Heavy-ion physics

M. Nguyen
Laboratoire Leprince-Ringuet (CNRS), Palaiseau, France

Abstract
These proceedings cover two lectures on heavy-ion physics given at the 2018 edition of the Asia-Europe-Pacific School of High-Energy Physics. After a brief introduction of basic concepts that are relevant for heavy-ion collisions, I discuss several interesting phenomena that have been observed, and mention currently still open questions.

Keywords
Lectures; heavy ion collisions; quark-gluon plasma; jet quenching; strangeness enhancement; collective flow.

1 Introduction

These lectures are an introduction to the physics of ultra-relativistic heavy-ion collisions and the quark-gluon plasma (QGP). In the first lecture, I start with a brief review of some salient features of quantum chromodynamics (QCD), which is the underlying theory that describes the dynamics of the QGP, the calculational difficulties of that theory notwithstanding. I then discuss properties of the QGP as seen through the lens of thermodynamics, which turns out to be an appropriate language with which to describe heavy-ion collisions, due to the typically large number of individual particles in such collisions. I then trace out the space-time evolution of a heavy-ion collision focusing first on collision centrality, a key collision property employed in nearly all heavy-ion measurements. Then I look at a simple picture attributed to J.D. Bjorken, who estimated the energy density of the QGP from mid-rapidity particle yields. I then turn to hadro-chemistry, first looking at the statistical hadroproduction model, which connects the abundance of the various particle species to the properties of the QCD matter prior to hadronization. Then I focus on the particular role of strangeness in this type of model, and what it tells us about chemical equilibrium.

In the second lecture I focus on the various phenomena that we use to infer the properties of the QGP. There are many, and I have not attempted a complete survey. Rather I focused on three types of measurements, which have been of particular interest of late, namely those related to collective flow, jet quenching and quarkonia dissociation. Finally, I close the second set of lectures with a discussion of "small systems", mainly referring to proton-proton and proton-ion collisions, which surprisingly exhibit certain effects previously only expected in ion-ion collisions.

2 Generalities

Let us begin by defining the term "heavy-ion collision". Ions are of course nuclei that have been stripped of their electrons, which is necessary to accelerate them to large energies. The Large Hadron Collider typically collides the lightest possible ion, the hydrogen ion, also known as the proton. By heavy ions, we mean nuclei with O(100) nucleons. At the Relativistic Heavy-Ion Collider (RHIC) at Brookhaven, the most commonly collided ion is gold (Au), although quite a few other species, the heaviest being the oblong uranium (U) ion, have also been collided. The Au isotope which is collided (the only stable one) consists of 197 nucleons, 79 protons and 118 neutrons. At the Large Hadron Collider the most frequently employed ion is Pb(208), with 82 protons and 126 neutrons. Due to the difficulty of isolating this particular isotope with high purity, the Pb source is perhaps surprisingly *more* expensive than Au.

 doi:10.23730/CYRSP-2020-005.129, ISSN 0531-4283.

The center of mass energy of the highest energy ion collisions at the LHC is about 5 TeV per nucleon pair. That corresponds to roughly a PeV of total energy! Not all of this energy is however available for particle production. This energy will be estimated in Bjorken's picture in Section 5.

The primary aim of colliding heavy ions is to observe matter in the deconfined phase, that of the quark-gluon plasma. Such a phase transition was predicted not long after the discovery of asymptotic freedom, first at large density [1], for example at the core of compact astrophysical objects, and then at large temperature [2], as in the first few microseconds of the universe. Although finite density calculations are complicated by the so-called sign problem, the temperature of the phase transition at vanishing net baryon density is amenable to calculation in lattice QCD. State-of-the-art calculations give a critical temperature of around 160 MeV [3, 4], which corresponds to about 10^{12} K! As it turns out, this temperature coincides well with an estimate for a "limiting temperature" of hadronic matter due to Hagedorn, that pre-dates the advent of QCD [5]. Hagedorn's estimate was based on the exponential growth in the number of bound states with increasing energy, and is intimately related to the statistical model of hadroproduction that will be discussed in Section 6. We will see in Section 5 that the energy delivered by heavy-ion collisions is sufficient to reach the critical temperature.

