree of polarization, expressed as the probability for a certain spin orientation $\hat{S}$, of Λ-hyperons within the final-state hadronic gas is [382]:

$$P_\Lambda(\hat{S}) \propto \exp\left(\frac{(\vec{\omega} + \mu_\Lambda \vec{B}) \cdot \hat{S}}{T}\right), \tag{61}$$

where $\vec{\omega}$ is the vorticity vector of the matter, $\mu_\Lambda = (-0.6138 \pm 0.0047)\mu_N$ is the Λ magnetic moment in nuclear magnetons μ_N, and $\vec{B}$ is the magnetic field present at emission. Since the magnetic moments of the Λ and $\overline{\Lambda}$ differ by their sign, a magnetic field would cause them to be oppositely polarized, whereas vorticity of the medium results in identical polarizations. Λ-hyperons are ideal probes of polarization, because the direction of their parity violating weak decay $\Lambda \to p + \pi^-$ is strongly aligned with the direction of the spin of the hyperon. The angular distribution of the decay proton momentum $\vec{p} = p\hat{n}$ in the Λ rest frame is given by [382]

$$\frac{\mathrm{d}w}{\mathrm{d}\Omega} = \frac{1 + \alpha_\Lambda \vec{P}_\Lambda \cdot \hat{n}}{4\pi} \tag{62}$$

with $\alpha_\Lambda = 0.732 \pm 0.014$ [383].

The global polarization of Λ and $\overline{\Lambda}$ have been measured in Au+Au collisions at RHIC over a wide energy range [384, 385] and in Pb+Pb collisions at LHC [386]. The polarization is found to be along the direction of the angular momentum in the collision, perpendicular to the reaction plane, and to grow with decreasing collision energy, presumably because a larger fraction of the angular momentum carried by the incident nucleons ends up at midrapidity. A recent compilation of data is shown in figure 77. The measured average polarizations can be converted into an estimate of the vorticity. From (61) one finds

$$\langle \omega \rangle = (\langle P_\Lambda \rangle + \langle P_{\overline{\Lambda}} \rangle)T, \tag{63}$$

which gives $\langle \omega \rangle \approx 10^{22}\,\mathrm{s}^{-1}$ for $\langle P_\Lambda \rangle \approx \langle P_{\overline{\Lambda}} \rangle \approx 0.02$. STAR has also measured the global polarizations of Ξ and Ω hyperons at $\sqrt{s_{NN}} = 200\,\mathrm{GeV}$, which are found to be of similar magnitude as the Λ-polarization supporting the interpretation of the phenomenon as a universal effect of QGP vorticity [388].
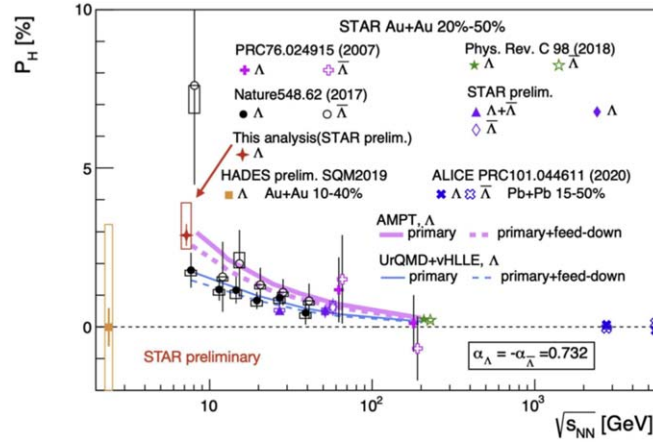
Figure 77. Global Λ and $\overline{\Lambda}$ polarization with respect to the collision plane in semicentral Au+Au (Pb+Pb) collisions as a function of the collision energy. Reproduced from [387]. CC BY 4.0.

The average polarizations $\langle P_\Lambda \rangle$ and $\langle P_{\overline{\Lambda}} \rangle$ agree with each other at all collision energies within the experimental errors. The most precise values have been measured in Au+Au at $\sqrt{s_{NN}} = 200$ GeV, where $\langle P_\Lambda \rangle - \langle P_{\overline{\Lambda}} \rangle = 0.037 \pm 0.07$ [389]. This measurement allows to set an upper limit on the magnetic field at the emission time: $|B| < 10^{12}$ T [390]. This is less than 10^{-3} of the maximal magnetic field generated during the collision of the two nuclei.

In addition to the global polarization of Λ and $\overline{\Lambda}$, the experiments have also observed local polarization of hyperons along the beam direction. The orientation of the polarization vector depends on the direction of the hyperon transverse momentum and shows a quadrupole pattern with respect to the beam axis [391]. This effect is now understood as a result of the shear caused by the anisotropy of the transverse flow (see [392] for a review).

Vector meson alignment is another phenomenon related to spin that has been experimentally observed. Both STAR [393] and ALICE [394] have found evidence for a global alignment of the spins of K^{*0} and ϕ mesons with respect to the collision plane. Alignment is defined as the deviation of the $m = m' = 0$ component of the spin–1 density matrix $\rho_{m'm} = \langle m'|\rho|m \rangle$ from its equilibrium value $\rho_{00} = 1/3$ when all spin orientations are equally likely. Different from the concept of polarization, which distinguishes between spin orientation up and down with respect to the collision plane, alignment makes no such distinction. The mechanisms that can cause a nonzero aligment are thus less constrained by symmetry than those that can cause polarization. Indeed, the experiments find much larger values $|\rho_{00} - 1/3| \approx 0.1 - 0.2$ (compared with 10^{-2} for global polarization), but there is currently no generally accepted explanation for the origin of this effect.

Particle spin can be added to the hydrodynamical model of quark-gluon plasma expansion by introducing spin degrees of freedom to the fluid [395]. The equations for such a spinning, viscous fluid can be derived from kinetic theory in the usual way by applying a reduction to a limited number of moments of the momentum space distribution [140, 396] or by symmetry-based analysis of the allowed perturbations of the energy-momentum tensor and the spin current around equilibrium [397]. In such an approach the observed collision energy dependence of the global Λ polarization can be explained with reasonable assumptions about the initial conditions (see figure 78).
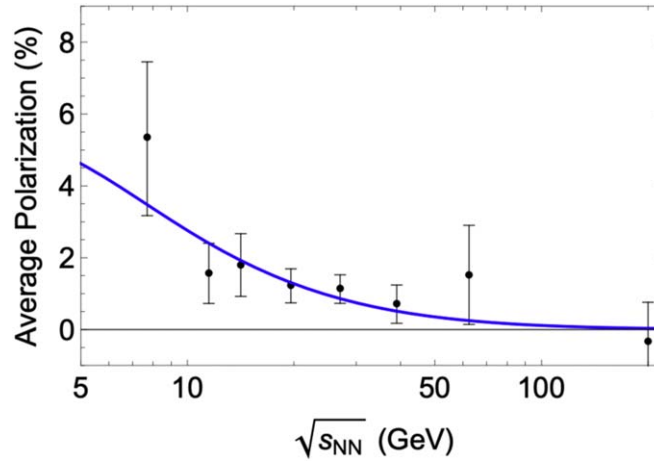
Figure 78. Comparison of the collision energy dependence of the average hyperon polarization predicted by a hydrodynamical model with spin degrees of freedom [397] (solid blue curve) with the STAR data shown in figure 77. Reproduced from [397]. CC BY 4.0.

5.14. Chiral magnetic effect

QCD gauge fields are characterized by a topological quantum number, called winding number. Field configurations with different winding number are separated by an energy barrier. In the vacuum, transitions between field configurations with different winding number are possible by tunneling processes, which can be described semiclassically. The gauge field fluctuations involved in the tunneling process are nonperturbative (i.e. they cannot be expanded in a power series of the gauge coupling g) and commonly known as instantons.

While field configurations with a definite winding number break CP invariance, experimental data tell us that the QCD vacuum realized in nature contains a superposition of such field configurations that conserves the global CP symmetry of QCD for yet unknown reasons. (An axion field, if it exists, would explain this mystery in a natural way through the so-called Peccei–Quinn mechanism [398]). In the presence of electromagnetic fields, however, these winding number fluctuations can generate local CP violations via the chiral anomaly [399] of the axial current

$$\partial_\mu j_5^\mu = -\frac{N_c}{2\pi^2} \sum_f Q_f^2 \vec{E} \cdot \vec{B}, \tag{64}$$

where Q_f are the electric charges of the light quark flavors. Anomalous hydrodynamics [400, 401] adds (64) to the conservation laws for energy-momentum, baryon number, and electric charge as a fourth macroscopic equation. As usual, the conservation laws must be supportlemented with constitutive equations, in this case for the vector and axial vector currents:

$$j^\mu = nu^\mu + \sigma E^\mu + \sigma_B B^\mu$$
$$j_5^\mu = n_5 u^\mu + \xi_E E^\mu + \xi_B B^\mu, \tag{65}$$
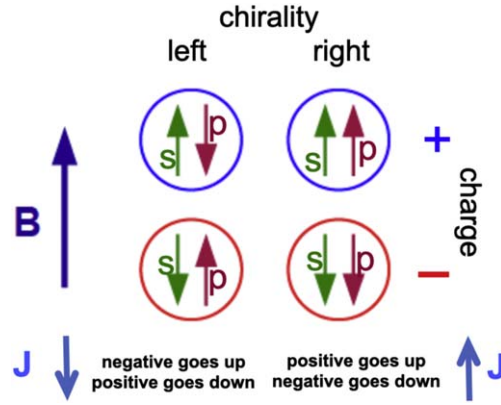
Figure 79. Schematic illustration of the mechanism behind the chiral magnetic effect. The magnetic field aligns the quark spins (*s*) along the field lines according to the quark charge. Chirality associates associates a momentum direction *p* with the quark spin; positive chirality (right-handed) quark spins are aligned with the momentum, negative (left-handed) quarks spins are anti-aligned. This creates a net electric current if the QGP contains a different number of left- and right-handed quarks.

where σ, σ_B, ξ_E, ξ_B are transport coefficients (σ is the usual electric conductivity), n, n_5 denote the usual and axial quark densities, and $E^\mu = F^{\mu\nu} u_\nu$, $B^\mu = \varepsilon^{\mu\nu\alpha\beta} u_\nu F_{\alpha\beta}$ are Lorentz covariant expressions for the electric and magnetic field. It is convenient to express these transport coefficients through vector and axial vector chemical potentials, μ and μ_5 [401].

In the absence of explicit parity violation, the global axial chemical potential is $\mu_5 = 0$, which implies that the magnetic conductivity σ_B vanishes. However, the presence of instantons implies that μ_5 fluctuates locally, which means that the electric current receives locally fluctuating contributions from the magnetic field. Owing to the motion of the colliding nuclei heavy ion collisions generate very strong, short-lived magnetic fields of the order $eB \sim m_\pi^2$ that point perpendicular to the collision plane. One thus expects electric current fluctuations perpendicular to the collision plane, which result in event-by-event separation of the net electric charge of particles emitted into the upper and lower hemispheres. This is called the chiral magnetic effect (CME) [399].

The charge separation can be understood as a direct kinematic consequence of the alignment of quark spins along (or against) the magnetic field and the alignment of spin and momentum encoded in the chirality of the quark. If the magnetic field aligns a quark spin in the direction of the field, then a right-handed quark (positive chirality) will move in the direction of the field, and a left-handed quark (negative chirality) will move against the field. A positive axial density implies a preponderance of quarks with positive chirality and vice versa. Since the magnetic moment of a quark depends on the sign of its electric charge, this leads to a current in the direction of the magnetic field, if the axial density $n_5 > 0$, and against the field if $n_5 < 0$, as is illustrated in figure 79.

Since the coefficients σ_B and ξ_B in (65) are proportional to the densities n_5 and n, respectively, the anomalous hydrodynamic equations sustain a low-energy propagating mode corresponding to propagating coupled density-axial density fluctuations similar to the sound mode [402]. This excitation, which propagates with a speed proportional to the magnetic field B, is known under the name 'chiral magnetic wave' (CMW). Like the current fluctuations

induced by the CME, the CMW could be seeded by the topological charge density fluctuations in the glasma during the earliest stage of the collision.

Another phenomenon worth mentioning is the 'chiral vortical effect' (CVE). A complete analysis shows that the constitutive equations (65) also receive contributions proportional to the vorticity vector ω^μ of the QGP. Their effect is similar to the chiral magnetic effect with axial charge fluctuations driving a fluctuating electric current along the direction of the vorticity vector. It can be similarly understood as the CME as a result of the alignment of quark spins along the QGP vorticity vector in thermal equilibrium, which we discussed in the context of global hyperon polarization. Thus, the illustration in figure 79 applies, except that the vorticity ω replaces the magnetic field B. In an off-central heavy ion collision both vectors, $\vec{\omega}$ and $\vec{B}$, point in the same direction perpendicular to the reaction plane. An overview of these phenomena can be found in [403].

A number of observables are specifically sensitive to such fluctuating electric charge separation phenomena. A possible search strategy in the context of known background effects is discussed in [404]. Quantitative predictions based on solutions of the anomalous hydrodynamical equations with reasonable assumptions for the initial axial density fluctuations can be found in [405]. The magnitude of the expected event-by-event fluctuations also depends strongly on the longevity of the magnetic field. For the parameters used in [405] the experimental signals for the chiral magnetic effect are in the range of 10^{-4}, but other assumptions may lead to much smaller predicted values [406].

As theoretical predictions of the magnitude of observables for the chiral magnetic effect are beset with large uncertainties, experimental searches for it are of paramount importance. Owing to the parity conserving nature of QCD, the anomalous electric current must fluctuate event by event, thus all signals involve the measurement of event-by-event fluctuations. This means that other 'normal' sources that are sensitive to the orientation of the reaction plane can contribute, in particular, those involving charged resonance decays modulated by the elliptic flow of the final-state hadron distribution (see [404] for an in-depth discussion and references). Experimental studies at RHIC [407] and LHC [408] have concluded that at most a small fraction (less than 10% for RHIC) of the observed signals can be attributed to the CME.

The comparison of measurements of observables sensitive to the chiral magnetic effect in p+Pb and Pb+Pb collisions is another way to assess the size of background effects. Any magnetic field-driven effect, such as the CME, should be greatly suppressed in p+Pb collisions in comparison with Pb+Pb collisions (by a factor $(1/Z_{Pb}^2 = 1.5 \times 10^{-4})$. Data from the CMS experiment constrain possible contributions of the CME to the Pb+Pb data to less than 7% [409].

A more sensitive search for CME signals requires a suppression or cancellation of such background effects. This motivated a comparative study of two collision systems involving nuclear isobars, ^{96}Zr and ^{96}Ru, which was carried out at RHIC. The magnetic field produced in Ru+Ru collisions is larger than that produced in Zr+Zr collisions under otherwise identical conditions, because a ^{96}Zr nucleus contains 40 protons, while a ^{96}Ru nucleus contains 44 protons. As the CME observables are proportional to the square of the magnetic field, one expects roughly a 15% difference between the two systems for the CME contribution to any observable. Great care was taken to ensure that the experimental conditions for collisions in the two isobar systems were identical, the data were subjected to a sophisticated blind analysis protocol [410], the first of its kind in the field of relativistic heavy ion physics, and the analysis was performed independently by several groups.

The results published by the STAR collaboration [411] showed no evidence for the presence of a CME contribution to any of the predefined observables with an experimental precision of $\pm 4 \times 10^{-3}$. Figure 80 shows the measured ratios $S(\text{Ru})/S(\text{Zr})$ for each of the
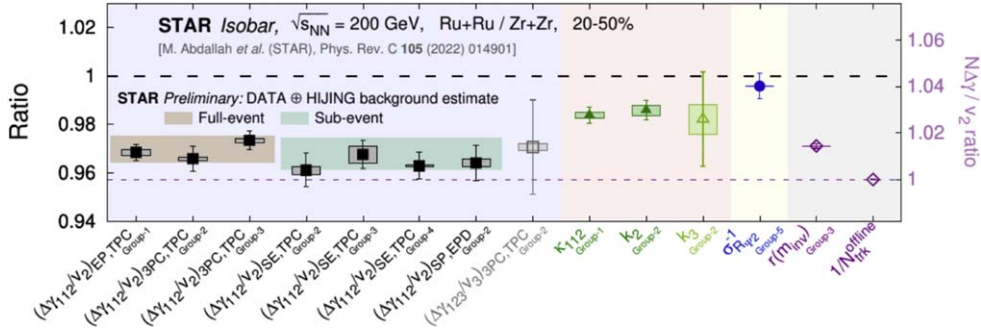
Figure 80. Results of the blind analysis of observables S sensitive to the CME in the ^{96}Zr$-^{96}$Ru isobar system. Shown is the ratio $S(\text{Ru})/S(\text{Zr})$. The horizontal shaded bars indicate the baseline without CME contribution, corrected for the nuclear shape difference and minor efficiency effects. A possible CME contribution would cause this ratio to be higher than the baseline. For details see [411, 412]. Reprinted (figure) with permission from [411], Copyright (2022) by the American Physical Society.

signature observables S considered in the analysis. All signals are in some way related to the difference $\Delta\gamma$ for same-sign and opposite-sign charged pairs of particles of the quantity γ defined in (8), normalized to the elliptic flow anisotropy v_2 that drives the background effects. A contribution from the CME would cause this ratio to be larger than unity. Clearly, all measured ratios lie well below one, which means that they do not provide evidence for a CME contribution.

The fact that all ratios related to $\Delta\gamma/v_2$ cluster around a value of 0.97 suggests that they have a common source that can be traced to a difference between the two isobars, which is not related to the nuclear charge. Indeed, such differences are known to exist: ^{96}Zr has a thicker neutron skin than ^{96}Ru (the ^{96}Zr nucleus has four additional neutrons), and the two nuclei have different quadrupole deformations. The experiment clearly revealed these differences in the centrality dependence of the final-state multiplicity distributions. This means that the shape differences of the two isobar nuclei are the source of the largest systematic uncertainty as background effects associated with elliptic flow do not cancel exactly. The resulting estimate for the background ratio of $\Delta\gamma/v_2$ is indicated by the shaded bars in figure 80. The measured value for the ratio, $R_{\text{Ru/Zr}} = 1.013 \pm 0.003 \pm 0.005$, is in excellent agreement with this baseline [412] and does not show evidence for a CME contribution.

6. Future opportunities

The physics of heavy ion collisions is a vibrant field of research, and there are many different avenues for future progress in understanding strongly interacting matter under extreme conditions. In general, the field is driven by experimental measurements and their theoretical interpretation, since only very few predictions from first principle calculations have been possible. It is important to identify questions that require more precise or more comprehensive data on well-studied observables or completely new types of measurements and analysis.