3 Brief reminders of QCD

The Lagrangian of QCD is similar in its form to that of QED, but with different constants related to the SU(3) structure, which are defined by the commutation relations of the generators of the group. The non-trivial part of the theory, which limits the applicability of perturbative calculations to hard processes, is the self-coupling of the gauge field, in other words the interaction between gluons. The particle content of the theory is given by the gluons, which come in 8 different linearly independent color combinations, and by the quarks, which come in 3 colors, but also in 6 different flavors, which have very different masses. Most of the matter in the nucleon has nothing to do with the bare quark mass, but is rather generated dynamically by virtual quarks and gluons that buzz around the valence quarks. The bare quark masses imparted by the Higgs mechanisms are basically free parameters as far as QCD is concerned. These masses will turn out to be important for the discussion of the QGP. The light quarks, up and down, have masses on the order of a few MeV, that are small compared to the temperature of the QGP, such that the energy can easily be converted to creation of these particles. The strange quark mass of around 100 MeV is of the same order as the transition temperature. The heavy quarks, charm and bottom, are interesting precisely because they cannot be created easily by the plasma. In Section 10 I will discuss how bound states of these quarks are modified by the QGP. Top quarks are so heavy that they decay before the QGP is formed, so I will not discuss them in these lectures.

The reason that the gluon self-interaction is important has to do with the effect it has on the QCD coupling constant. In QED virtual particles screen the (electric) charge, such that the coupling grows weaker with distance. In QCD the self-interaction of gluons causes the opposite effect, such that the coupling increases with distance. The small coupling at short distance scales is what allows us to probe the quarks inside the nucleons with deep inelastic scattering experiments. This is also what allows us to solve QCD using a perturbative expansion. But this linearly increasing potential, rather than one which goes to zero in QED, is what keeps quarks confined inside hadrons. It is also what makes the low energy behavior of QCD difficult to calculate analytically. In this regime, one has to resort to lattice QCD calculations, which we'll come back to in a moment.

4 QCD thermodynamics

First we can get some guidance on what to expect from QCD matter via thermodynamics. The idea of a QGP emerges directly from the concept of asymptotic freedom. In the limit of large temperature, the coupling constant approaches zero, so one ends up with a weakly interacting gas of quarks and gluons. That is not to say that the QGP we create in the laboratory satisfies this condition. As we'll see, this is

quite far from what actually happens. Rather the QGP just above the confinement temperature is strongly coupled. However, it's still instructive to consider this limit, where we can derive the equation of state directly from kinetic theory.

The equation of state is usually expressed as the relationship between the pressure and the temperature. Equivalently we may express the E.O.S. in terms of the energy density and temperature. From kinetic theory, this is simply given by the density of states multiplied by their energy:

$$\epsilon(T) = \int \frac{d^3p}{(2\pi)^3} \sum_i \frac{E_i}{e^{E_i/T} \pm 1} \quad . \tag{1}$$

Recall that the plus-minus takes into account quantum effects. It would be zero for a classical system, i.e., Maxwell-Boltzmann statistics. The plus corresponds to fermions, i.e., Fermi-Dirac statistics, and the minus to bosons, i.e., Bose-Einstein statistics. The sum is over the different species of particles. Performing the integration (by parts), one obtains

$$\epsilon(T) = \frac{\pi}{30} N T^4 \quad . \tag{2}$$

If one performs the integration, one finds that the energy density is proportional to T^4, the familiar dependence of blackbody radiation. N is the number of different states. For matter below the critical temperature, one may approximate matter to be a pion gas, as pions are the lightest and therefore most abundant particle. The pions form an isospin triplet, so N=3. For the QGP, one obtains the following result, which is identical save for a slightly different numerical pre-factor for fermions:

$$\epsilon(T) = \frac{\pi}{30} \left(N_B + \frac{7}{8} N_F \right) T^4 \quad . \tag{3}$$

The number of degrees of freedom is much larger for the case of the QGP, however. For the bosons one has 8 different linearly independent color combinations. As a massless spin-1 boson, the gluon also has two allowable spin states (compare to 3 spin states for the massive pion). As fermions, the quarks each contribute two spin states. Another factor of two comes from the anti-quarks, as well as a factor of 3 for the three colors of the SU(3) group. Finally we have flavor. Here we can consider that 3 colors are available to be created, up, down and strange. In practice, strange production is somewhat suppressed by its mass, which we've neglected, but it doesn't change the result dramatically. In total, instead of 3, we find 47.5 for the number of degrees of freedom, which is about a factor 15 jump in the energy density at T_c.

Now we can compare what we get from kinetic theory for an ideal gas to what we get from the lattice QCD. The equation of state from the lattice is shown in Fig. 1. Indeed one finds a large jump in the ratio of ϵ/T^4, indicative of a jump in the number of degrees of freedom. Inserting a critical temperature of 160 MeV into equation 3, one obtains an energy density of around 1 GeV/fm^3[1].

Although the height of the plateau above T_c in Fig. 1 depends a bit on how exactly one treats the strange quark, one observes a jump all the same. It's reasonable to ask, however, whether the ideal gas E.O.S. is actually a good description of the QGP we observe in the lab. Note that although this jump is a large factor, it does not quite saturate the Stefan-Boltzmann limit, indicated by arrows on the figure. To get some intuition, it's useful to look at N=4 Super Yang Mills, which is a 10 dimensional theory that has nothing to do with QCD. But this theory is a dual to a conformal field theory, which is a particular kind of quantum field theory that is scale invariant. Such a theory is very much the opposite of an ideal gas, as it has infinitely strong coupling. So even though this is not QCD, one may reasonably expect that the details of the theory aren't important in the strong coupling limit. And so it turns out that the scaled

[1] Keep in mind the handy conversion factor for natural units: $\hbar c = 0.197$ GeV fm

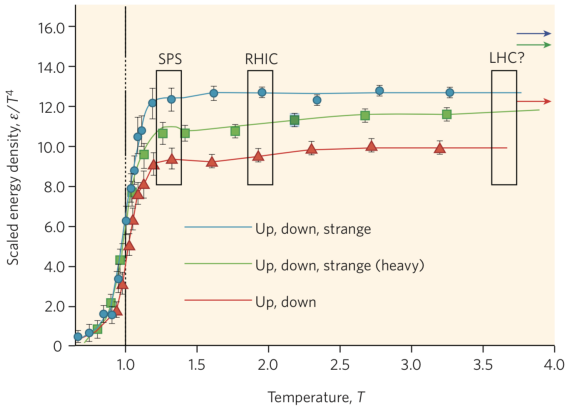

Fig. 1: A lattice QCD calculation of the equation of state of matter in terms of the scaled energy density, as a function of temperature (normalized by the critical temperature for a phase transition) [6, 7]. The lattice data demonstrate a precipitous jump near T_c, as expected for a phase transition from hadronic to partonic degrees of freedom.

energy density is rather close to what one gets in such a theory. So this is by no means a proof that the QGP is strongly coupled, however it gives a hint that this departure from the weak-coupling limit could be significant. We will see evidence of the strongly coupled behavior of the plasma in Section 8.

5 Space-time evolution

So far we've been discussing the quark-gluon plasma as if it's a static object, when in fact it's a very transient state. Before the collision, if we are standing in the lab frame, the two nuclei appear as flattened pancakes. These pancakes pass through one another, and the particles start to scatter. The time taken for the scattering to occur can be roughly estimated from the uncertainty principle, as one over the momentum of the particles. This implies that by a time of 1 fm, the bulk of particles have been created and can rescatter with each other. We'll see later from measurements of collectivity that we do in fact have evidence that thermalization is achieved over such a short timescale. Then the QGP expands and cools until about 10 fm. At this point the temperature falls below the critical one and hadronization occurs. This is referred to as chemical freeze-out as the abundances of different particle species are fixed, barring long-lived decays. There's a short period of additional rescattering before the particles stop interacting and free stream to our detectors. So what we end up detecting is a mess of thousands of final state hadrons. What I'll try to show you is how we can learn about the QGP, which lasts shorter than 10 fm, which is something like 3×10^{-23} seconds!