There are four areas which are going to be explored in the near-to-midterm future:

- Precision measurements with ultra-relativistic nucleus-nucleus collisions at RHIC with sPHENIX [413] and STAR and at LHC with ALICE, ATLAS, and CMS and LHCb;

- A whole new level of heavy flavor and electromagnetic probes measurements with ALICE3 [414];
- Ultraperipheral collisions and the transition to the electron-ion collider [415];
- Measurements at low beam energies with final results from the RHIC beam energy scan, HADES at GSI and future FAIR.

Let us point out the main physics case and opportunities for each of the four broad directions mentioned above.

The upcoming LHC runs 3 and 4, where much higher luminosities are going to be achieved, will generate excellent data sets for proton–proton, proton–nucleus, and nucleus–nucleus collisions in the TeV energy range. In addition to Pb+Pb collisions, there is the proposal to add O+O collisions to enhance the understanding of the transition between very small systems and large collision systems. Collisions of different species of ions might be employed to explore nuclear structure in heavy ion collisions.

All four experimental collaborations at the LHC (ALICE, ATLAS, CMS, and LHCb) are planning to participate in the heavy ion runs. Very precise measurements of bulk observables and hard probes are expected. In a complementary campaign the sPHENIX and STAR will be taking data at the lower beam energy of RHIC, but with comparable precision for hard probes. To understand the temperature dependence of transport coefficients like $\hat{q}$ the lever arm in beam energy between $\sqrt{s_{NN}} = 200$ GeV and a few TeV is going to be important.

For Run 5 expected in the 2030s the ALICE collaboration has proposed to construct a completely new detector named ALICE3. This device will allow for an entirely new level of precision for measurements of heavy flavor probes in the soft and hard sector as well as electromagnetic probes down to very low transverse momenta. In addition, new capabilities to identify multi-charm hadronic states will become available. ALICE3 will provide detailed constraints for the more refined theoretical modeling that is going to ready by the time the new detector will come online.

In addition to the study of the midrapidity region, where the quark-gluon plasma production can be studied, it is going to be of interest to explore the whole longitudinal phase space. For certain investigations a fixed target setup at collider facilities provides advantages and is being explored at LHC to study particle production in the fragmentation region [416–418]. Forward rapidities may also be useful to investigate the dependence on net baryon content of the system.

Extremely peripheral ('ultraperipheral') events also receive increasing attention [419]. The idea is that in such collisions the nuclei pass without coming into direct contact but interact via electromagnetic interactions that are enhanced by the Lorentz factor γ and powers of the nuclear charge Z (Z^2 for processes involving single initial-state photons and Z^4 for photon–photon collisions). Exploring the interplay between electromagnetic and strong interactions under such conditions paves the path to the planned electron-ion collider (EIC). In 2020, the decision was been made to construct such a facilty at Brookhaven National Laboratory, which will allows for collisions of electrons with light and heavy ions at $\sqrt{s_{eN}}$ up to 140 GeV. The electron-ion collider will generate unprecedented insights into the quark-gluon structure of the proton and complex nuclei. This will provide very detailed information for the initial state of a heavy ion collision.

Moving forward with the exploration of the QCD phase diagram at nonzero net baryon chemical potential, several experimental efforts are worth mentioning. The beam energy scan program at RHIC has just finished Phase 2, and precision measurements are expected to be published by the STAR collaboration in the near future for energies down to $\sqrt{s_{NN}} = 3$ GeV. The HADES experiment [420] at GSI is running within the FAIR Phase-0 program and will

measure the excitation function of observables in the fixed target beam energy regime around $E_{\text{lab}}/A \approx 1\,\text{GeV}$ per nucleon ($\sqrt{s_{\text{NN}}} \approx 2.4\,\text{GeV}$). In the later part of the decade, SIS-100 will be completed and the CBM (Compressed Baryonic Matter) experiment [421] will measure rare probes with unprecedented precision in the beam energy region from $E_{\text{lab}}/A = 3.4 - 12\,\text{GeV}$ ($\sqrt{s_{\text{NN}}} \approx 3.3 - 5.3\,\text{GeV}$). There are also other projects proposed around the world, like a heavy ion extension of J-PARC in Japan [422], HIAF in China [423], and NICA in Dubna [424].

These lower-energy facilities aim to determine the structure of the QCD phase diagram, in particular, to find out whether there is a first-order phase transition between the hadron gas and the quark-gluon plasma at nonzero baryon chemical potential with a critical endpoint or whether the transition is a cross-over everywhere [6]. Knowing the nuclear matter equation of state at high net baryon density is also important for the understanding of neutron star mergers. Since the first detection of gravitational waves from such a merger event in 2017, there has been increasing interaction between astrophysicists and the heavy ion physics community, which is likely to further intensify in the years to come.

All of the expected measurements will need to be accompanied by theoretical progress in the understanding of QCD matter and high-energy nuclear reactions. Such progress relies on fundamental theory developments based on lattice QCD techniques, functional methods, and effective field theories as well as on sophisticated dynamical modeling that connects properties of strongly interacting matter to experimental measurements. More standardized ways to compare data to calculations will be helpful for quantitative conclusions, e.g. based on HepMC and RIVET adapted for heavy ions. Modern analysis tools based on machine learning and deep learning methods as well as potential applications of quantum computing will complement the more traditional efforts.

7. Summary

Relativistic heavy ion collisions produce matter with the highest energy density known in nature, thereby recreating conditions similar to those in the early universe or in neutron star mergers. We now know that this matter, the quark-gluon plasma, is also the most 'perfect' fluid and endowed with high vorticity. This conclusion has been reached by a concerted theoretical and experimental effort over the last three decades. Many detailed measurements and sophisticated calculations enabled by technological advances have led to a 'standard model' for relativistic heavy ion collisions that is based on non-equilibrium initial conditions, viscous hydrodynamics and hadronic transport. While the quantitative insight into the properties of the quark-gluon plasma has lately seen remarkable progress by the application of Bayesian multi-parameter model-to-data comparisons, a more complete understanding of the structures in the QCD phase diagram—a potential first-order phase transition and critical endpoint—requires further theoretical developments and a new level of experimental precision.

Given the multitude of available beam energies, collision systems, and experimental probes it is important not to lose one's overview. Everyone working in this field must from time to time ask themselves how their current project connects to the major physics questions of the field and what can be learned by looking from a broader perspective. This is especially true for scientists at the beginning of their career, who have not yet developed the breadth of knowledge and insight of more experienced scientists. In a field as complex as this it is easy to be misled to conclusions based on a limited set of data and observables. To avoid going down the wrong path, the sophisticated dynamical models now available need to be consistently applied to as many observables as possible with the same parameter settings. In

doing so, it is crucial to apply the right methods in the right places and to be aware of the limits of applicability of each of them. Getting a prediction from a complex numerical code does not guarantee that it is physically meaningful!

Relativistic heavy ion physics is attractive to young researchers owing to its mode of global collaboration, on the experimental as well as increasingly on the theoretical side. The multitude of different methods that are being applied to further our knowledge makes it an ideal training ground.

There are close connections to other fields within nuclear physics, for example, hadron physics, nuclear structure physics, and nuclear astrophysics. More recently, the connection to the astrophysics community has intensified since it was realized that heavy ion collisions at low beam energy allows us to produce and study conditions in the laboratory that resemble those occurring in neutron star mergers. The hadronic interactions that are relevant in heavy ion physics, are also of interest to the astroparticle physics community for the description of cosmic air showers. The non-equilibrium phase transitions and chiral phenomena encountered in heavy ion collisions have connections to phenomena of interest to the condensed matter physics community. Modern computational techniques applied to data acquisition, model-to-data comparison, and simulation of processes at the quantum level are closely linked to current developments in computer science and statistics.

We hope that this review will help beginning and more experienced scientists alike to get a more complete appreciation for the wealth of phenomena and approaches that are currently available to study and understand relativistic heavy ion collisions. Impressive progress is being made, and the future opportunities that await those who venture into this field of research are great.

Acknowledgments

Data availability statement

No new data were created or analyzed in this study.

ORCID iDs

Hannah Elfner ⓘ https://orcid.org/0000-0002-6213-3613
Berndt Müller ⓘ https://orcid.org/0000-0003-2468-3996

References

[1] Schwarz D J 2003 *Ann. Phys.* **12** 220–70
[2] Aoki Y, Fodor Z, Katz S D and Szabo K K 2006 *Phys. Lett.* B **643** 46–54

[3]   Aoki Y, Endrodi G, Fodor Z, Katz S D and Szabo K K 2006 *Nature* **443** 675–8
[4]   Fukushima K 2004 *Phys. Lett.* B **591** 277–84
[5]   Roessner S, Hell T, Ratti C and Weise W 2008 *Nucl. Phys. A* **814** 118–43
[6]   Luo X, Wang Q, Xu N and Zhuang P 2022 *Properties of QCD Matter at High Baryon Density*
      (Singapore: Springer) (https://doi.org/10.1007/978-981-19-4441-3)
[7]   Siemens P J 1983 *Nature* **305** 410–2
[8]   Chomaz P 2002 *AIP Conf. Proc.* **610** 167
[9]   Alford M G, Schmitt A, Rajagopal K and Schäfer T 2008 *Rev. Mod. Phys.* **80** 1455–515
[10] Stephanov M A 2004 *Prog. Theor. Phys. Suppl.* **153** 139–56
[11] Arsene I *et al* (BRAHMS) 2005 *Nucl. Phys. A* **757** 1–27
[12] Adcox K *et al* (PHENIX) 2005 *Nucl. Phys. A* **757** 184–283
[13] Back B B *et al* 2005 *Nucl. Phys. A* **757** 28–101
[14] Adams J *et al* (STAR) 2005 *Nucl. Phys. A* **757** 102–83
[15] Müller B and Nagle J L 2006 *Ann. Rev. Nucl. Part. Sci.* **56** 93–135
[16] Bannur V M 1999 *Eur. Phys. J.* C **11** 169–71
[17] Gyulassy M and McLerran L 2005 *Nucl. Phys. A* **750** 30–63
[18] Schukraft J (ALICE) 2011 *J. Phys.* G **38** 124003
[19] Steinberg P (ATLAS) 2011 *J. Phys.* G **38** 124004
[20] Wyslouch B (CMS) 2011 *J. Phys.* G **38** 124005
[21] Müller B, Schukraft J and Wyslouch B 2012 *Ann. Rev. Nucl. Part. Sci.* **62** 361–86
[22] Braun-Munzinger P, Koch V, Schäfer T and Stachel J 2016 *Phys. Rept.* **621** 76–126
[23] Busza W, Rajagopal K and van der Schee W 2018 *Ann. Rev. Nucl. Part. Sci.* **68** 339–76
[24] Casalderrey-Solana J, Liu H, Mateos D, Rajagopal K and Wiedemann U A 2014 *Gauge/String
      Duality, Hot QCD and Heavy Ion Collisions* (Cambridge: Cambridge University Press)
      (https://doi.org/10.1017/CBO9781139136747)
[25] Kapusta J, Müller B and Rafelski J 2003 *Quark-Gluon Plasma: Theoretical Foundations* (New
      York: Elsevier)
[26] Yagi K, Hatsuda T and Miake Y 2005 *Cambridge Monogr. Part. Phys. Nucl. Phys. Cosmol.* **23**
      1–446
[27] Letessier J and Rafelski J 2022 *Hadrons and Quark–Gluon Plasma* (Cambridge University
      Press) 18, 1–397
[28] Czajka A and Jeon S 2017 *Phys. Rev.* C **95** 064906
[29] Majumder A and Van Leeuwen M 2011 *Prog. Part. Nucl. Phys.* **66** 41–92
[30] Hanbury Brown R and Twiss R Q 1954 *Phil. Mag. Ser. 7* **45** 663–82
[31] Lisa M A, Pratt S, Soltz R and Wiedemann U 2005 *Ann. Rev. Nucl. Part. Sci.* **55** 357–402
[32] Sveshnikov N A and Tkachov F V 1996 *Phys. Lett.* B **382** 403–8
[33] Komiske P T, Moult I, Thaler J and Zhu H X 2023 *Phys. Rev. Lett.* **130** 051901
[34] Chang H M, Procura M, Thaler J and Waalewijn W J 2013 *Phys. Rev. Lett.* **111** 102002
[35] Chang H M, Procura M, Thaler J and Waalewijn W J 2013 *Phys. Rev.* D **88** 034030
[36] Li Y, Moult I, van Velzen S S, Waalewijn W J and Zhu H X 2022 *Phys. Rev. Lett.* **128** 182001
[37] Kaplan D B 1992 *Phys. Lett.* B **288** 342–7
[38] Narayanan R and Neuberger H 1994 *Nucl. Phys.* B **412** 574–606
[39] Kuti J, Polonyi J and Szlachanyi K 1981 *Phys. Lett.* B **98** 199
[40] McLerran L D and Svetitsky B 1981 *Phys. Rev.* D **24** 450
[41] Borsanyi S, Endrodi G, Fodor Z, Jakovac A, Katz S D, Krieg S, Ratti C and Szabo K K 2010
      J. High Energy Phys. JHEP11(2010)077
[42] Borsanyi S, Fodor Z, Hoelbling C, Katz S D, Krieg S and Szabo K K 2014 *Phys. Lett.* B **730**
      99–104
[43] Bazavov A *et al* (HotQCD) 2014 *Phys. Rev.* D **90** 094503
[44] Allton C R *et al* 2005 *Phys. Rev.* D **71** 054508
[45] de Forcrand P and Philipsen O 2002 *Nucl. Phys.* B **642** 290–306
[46] Fodor Z and Katz S D 2002 *J. High Energy Phys.* JHEP03(2002)014
[47] de Forcrand P, Kim S and Philipsen O 2007 *PoS* **178** LAT2007
[48] Borsanyi S, Fodor Z, Hoelbling C, Katz S D, Krieg S, Ratti C, Szabo K K and (Wuppertal-
      Budapest) 2010 *J. High Energy Phys.* JHEP09(2010)073
[49] Borsanyi S, Fodor Z, Guenther J N, Kara R, Katz S D, Parotto P, Pasztor A, Ratti C and
      Szabo K K 2020 *Phys. Rev. Lett.* **125** 052001
[50] Majumder A and Müller B 2010 *Phys. Rev. Lett.* **105** 252002

[51] Bazavov A *et al* 2014 *Phys. Rev. Lett.* **113** 072001
[52] Andronic A, Braun-Munzinger P, Friman B, Lo P M, Redlich K and Stachel J 2019 *Phys. Lett.* B **792** 304–9
[53] Andersen J O, Leganger L E, Strickland M and Su N 2011 *Phys. Lett.* B **696** 468–72
[54] Andersen J O, Leganger L E, Strickland M and Su N 2011 *J. High Energy Phys.* JHEP08 (2011)053
[55] Peshier A, Kämpfer B and Soff G 2000 *Phys. Rev.* C **61** 045203
[56] Ivanov Y B, Skokov V V and Toneev V D 2005 *Phys. Rev.* D **71** 014005
[57] Braaten E and Pisarski R D 1990 *Nucl. Phys.* B **337** 569
[58] Kapusta J I and Gale C 2011 Finite-temperature field theory: Principles and applications *Cambridge Monographs on Mathematical Physics* (Cambridge: Cambridge University Press) (https://doi.org/10.1017/CBO9780511535130)
[59] Le Bellac M 2000 *Thermal Field Theory* (Cambridge: Cambridge University Press)
[60] Haque N, Bandyopadhyay A, Andersen J O, Mustafa M G, Strickland M and Su N 2014 *J. High Energy Phys.* JHEP02(2014)027
[61] Arnold P, Moore G D and Yaffe L G 2000 *J. High Energy Phys.* JHEP11(2000)001
[62] Arnold P, Moore G D and Yaffe L G 2003 *J. High Energy Phys.* JHEP05(2003)051
[63] Ghiglieri J, Moore G D and Teaney D 2018 *J. High Energy Phys.* JHEP03(2018)179
[64] Caron-Huot S 2009 *Phys. Rev.* D **79** 065039
[65] Ghiglieri J and Teaney D 2015 *Int. J. Mod. Phys.* E **24** 1530013
[66] Ghiglieri J, Moore G D and Teaney D 2016 *J. High Energy Phys.* JHEP03(2016)095
[67] Thoma M H 1991 *Phys. Lett.* B **273** 128–32
[68] Caron-Huot S and Moore G D 2008 *Phys. Rev. Lett.* **100** 052301
[69] Arnold P, Moore G D and Yaffe L G 2001 *J. High Energy Phys.* JHEP12(2001)009
[70] Ghiglieri J, Hong J, Kurkela A, Lu E, Moore G D and Teaney D 2013 *J. High Energy Phys.* JHEP05(2013)010
[71] Braaten E, Pisarski R D and Yuan T C 1990 *Phys. Rev. Lett.* **64** 2242
[72] Ghiglieri J and Moore G D 2014 *J. High Energy Phys.* JHEP12(2014)029
[73] Ghiglieri J 2016 *Nucl. Part. Phys. Proc.* **276-278** 305–8
[74] Maldacena J M 1998 *Adv. Theor. Math. Phys.* **2** 231–52
[75] Aharony O, Gubser S S, Maldacena J M, Ooguri H and Oz Y 2000 *Phys. Rept.* **323** 183–386
[76] Natsuume M 2015 *AdS/CFT Duality User Guide* (Tokyo: Springer) (https://doi.org/10.1007/978-4-431-55441-7)
[77] Gursoy U and Kiritsis E 2008 *J. High Energy Phys.* JHEP02(2008)032
[78] Karch A, Katz E, Son D T and Stephanov M A 2006 *Phys. Rev.* D **74** 015005
[79] de Teramond G F and Brodsky S J 2009 *Phys. Rev. Lett.* **102** 081601
[80] Son D T and Starinets A O 2002 *J. High Energy Phys.* JHEP09(2002)042
[81] Balasubramanian V, Bernamonti A, de Boer J, Copland N, Craps B, Keski-Vakkuri E, Müller B, Schäfer A, Shigemori M and Staessens W 2011 *Phys. Rev. Lett.* **106** 191601
[82] Balasubramanian V, Bernamonti A, de Boer J, Copland N, Craps B, Keski-Vakkuri E, Müller B, Schäfer A, Shigemori M and Staessens W 2011 *Phys. Rev.* D **84** 026010
[83] Price R H and Thorne K S 1986 *Phys. Rev.* D **33** 915–41
[84] Policastro G, Son D T and Starinets A O 2001 *Phys. Rev. Lett.* **87** 081601
[85] Kovtun P, Son D T and Starinets A O 2005 *Phys. Rev. Lett.* **94** 111601
[86] Ecker C, Grumiller D, Stanzer P, Stricker S A and van der Schee W 2016 *J. High Energy Phys.* JHEP11(2016)054
[87] Herzog C P, Karch A, Kovtun P, Kozcaz C and Yaffe L G 2006 *J. High Energy Phys.* JHEP07 (2006)013
[88] Liu H, Rajagopal K and Wiedemann U A 2006 *Phys. Rev. Lett.* **97** 182301
[89] Chesler P M, Jensen K, Karch A and Yaffe L G 2009 *Phys. Rev.* D **79** 125015
[90] Janik R A and Peschanski R B 2006 *Phys. Rev.* D **73** 045013
[91] Heller M P, Janik R A and Witaszczyk P 2012 *Phys. Rev. Lett.* **108** 201602
[92] Shuryak E 2012 *J. Phys.* G **39** 054001
[93] Chesler P M and Yaffe L G 2010 *Phys. Rev.* D **82** 026006
[94] Chesler P M and Yaffe L G 2011 *Phys. Rev. Lett.* **106** 021601
[95] Chesler P M and Yaffe L G 2014 *J. High Energy Phys.* JHEP07(2014)086
[96] Casalderrey-Solana J, Heller M P, Mateos D and van der Schee W 2013 *Phys. Rev. Lett.* **111** 181601