In addition to the final state particles, we also need to consider the role of the initial state of the colliding nuclei. In particular, it's important to consider the special role that geometry plays in heavy-ion collisions. When collisions have a relatively small impact parameter, the overlap region is going to be relatively large. This is referred to as a central event. In more glancing or "peripheral" collisions, one expects nuclear effects to be smaller, eventually looking very much like proton collisions in the limit of two nucleon scattering. Nucleons that scatter are referred to as participants, while others are spectators. Generally in a collider experiment one cannot directly measure how many nucleons participate in the collision, event-by-event. So if the energy density of the matter depends crucially on impact parameter, i.e., collision "centrality", how can we access this in our experiments?

The approach used is to make use of a so-called Glauber model, named after Nobel prize winner Roy Glauber, who passed away since these lectures were given. The model is based on work he did in the late 1950's [8], although this description follows a modern review [9]. The basic assumptions are: 1) the probability of scattering is independent of the scatterings which came before, and 2) the struck

nucleons continue along straight line trajectories. The starting point is to take the nuclear charge density, which is well-known from all sorts of low energy scattering experiments. This is typically parameterized by a Woods-Saxon distribution, which for spherical nuclei is expressed as

$$\rho(r) = \frac{\rho_0}{1 + \exp[(r - R)/a]} \quad , \tag{4}$$

where ρ_0 is the density at the core. However, we're only interested in how the density falls off as a function of the radius (r), such that we can look at the ratio ρ/ρ_0. As shown for Pb(208) in Fig. 2 (left), this density is essentially flat until you get close to the surface, where it falls off exponentially. The parameter R corresponds to the nuclear radius, while the surface thickness a describes how rapidly the nuclear density falls off. For the example of Pb(208), R = 6.7 fm and a is about 0.5 fm. The nuclear thickness of nucleus A is defined by the following integral over the nuclear density, $T_{\rm A}(\vec{s}) = \int \rho_A(\vec{s}, z)\mathrm{d}z$, as shown n Fig 2 (right). T is normalized such that the integral over the entire nucleus gives back the number of nucleons: $\int T_{\rm A}(\vec{s})\mathrm{d}^2s = A$. The thickness function is then the product of the two individual nuclear thicknesses integrated over the impact parameter of the collision: $T_{AB}(b) = \int T_{\rm A}(\vec{s}) \cdot T_{\rm B}(\vec{s} - \vec{b})\mathrm{d}^2s$.

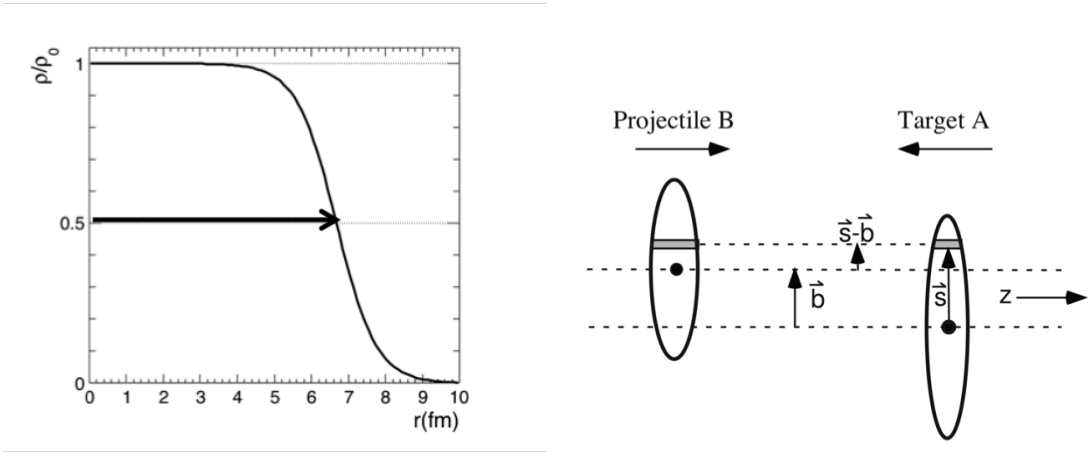

Fig. 2: Left: Woods-Saxon profile for Pb(208). Right: Diagram of nuclear scattering in the Glauber model [9].