[97]   Casalderrey-Solana J, Heller M P, Mateos D and van der Schee W 2014 *Phys. Rev. Lett.* **112** 221602

[98]   Waeber S, Rabenstein A, Schäfer A and Yaffe L G 2019 *J. High Energy Phys.* JHEP08(2019)005

[99]   Attems M, Casalderrey-Solana J, Mateos D, Santos-Oliván D, Sopuerta C F, Triana M and Zilhão M 2017 *J. High Energy Phys.* JHEP06(2017)154

[100]   van der Schee W, Romatschke P and Pratt S 2013 *Phys. Rev. Lett.* **111** 222302

[101]   Shukla P 2001 arXiv:nucl-th/0112039

[102]   Stock R 2020 *Springer Proc. Phys.* **250** 3–12

[103]   Kharzeev D and Nardi M 2001 *Phys. Lett.* B **507** 121–8

[104]   Blaizot J P, Broniowski W and Ollitrault J Y 2014 *Phys. Rev.* C **90** 034906

[105]   d'Enterria D and Loizides C 2021 *Ann. Rev. Nucl. Part. Sci.* **71** 315–44

[106]   Alver B and Roland G 2010 *Phys. Rev.* **81** 054905 [Erratum: Phys. Rev. C 82, 039903 (2010)]

[107]   Gelis F, Iancu E, Jalilian-Marian J and Venugopalan R 2010 *Ann. Rev. Nucl. Part. Sci.* **60** 463–89

[108]   McLerran L D and Venugopalan R 1994 *Phys. Rev.* D **49** 2233–41

[109]   McLerran L D and Venugopalan R 1994 *Phys. Rev.* D **49** 3352–5

[110]   Lappi T 2008 *Eur. Phys. J.* C **55** 285–92

[111]   Mueller A H 2001 *Phys. Lett.* B **523** 243–8

[112]   Rummukainen K and Weigert H 2004 *Nucl. Phys.* A **739** 183–226

[113]   Praszalowicz M 2011 *Phys. Lett.* B **704** 566–9

[114]   Kowalski H and Teaney D 2003 *Phys. Rev.* D **68** 114005

[115]   Gelis F, Lappi T and McLerran L 2009 *Nucl. Phys.* A **828** 149–60

[116]   Schenke B, Tribedy P and Venugopalan R 2012 *Phys. Rev. Lett.* **108** 252301

[117]   Moreland J S, Bernhard J E and Bass S A 2015 *Phys. Rev.* C **92** 011901

[118]   Baier R, Mueller A H, Schiff D and Son D T 2001 *Phys. Lett.* B **502** 51–8

[119]   Kurkela A, Mazeliauskas A, Paquet J F, Schlichting S and Teaney D 2019 *Phys. Rev.* C **99** 034910

[120]   Huovinen P, Kolb P F, Heinz U W, Ruuskanen P V and Voloshin S A 2001 *Phys. Lett.* B **503** 58–64

[121]   Kolb P F and Heinz U W 2003 634–714 arXiv:nucl-th/0305084

[122]   Denicol G S, Jeon S and Gale C 2014 *Phys. Rev.* C **90** 024912

[123]   Ollitrault J Y 2008 *Eur. J. Phys.* **29** 275–302

[124]   Romatschke P 2010 *Int. J. Mod. Phys.* E **19** 1–53

[125]   Gale C, Jeon S and Schenke B 2013 *Int. J. Mod. Phys.* A **28** 1340011

[126]   Huovinen P and Petreczky P 2010 *Nucl. Phys.* A **837** 26–53

[127]   Noronha-Hostler J, Parotto P, Ratti C and Stafford J M 2019 *Phys. Rev.* C **100** 064910

[128]   Monnai A, Schenke B and Shen C 2019 *Phys. Rev.* C **100** 024907

[129]   Grefa J, Noronha J, Noronha-Hostler J, Portillo I, Ratti C and Rougemont R 2021 *Phys. Rev.* D **104** 034002

[130]   Bjorken J D 1983 *Phys. Rev.* D **27** 140–51

[131]   Gale C, Paquet J F, Schenke B and Shen C 2021 *PoS* **039** HardProbes2020

[132]   Fotakis J A, Greif M, Greiner C, Denicol G S and Niemi H 2020 *Phys. Rev.* D **101** 076007

[133]   Schenke B, Jeon S and Gale C 2010 *Phys. Rev.* C **82** 014903

[134]   Karpenko I, Huovinen P and Bleicher M 2014 *Comput. Phys. Commun.* **185** 3016–27

[135]   Shen C, Qiu Z, Song H, Bernhard J, Bass S and Heinz U 2016 *Comput. Phys. Commun.* **199** 61–85

[136]   Pang L G, Petersen H and Wang X N 2018 *Phys. Rev.* C **97** 064918

[137]   Denicol G S, Niemi H, Molnar E and Rischke D H 2012 *Phys. Rev.* D **85** 114047 [Erratum: Phys. Rev.D 91, 039902 (2015)]

[138]   McNelis M, Bazow D and Heinz U 2018 *Phys. Rev.* C **97** 054912

[139]   Alqahtani M, Nopoush M and Strickland M 2018 *Prog. Part. Nucl. Phys.* **101** 204–48

[140]   Weickgenannt N, Wagner D, Speranza E and Rischke D 2022 *Phys. Rev. D* **106** 096014

[141]   Bhadury S, Florkowski W, Jaiswal A, Kumar A and Ryblewski R 2021 *Phys. Rev. D* **103** 014030

[142]   Ammon M, Grieninger S, Hernandez J, Kaminski M, Koirala R, Leiber J and Wu J 2021 *J. High Energy Phys.* JHEP04(2021)078

[143]   Xu Z and Greiner C 2005 *Phys. Rev.* C **71** 064901

[144]   Lin Z W, Ko C M, Li B A, Zhang B and Pal S 2005 *Phys. Rev.* C **72** 064901

[145]   Cassing W and Bratkovskaya E L 2009 *Nucl. Phys.* A **831** 215–42

[146] Petersen H 2014 *J. Phys.* G **41** 124005
[147] Huovinen P and Petersen H 2012 *Eur. Phys. J.* A **48** 171
[148] Cooper F and Frye G 1974 *Phys. Rev.* D **10** 186
[149] Everett D *et al* (JETSCAPE) 2021 *Phys. Rev. Lett.* **126** 242301
[150] Bass S A *et al* 1998 *Prog. Part. Nucl. Phys.* **41** 255–369
[151] Bleicher M *et al* 1999 *J. Phys.* G **25** 1859–96
[152] Weil J *et al* 2016 *Phys. Rev.* C **94** 054905
[153] Plumberg C and Heinz U 2015 *Phys. Rev.* C **91** 054905
[154] Oliinychenko D, Huovinen P and Petersen H 2015 *Phys. Rev.* C **91** 024906
[155] Oliinychenko D and Koch V 2019 *Phys. Rev. Lett.* **123** 182302
[156] Majumder A, Müller B and Mrowczynski S 2009 *Phys. Rev.* D **80** 125020
[157] Guo X F and Wang X N 2000 *Phys. Rev. Lett.* **85** 3591–4
[158] Fries R J, Müller B and Srivastava D K 2003 *Phys. Rev. Lett.* **90** 132301
[159] Zapp K C 2014 *Eur. Phys. J.* C **74** 2762
[160] Putschke J H *et al* 2019 arXiv:1903.07706
[161] Dokshitzer Y L, Khoze V A and Troian S I 1991 *J. Phys.* G **17** 1602–4
[162] Petreczky P and Teaney D 2006 *Phys. Rev.* D **73** 014508
[163] Casalderrey-Solana J and Teaney D 2006 *Phys. Rev.* D **74** 085012
[164] Ke W, Xu Y and Bass S A 2018 *Phys. Rev.* C **98** 064901
[165] Dai T, Paquet J F, Teaney D and Bass S A 2022 *Phys. Rev.* C **105** 034905
[166] Burke K M *et al* (JET) 2014 *Phys. Rev.* C **90** 014909
[167] Neufeld R B and Renk T 2010 *Phys. Rev.* C **82** 044903
[168] Tachibana Y, Chang N B and Qin G Y 2017 *Phys. Rev.* C **95** 044909ace
[169] Kunnawalkam Elayavalli R and Zapp K C 2017 *J. High Energy Phys.* JHEP07(2017)141
[170] Floerchinger S, Gebhardt C and Reygers K 2021 *Phys. Lett.* B **837** 137647 arXiv:2112.12497
[171] Gale C and Haglin K L 2003 364–429 arXiv:hep-ph/0306098
[172] Paquet J F 2017 *Nucl. Part. Phys. Proc.* **289-290** 89–94
[173] Reygers K 2022 Experimental overview of electromagnetic probes in ultra-relativistic nucleus-
      nucleus collisions *29th Int. Conf. on Ultra-relativistic Nucleus-Nucleus Collisions*
[174] David G 2020 *Rept. Prog. Phys.* **83** 046301
[175] Rapp R and van Hees H 2016 *Eur. Phys. J.* A **52** 257
[176] Ding H T, Kaczmarek O and Meyer F 2016 *Phys. Rev.* D **94** 034504
[177] Ghiglieri J, Kaczmarek O, Laine M and Meyer F 2016 *Phys. Rev.* D **94** 016005
[178] Kapusta J I and Shuryak E V 1994 *Phys. Rev.* D **49** 4694–704
[179] Holt N P M, Hohler P M and Rapp R 2013 *Phys. Rev.* D **87** 076010
[180] Adam J *et al* (ALICE) 2017 *Phys. Rev.* C **95** 064606
[181] Becattini F, Castorina P, Manninen J and Satz H 2008 *Eur. Phys. J.* C **56** 493–510
[182] Andronic A, Braun-Munzinger P, Redlich K and Stachel J 2018 *Nature* **561** 321–30
[183] Zhang Z and Ko C M 2018 *Phys. Lett.* B **780** 191–5
[184] Müller B and Schäfer A 2017 arXiv:1712.03567
[185] Cleymans J, Oeschler H, Redlich K and Wheaton S 2006 *Phys. Rev.* C **73** 034905
[186] Agakishiev G *et al* (HADES) 2016 *Eur. Phys. J.* A **52** 178
[187] Bellwied R, Noronha-Hostler J, Parotto P, Portillo Vazquez I, Ratti C and Stafford J M 2019
      Phys. Rev. C **99** 034912
[188] Rafelski J and Müller B 1982 *Phys. Rev. Lett.* **48** 1066 [Erratum: Phys.Rev.Lett. 56, 2334 (1986)]
[189] Afanasiev S V *et al* (NA49) 2002 *Phys. Rev.* C **66** 054902
[190] Steinheimer J and Bleicher M 2011 *Phys. Rev.* C **84** 024905
[191] Petersen H, Steinheimer J, Bleicher M and Stocker H 2009 *J. Phys.* G **36** 055104
[192] Vovchenko V and Stoecker H 2019 *Comput. Phys. Commun.* **244** 295–310
[193] Gorenstein M I, Hauer M and Moroz O N 2008 *Phys. Rev.* C **77** 024911
[194] Steinheimer J, Aichelin J and Bleicher M 2013 *Phys. Rev. Lett.* **110** 042501
[195] Blume C 2005 *J. Phys.* G **31** S57–68
[196] Anishetty R, Koehler P and McLerran L D 1980 *Phys. Rev.* D **22** 2793
[197] Abelev B *et al* (ALICE) 2013 *Phys. Rev.* C **88** 044910
[198] Adamczyk L *et al* (STAR) 2017 *Phys. Rev.* C **96** 044904
[199] Schnedermann E, Sollfrank J and Heinz U W 1993 *Phys. Rev.* C **48** 2462–75
[200] Song H, Bass S A, Heinz U, Hirano T and Shen C 2011 *Phys. Rev.* **83** 054910 [Erratum: Phys.
      Rev. C 86, 059903 (2012)]

[201] Ryu S, Paquet J F, Shen C, Denicol G S, Schenke B, Jeon S and Gale C 2015 *Phys. Rev. Lett.* **115** 132301

[202] Ryu S, Paquet J F, Shen C, Denicol G, Schenke B, Jeon S and Gale C 2018 *Phys. Rev. C* **97** 034910

[203] Everett D *et al* (JETSCAPE) 2022 *Phys. Rev. C* **106** 064901 arXiv:2203.08286

[204] Adamczyk L *et al* (STAR) 2015 *Phys. Rev. C* **92** 014904

[205] Adam J *et al* (ALICE) 2016 *Phys. Rev. C* **93** 024905

[206] Adam J *et al* (ALICE) 2015 *Phys. Rev. C* **92** 054908

[207] Acharya S *et al* (ALICE) 2021 *Phys. Lett. B* **813** 136030

[208] Pratt S 2009 *Phys. Rev. Lett.* **102** 232301

[209] Li Q f, Steinheimer J, Petersen H, Bleicher M and Stocker H 2009 *Phys. Lett. B* **674** 111–6

[210] Rischke D H and Gyulassy M 1996 *Nucl. Phys. A* **608** 479–512

[211] Heinz U W and Kolb P F 2002 Two RHIC puzzles: early thermalization and the HBT problem *18th Winter Workshop on Nuclear Dynamics*

[212] O'Hara K M, Hemmer S L, Gehm M E, Granade S R and Thomas J E 2002 *Science* **298** 2179–82

[213] Teaney D and Yan L 2012 *Phys. Rev. C* **86** 044908

[214] Schenke B, Jeon S and Gale C 2012 *Phys. Rev. C* **85** 024901

[215] Luzum M and Petersen H 2014 *J. Phys. G* **41** 063102

[216] Alver B H, Gombeaud C, Luzum M and Ollitrault J Y 2010 *Phys. Rev. C* **82** 034913

[217] Schenke B, Jeon S and Gale C 2011 *Phys. Rev. Lett.* **106** 042301

[218] Petersen H, Qin G Y, Bass S A and Müller B 2010 *Phys. Rev. C* **82** 041901

[219] Adare A *et al* (PHENIX) 2011 *Phys. Rev. Lett.* **107** 252301

[220] Aad G *et al* (ATLAS) 2012 *Phys. Rev. C* **86** 014907

[221] Aad G *et al* (ATLAS) 2021 *Phys. Rev. Lett.* **126** 122301

[222] Nie M (STAR) 2021 *Nucl. Phys. A* **1005** 121783

[223] Sakai A, Murase K and Hirano T 2022 *Phys. Lett. B* **829** 137053 arXiv:2111.08963

[224] Auvinen J and Petersen H 2013 *Phys. Rev. C* **88** 064908

[225] Brachmann J, Soff S, Dumitru A, Stoecker H, Maruhn J A, Greiner W, Bravina L V and Rischke D H 2000 *Phys. Rev. C* **61** 024909

[226] Adamczyk L *et al* (STAR) 2014 *Phys. Rev. Lett.* **112** 162301

[227] Khachatryan V *et al* (CMS) 2010 *J. High Energy Phys.* JHEP09(2010)091

[228] Chatrchyan S *et al* (CMS) 2013 *Phys. Lett. B* **718** 795–814

[229] Aad G *et al* (ATLAS) 2013 *Phys. Rev. Lett.* **110** 182302

[230] Abelev B *et al* (ALICE) 2013 *Phys. Lett. B* **719** 29–41

[231] Aidala C *et al* (PHENIX) 2019 *Nature Phys.* **15** 214–20

[232] Sirunyan A M *et al* (CMS) 2018 *Phys. Rev. Lett.* **121** 082301

[233] Adam J *et al* (ALICE) 2017 *Nature Phys.* **13** 535–9

[234] Šefčík M (ALICE) 2020 *Springer Proc. Phys.* **250** 173–6

[235] Acharya S *et al* (ALICE) 2018 *Phys. Lett. B* **783** 95–113

[236] Tywoniuk K 2014 *Nucl. Phys. A* **926** 85–91

[237] Soloviev A 2022 *Eur. Phys. J. C* **82** 319 arXiv:2109.15081

[238] Baier R, Romatschke P, Son D T, Starinets A O and Stephanov M A 2008 *J. High Energy Phys.* JHEP04(2008)100

[239] Romatschke P 2018 *Phys. Rev. Lett.* **120** 012301

[240] Kurkela A, Mazeliauskas A, Paquet J F, Schlichting S and Teaney D 2019 *Phys. Rev. Lett.* **122** 122302

[241] Petersen H, Coleman-Smith C, Bass S A and Wolpert R 2011 *J. Phys. G* **38** 045102

[242] Novak J, Novak K, Pratt S, Vredevoogd J, Coleman-Smith C and Wolpert R 2014 *Phys. Rev. C* **89** 034917

[243] Pratt S, Sangaline E, Sorensen P and Wang H 2015 *Phys. Rev. Lett.* **114** 202301

[244] Bernhard J E, Moreland J S and Bass S A 2019 *Nature Phys.* **15** 1113–7

[245] Moreland J S, Bernhard J E and Bass S A 2020 *Phys. Rev. C* **101** 024911

[246] Nijs G, van der Schee W, Gürsoy U and Snellings R 2021 *Phys. Rev. Lett.* **126** 202301

[247] Parkkila J E, Onnerstad A, Taghavi F, Mordasini C, Bilandzic A and Kim D J 2022 *Phys. Lett. B* **835** 137485 arXiv:2111.08145

[248] Nijs G and van der Schee W 2022 *Phys. Rev. C* **106** 044903 arXiv:2110.13153

[249] Cao S *et al* (JETSCAPE) 2021 *Phys. Rev. C* **104** 024905

[250] Xu Y, Bernhard J E, Bass S A, Nahrgang M and Cao S 2018 *Phys. Rev. C* **97** 014907