There are two numbers that are key for describing heavy-ion collisions. One is the mean number of inelastic collisions for a given impact parameter. Elastic collisions don't contribute much to the mid-rapidity energy density, so they're not taken into account. The total number of binary inelastic collisions in a heavy-ion collision is simply the nuclear thickness function times the nucleon-nucleon cross section: $N_{\rm coll} = T_{\rm AB}(b) \cdot \sigma^{\rm nn}_{\rm inel}$. The other important number is the number of participating nucleons, sometimes called the number of wounded nucleons. This is the number of nucleons that are struck at least once. Given a nucleon in nucleus A, the probability for it to scatter with a nucleon from B is $p_{\rm int} = T_{\rm B}(\vec{s} - \vec{b}) \cdot \sigma^{\rm nn}_{\rm inel}/B$. The probability for it not to interact with any nucleon in B is then just $(1 - p_{\rm int})^B$. The probability that at least one nucleon interacts is $1 - (1 - p_{\rm int})$. To get the total number of participants in A, we integrate this probability over all nucleons in A, taking into account the thickness: $N^{\rm A}_{\rm part} = \int T_{\rm A}(\vec{s}) \cdot \left(1 - [1 - T_B(\vec{s} - \vec{b}] \cdot \sigma^{\rm nn}_{\rm inel}/B]^B\right) \mathrm{d}^2s$. The total number of participants is then the sum over the two nuclei: $N_{\rm part}(b) = N^{\rm A}_{\rm part}(b) + N^{\rm B}_{\rm part}(b)$.

The Glauber model gives us a connection between the geometry, i.e., the impact parameter, on one hand, and which nucleons are struck on the other. Still missing is a connection to what we measure in our detectors. Centrality is defined based on some observable, usually multiplicity in some detector, that has a monotonic dependence on the Glauber quantities, most importantly on $N_{\rm part}$. To make the connection, one has to model the multiplicity distribution for a nn scattering. We know from pp and $p\bar{p}$

experiments that the multiplicity distribution is well described by a negative binomial distribution. One can randomly sample this distribution $N_{\rm part}$ times to get the total multiplicity. One then divides the total cross section in equal bins or percentiles. The top 5% highest multiplicity events become centrality class 0-5% for example. One then quotes Glauber quantities, typically $\langle N_{\rm part}\rangle$, for a given centrality class, which has the advantage of being independent of one's detector, to good approximation.

Although centrality determination varies by experiment, it's instructive to look at an example. Figure 3 shows how ALICE determines their centrality from the multiplicity measured in a forward detector called the VZERO [10]. Using forward detectors is rather typical, as one wants to avoid correlations between the centrality determination and mid-rapidity measurements. Their Glauber calculation is shown in red, which is actually a fit to the data. They assume that particle production scales like a linear sum of a component proportional to $N_{\rm part}$ and one proportional to $N_{\rm coll}$. The $N_{\rm part}$ dominates, as expected if the multiplicity is dominated by particle production. To model the multiplicity this linear combination of $N_{\rm part}$ and $N_{\rm coll}$ is convoluted with a negative binomial distribution where the mean and variance are also free parameters in the fit. The fit turns out to give an excellent description of the data, as shown in the figure, giving confidence that the extracted Glauber parameters are meaningful.

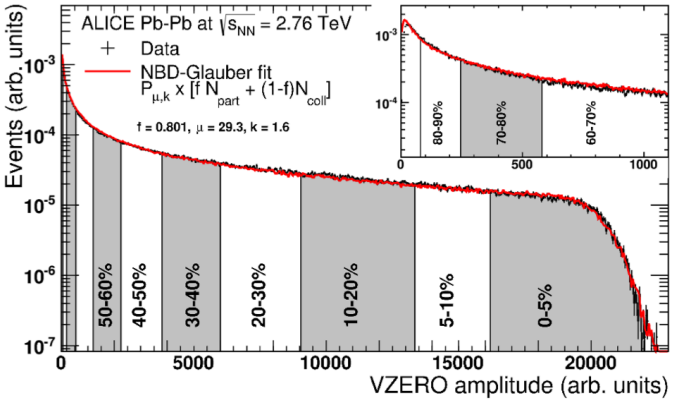

Fig. 3: Centrality determination from the VZERO amplitude in ALICE [10]. The red line is a fit to a linear combination of $N_{\rm part}$ and $N_{\rm coll}$, which is smeared by a negative binomial distribution, as described in the text.