[251] Auvinen J, Bernhard J E, Bass S A and Karpenko I 2018 *Phys. Rev.* C **97** 044905
[252] Jeon S and Koch V 2004 Event by event fluctuations *Quark-Gluon Plasma* ed R C Hwa and X N Wang 3rd edn (Singapore: World Scientific) pp 430–90
[253] Asakawa M, Heinz U W and Müller B 2000 *Phys. Rev. Lett.* **85** 2072–5
[254] Haussler S, Scherer S and Bleicher M 2008 *Phys. Lett.* B **660** 197–201
[255] Pratt S 2012 *Phys. Rev. Lett.* **108** 212301
[256] Pratt S, McCormack W P and Ratti C 2015 *Phys. Rev.* C **92** 064905
[257] Ratti C 2018 *Rept. Prog. Phys.* **81** 084301
[258] Ratti C and Bellwied R 2021 The deconfinement transition of QCD: theory meets experiment *Lecture Notes in Physics* (Cham: Springer) vol. 981
[259] Shuryak E V and Stephanov M A 2001 *Phys. Rev.* C **63** 064903
[260] Skokov V, Friman B and Redlich K 2013 *Phys. Rev.* C **88** 034911
[261] Alba P, Bellwied R, Bluhm M, Mantovani Sarti V, Nahrgang M and Ratti C 2015 *Phys. Rev.* C **92** 064910
[262] Mogliacci S, Andersen J O, Strickland M, Su N and Vuorinen A 2013 *J. High Energy Phys.* JHEP12(2013)055
[263] Bellwied R, Borsanyi S, Fodor Z, Katz S D, Pasztor A, Ratti C and Szabo K K 2015 *Phys. Rev.* D **92** 114505
[264] Alba P, Alberico W, Bellwied R, Bluhm M, Mantovani Sarti V, Nahrgang M and Ratti C 2014 *Phys. Lett.* B **738** 305–10
[265] Berdnikov B and Rajagopal K 2000 *Phys. Rev.* D **61** 105017
[266] Vovchenko V, Anchishkin D V, Gorenstein M I and Poberezhnyuk R V 2015 *Phys. Rev.* C **92** 054901
[267] Stephanov M A 2011 *Phys. Rev. Lett.* **107** 052301
[268] Adam J *et al* (STAR) 2021 *Phys. Rev. Lett.* **126** 092301
[269] Abdallah M *et al* (STAR) 2021 *Phys. Rev.* C **104** 024902
[270] Adamczewski-Musch J *et al* (HADES) 2020 *Phys. Rev.* C **102** 024914
[271] Abdallah M S *et al* (STAR) 2022 *Phys. Rev. Lett.* **128** 202303 arXiv:2112.00240
[272] Bluhm M *et al* 2020 *Nucl. Phys.* A **1003** 122016
[273] Stephanov M and Yin Y 2018 *Phys. Rev.* D **98** 036006
[274] Adler S S *et al* (PHENIX) 2003 *Phys. Rev. Lett.* **91** 172301
[275] Adams J *et al* (STAR) 2004 *Phys. Rev. Lett.* **92** 052302
[276] Acharya S *et al* (ALICE) 2018 *J. High Energy Phys.* JHEP09(2018)006
[277] Fries R J, Müller B, Nonaka C and Bass S A 2003 *Phys. Rev. Lett.* **90** 202303
[278] Greco V, Ko C M and Levai P 2003 *Phys. Rev. Lett.* **90** 202302
[279] Fries R J, Müller B, Nonaka C and Bass S A 2003 *Phys. Rev.* C **68** 044902
[280] Molnar D and Voloshin S A 2003 *Phys. Rev. Lett.* **91** 092301
[281] Fries R J, Greco V and Sorensen P 2008 *Ann. Rev. Nucl. Part. Sci.* **58** 177–205
[282] Abdallah M *et al* (STAR) 2023 *Phys. Rev.* C **107** 024912
[283] Hwa R C and Yang C B 2004 *Phys. Rev.* C **70** 024905
[284] Werner K 2012 *Phys. Rev. Lett.* **109** 102301
[285] Werner K, Karpenko I, Bleicher M, Pierog T and Porteboeuf-Houssais S 2012 *Phys. Rev.* C **85** 064907
[286] Belikov I and (ALICE) 2011 *J. Phys.* G **38** 124078
[287] Lu Y, Bleicher M, Liu F, Liu Z, Petersen H, Sorensen P, Stoecker H, Xu N and Zhu X 2006 *J. Phys.* G **32** 1121–30
[288] Hirano T, Heinz U W, Kharzeev D, Lacey R and Nara Y 2008 *Phys. Rev.* C **77** 044909
[289] Ryu S, Staudenmaier J and Elfner H 2019 *MDPI Proc.* **10** 44
[290] Gyulassy M and Plumer M 1990 *Phys. Lett.* B **243** 432–8
[291] Wang X N and Gyulassy M 1992 *Phys. Rev. Lett.* **68** 1480–3
[292] Baier R, Dokshitzer Y L, Mueller A H, Peigne S and Schiff D 1997 *Nucl. Phys.* B **484** 265–82
[293] Zakharov B G 1996 *JETP Lett.* **63** 952–7
[294] Baier R, Schiff D and Zakharov B G 2000 *Ann. Rev. Nucl. Part. Sci.* **50** 37–69
[295] Caucal P, Iancu E and Soyez G 2019 *J. High Energy Phys.* JHEP10(2019)273
[296] Adler S S *et al* (PHENIX) 2006 *Phys. Rev. Lett.* **96** 202301
[297] Aad G *et al* (ATLAS) 2013 *Phys. Rev. Lett.* **110** 022301
[298] Chatrchyan S *et al* (CMS) 2015 *J. High Energy Phys.* JHEP03(2015)022
[299] Aad G *et al* (ATLAS) 2020 *Phys. Lett.* B **802** 135262

[300] Acharya S *et al* (ALICE) 2020 *J. High Energy Phys.* JHEP09(2020)076
[301] Aad G *et al* (ATLAS) 2019 *Eur. Phys. J.* C **79** 935
[302] Sirunyan A M *et al* (CMS) 2022 *Phys. Rev. Lett.* **128** 122301
[303] Adare A *et al* (PHENIX) 2012 *Phys. Rev. Lett.* **109** 152301 [Erratum: Phys. Rev. Lett. 125, 049901 (2020)]
[304] Abelev B B *et al* (ALICE) 2014 *Eur. Phys. J.* C **74** 3108
[305] Shibata M (PHENIX) 2022 *PoS* **249** PANIC2021
[306] Nishitani R (PHENIX) 2019 *MDPI Proc.* **10** 42
[307] Sirunyan A M *et al* (CMS) 2018 *Phys. Lett.* B **776** 195–216
[308] Aad G *et al* (ATLAS) 2013 *Phys. Lett.* B **719** 220–41
[309] Casalderrey-Solana J, Milhano J G and Wiedemann U A 2011 *J. Phys.* G **38** 035006
[310] Connors M, Nattrass C, Reed R and Salur S 2018 *Rev. Mod. Phys.* **90** 025005
[311] Chatrchyan S *et al* (CMS) 2014 *Phys. Rev.* C **90** 024908
[312] Mao Y (CMS) 2014 *Nucl. Phys.* A **932** 88–92
[313] ATLAS Collaboration 2023 arXiv:2303.10090
[314] Adamczyk L *et al* (STAR) 2013 *Phys. Rev.* C **87** 044903
[315] Ke W, Xu Y and Bass S A 2019 *Phys. Rev.* C **100** 064911
[316] ATLAS Collaboration 2023 *Phys. Rev.* C **107** 054908 arXiv:2205.00682
[317] Perepelitsa D V (ATLAS) 2016 *Nucl. Phys.* A **956** 653–6
[318] Andronic A *et al* 2016 *Eur. Phys. J.* C **76** 107
[319] Prino F and Rapp R 2016 *J. Phys.* G **43** 093002
[320] Karsch F and Satz H 1991 *Z. Phys.* C **51** 209–24
[321] Jakovac A, Petreczky P, Petrov K and Velytsky A 2007 *Phys. Rev.* D **75** 014506
[322] Petreczky P, Sharma S and Weber J H 2021 *Phys. Rev.* D **104** 054511
[323] Rothkopf A 2020 *Phys. Rept.* **858** 1–117
[324] Adler S S *et al* (PHENIX) 2006 *Phys. Rev. Lett.* **96** 032301
[325] Acharya U A *et al* (PHENIX) 2022 arXiv:2203.17058
[326] Acharya S *et al* (ALICE) 2021 *Phys. Lett.* B **820** 136558
[327] Abdallah M S *et al* (STAR) 2023 *Eur. Phys. J.* C **83** 455 [Erratum to: Eur. Phys. J. C 83, 455 (2023)]
[328] Acharya S *et al* (ALICE) 2022 *Nature* **605** 440–6
[329] Armesto N, Salgado C A and Wiedemann U A 2004 *Phys. Rev.* D **69** 114003
[330] Acharya S *et al* 2020 ALICE *Phys. Lett.* B **804** 135377
[331] Acharya S *et al* (ALICE) 2022 *J. High Energy Phys.* JHEP01(2022)174
[332] Acharya S *et al* (ALICE) 2021 *Phys. Lett.* B **813** 136054
[333] Acharya S *et al* (ALICE) 2022 *Phys. Lett.* B **827** 136986
[334] Matsui T and Satz H 1986 *Phys. Lett.* B **178** 416
[335] Hufner J, Klevansky S P and Rehberg P 1996 *Nucl. Phys.* A **606** 260–82
[336] Adare A *et al* (PHENIX) 2012 *Phys. Rev.* C **86** 064901
[337] Adam J *et al* (ALICE) 2017 *Phys. Lett.* B **766** 212–24
[338] Sirunyan A M *et al* (CMS) 2019 *Phys. Lett.* B **790** 270–93
[339] Abdallah M *et al* (STAR) 2022 *Phys. Lett.* B **825** 136865
[340] Collaboration A (ALICE) 2022 arXiv:2211.04384
[341] Chanfray G, Rapp R and Wambach J 1996 *Phys. Rev. Lett.* **76** 368–71
[342] Rapp R and Wambach J 2000 *Adv. Nucl. Phys.* **25** 1
[343] Brown G E and Rho M 2002 *Phys. Rept.* **363** 85–171
[344] Arnaldi R (NA60) *et al* 2006 *Phys. Rev. Lett.* **96** 162302
[345] Arnaldi R (NA60) *et al* 2008 *Phys. Rev. Lett.* **100** 022302
[346] Salabura P and Stroth J 2021 *Prog. Part. Nucl. Phys.* **120** 103869
[347] Adam J *et al* (STAR) 2018 arXiv:1810.10159
[348] van Hees H and Rapp R 2006 *Phys. Rev. Lett.* **97** 102301
[349] Rapp R 2013 *Adv. High Energy Phys.* **2013** 148253
[350] Endres S, van Hees H and Bleicher M 2016 *Phys. Rev.* C **94** 024912
[351] Endres S, van Hees H, Weil J and Bleicher M 2015 *Phys. Rev.* C **91** 054911
[352] Linnyk O, Bratkovskaya E L and Cassing W 2016 *Prog. Part. Nucl. Phys.* **87** 50–115
[353] Adamczewski-Musch J *et al* (HADES) 2019 *Nature Phys.* **15** 1040–5
[354] Staudenmaier J, Weil J, Steinberg V, Endres S and Petersen H 2018 *Phys. Rev.* C **98** 054908
[355] Afanasiev S *et al* (PHENIX) 2012 *Phys. Rev. Lett.* **109** 152302

[356] Adare A *et al* (PHENIX) 2010 *Phys. Rev. Lett.* **104** 132301
[357] Adare A *et al* (PHENIX) 2015 *Phys. Rev.* C **91** 064904
[358] Adare A *et al* (PHENIX) 2018 *Phys. Rev.* C **98** 054902
[359] Adam J *et al* (ALICE) 2016 *Phys. Lett.* B **754** 235–48
[360] Angelis A L S *et al* (CMOR) 1989 *Nucl. Phys.* B **327** 541–68
[361] Akesson T (Axial Field Spectrometer) *et al* 1990 *Sov. J. Nucl. Phys.* **51** 836–45
[362] Adare A *et al* (PHENIX) 2013 *Phys. Rev.* C **87** 054907
[363] Acharya U A *et al* (PHENIX) 2023 *Phys. Rev.* C **107** 024914 arXiv:2203.12354
[364] Aggarwal M M *et al* (WA98) 2000 *Phys. Rev. Lett.* **85** 3595–9
[365] Paquet J F, Shen C, Denicol G S, Luzum M, Schenke B, Jeon S and Gale C 2016 *Phys. Rev.* C **93** 044906
[366] Adam J *et al* (ALICE) 2016 *Phys. Lett.* B **754** 235–48
[367] Acharya S *et al* (ALICE) 2019 *Phys. Lett.* B **789** 308–22
[368] Gale C, Paquet J F, Schenke B and Shen C 2022 *Phys. Rev.* C **105** 014909
[369] Schäfer A, Garcia-Montero O, Paquet J F, Elfner H and Gale C 2022 *Phys. Rev.* C **105** 044910
[370] Steinheimer J, Gudima K, Botvina A, Mishustin I, Bleicher M and Stocker H 2012 *Phys. Lett.* B **714** 85–91
[371] Csernai L P and Kapusta J I 1986 *Phys. Rept.* **131** 223–318
[372] Scheibl R and Heinz U W 1999 *Phys. Rev.* C **59** 1585–602
[373] Vovchenko V, Gallmeister K, Schaffner-Bielich J and Greiner C 2020 *Phys. Lett.* B **800** 135131
[374] Neidig T, Gallmeister K, Greiner C, Bleicher M and Vovchenko V 2022 *Phys. Lett.* B **827** 136891
[375] Danielewicz P and Bertsch G F 1991 *Nucl. Phys.* A **533** 712–48
[376] Oliinychenko D, Pang L G, Elfner H and Koch V 2019 *Phys. Rev.* C **99** 044907
[377] Staudenmaier J, Oliinychenko D, Torres-Rincon J M and Elfner H 2021 *Phys. Rev.* C **104** 034908
[378] Dönigus B 2020 *Eur. Phys. J.* A **56** 280
[379] ALICE Collaboration 2022 arXiv:2209.07360
[380] Tolos L and Fabbietti L 2020 *Prog. Part. Nucl. Phys.* **112** 103770
[381] Fabbietti L, Mantovani Sarti V and Vazquez Doce O 2021 *Ann. Rev. Nucl. Part. Sci.* **71** 377–402
[382] Becattini F and Lisa M A 2020 *Ann. Rev. Nucl. Part. Sci.* **70** 395–423
[383] Zyla P A (Particle Data Group) *et al* 2020 *PTEP* **2020** 083C01
[384] Adamczyk L *et al* (STAR) 2017 *Nature* **548** 62–5
[385] Abdallah M S *et al* (STAR) 2021 *Phys. Rev.* C **104** L061901
[386] Acharya S *et al* (ALICE) 2020 *Phys. Rev.* C **101** 044611
[387] Okubo K (STAR) 2022 *EPJ Web Conf.* **259** 06003
[388] Adam J *et al* (STAR) 2021 *Phys. Rev. Lett.* **126** 162301
[389] Adam J *et al* (STAR) 2018 *Phys. Rev.* C **98** 014910
[390] Müller B and Schäfer A 2018 *Phys. Rev.* D **98** 071902
[391] Adam J *et al* (STAR) 2019 *Phys. Rev. Lett.* **123** 132301
[392] Becattini F 2022 *Rept. Prog. Phys.* **85** 122301 arXiv:2204.01144
[393] Singha S (STAR) 2021 *Nucl. Phys.* A **1005** 121733
[394] Mohanty B (ALICE) 2021 *PoS* **ICHEP2020** 555
[395] Florkowski W, Kumar A and Ryblewski R 2019 *Prog. Part. Nucl. Phys.* **108** 103709
[396] Weickgenannt N, Sheng X L, Speranza E, Wang Q and Rischke D H 2019 *Phys. Rev.* D **100** 056018
[397] Gallegos A D, Gürsoy U and Yarom A 2021 *SciPost Phys.* **11** 041
[398] Peccei R D and Quinn H R 1977 *Phys. Rev. Lett.* **38** 1440–3
[399] Kharzeev D E, McLerran L D and Warringa H J 2008 *Nucl. Phys.* A **803** 227–53
[400] Newman G M 2006 *J. High Energy Phys.* JHEP01(2006)158
[401] Son D T and Surowka P 2009 *Phys. Rev. Lett.* **103** 191601
[402] Kharzeev D E and Yee H U 2011 *Phys. Rev.* D **83** 085007
[403] Kharzeev D E, Liao J, Voloshin S A and Wang G 2016 *Prog. Part. Nucl. Phys.* **88** 1–28
[404] Voloshin S A 2018 *Phys. Rev.* C **98** 054911
[405] Shi S, Jiang Y, Lilleskov E and Liao J 2018 *Ann. Phys.* **394** 50–72
[406] Müller B and Schäfer A 2010 *Phys. Rev.* C **82** 057902
[407] Abdallah M *et al* 2022 STAR *Phys. Rev. Lett.* **128** 092301
[408] Acharya S *et al* (ALICE) 2020 *J. High Energy Phys.* JHEP09(2020)160

[409] Sirunyan A M *et al* (CMS) 2018 *Phys. Rev.* C **97** 044912
[410] Adam J *et al* (STAR) 2021 *Nucl. Sci. Tech.* **32** 48
[411] Abdallah M *et al* (STAR) 2022 *Phys. Rev.* C **105** 014901
[412] Feng Y (STAR) 2022 Estimate of a nonflow baseline for the chiral magnetic effect in isobar collisions at RHIC *20th Int. Conf. on Strangeness in Quark Matter 2022*
[413] Adare A *et al* (PHENIX) 2015 arXiv:1501.06197
[414] Adamová D *et al* 2019 arXiv:1902.01211
[415] Accardi A *et al* 2016 *Eur. Phys. J.* A **52** 268
[416] Brodsky S J, Fleuret F, Hadjidakis C and Lansberg J P 2013 *Phys. Rept.* **522** 239–55
[417] Aaij R (LHCb) *et al* 2019 *Phys. Rev. Lett.* **122** 132002
[418] Franzoso E (LHCb) 2022 *EPJ Web Conf.* **259** 13010
[419] Steinberg P A (ALICE, ATLAS, CMS, LHCb, STAR) 2021 *Nucl. Phys.* A **1005** 122007
[420] Agakishiev G *et al* (HADES) 2009 *Eur. Phys. J.* A **41** 243–77
[421] 2011 The CBM physics book: Compressed baryonic matter in laboratory experiments *Lecture Notes in Physics* Friman B *et al* (ed) (Berlin: Springer) vol 814 978-3-642-13293-3
[422] Hachiya T 2020 *Int. J. Mod. Phys.* E **29** 2040005
[423] Xiaohong Z 2018 *Nucl. Phys. Rev.* **35** 339–49
[424] Kekelidze V D (NICA) 2017 *JINST* **12** C06012