We've been discussing multiplicity as a way to characterize the initial state, but multiplicity can also be used to estimate how much energy is deposited in a heavy-ion collision. Recall that although the total collision energy is on the order of a PeV, a lot of this energy will just continue with the remnants of the nuclei down the beampipe. On the other hand, we're interested to know how much of this energy is made available for particle production. The multiplicities involved are quite large, around 1.9k particles per unit η at mid-rapidity in the 5% most central collisions at top LHC energy [11]. The distribution in η has a characteristic shape with a plateau region at mid-rapidity. This shape can be understood in a simple space-time picture, which is commonly used to estimate the energy density.

That picture is due to J.D. Bjorken [12], more famous for his work on deep inelastic scattering. He considered the nuclei to be infinitely thin pancakes moving at the speed of light. The beam remnants continue moving forward, but mid-rapidity is dominated by produced particles. He considered the produced matter to expand homogeneously in the longitudinal direction, until some formation time t_0 in the particle rest frame. At that point the particles will materialize and free stream towards your detector. The velocity of a particle in this system is then simply z/t, or $z\gamma/\tau$. Considering an infinitely thin slice at $z = 0$ and time = τ_0, the energy density is then just the area of the cylinder (A), so the size of the nucleus basically, times dE/dz (evaluated at $z = 0$):

$$\epsilon = \frac{E}{V} = \frac{1}{A}\frac{dE}{dz}\bigg|_{z=0} = \frac{1}{A}\frac{dy}{dz}\bigg|_{z=0}\frac{dE}{dy}\bigg|_{y=0} = \frac{1}{A\cdot\tau_0}\frac{dE}{dy}\bigg|_{y=0} \quad . \tag{5}$$

Here we have used the chain rule to get the energy density in terms of rapidity, as well as using a useful relation for rapidity: $\sinh y = \beta/\gamma = z/\tau$. This equation now relates the actual energy density in the produced matter to a quantity we can observe in our detector. Inserting some typical numbers, A is about 7 fm for the Pb nucleus. τ_0 is usually assumed to be on the order of 1 fm based on the uncertainty principle, although we'll see from flow measurements in Section 8 that such a short formation time is justified. The energy we measure at mid-rapidity can be up to 2 GeV per unit η in central collisions at the LHC. That gives an energy density of around 13 GeV/fm^3 which is well in excess of the energy density of 1 GeV/fm^3 predicated for a phase transition from the lattice.

6 Statistical hadroproduction

In addition to looking at the overall number of particles produced, we can look at what types of particles are produced. What actually determines this? For elementary collisions, we don't rely on a first principles theory. Instead, hadronization is described by fragmentation functions which are universal, but non-perturbative. There are various phenomenological models to describe hadronization such as the string model in the Pythia generator [13] and the cluster model in the Herwig generator [14]. In heavy ions we may however choose to think in terms of an ensemble, where the language of statistical mechanics may be suitable. In fact that was what Fermi had tried to do back in the 50's before QCD [15].

We postulate that the QGP turns into an equilibrated gas of hadrons. Equilibrium is assumed to be established over a large volume, such that particle numbers are conserved only over the entire volume, not locally. In the language of statistical mechanics this is a grand canonical ensemble. So, for example, the number density of each state (n_i) will be given by a derivative of the partition function as follows,

$$n_i = N/V = -\frac{T}{V}\frac{\partial \ln Z_i}{\partial \mu} = \frac{g_i}{2\pi^2}\int_0^\infty \frac{p^2 \mathrm{d}p}{e^{(E_i-\mu_i)/T}\pm 1} \quad . \tag{6}$$

The abundance of each state depends on its mass and the temperature. Each conserved quantum number gives you a chemical potential (μ) such that you can conserve charge, baryon number, etc. The key point is that, if this holds, you can determine the temperature via a fit to the abundance of different states. This temperature corresponds not to the critical temperature, but to the temperature at chemical freezeout, when the identity of particles is fixed.

Figure 4 (left) shows an example of such a fit in central PbPb collisions. The ratio of data to the statistical hadronization model is plotted for a large variety of hadron species. The free parameters are the freeezeout temperature, the baryon chemical potential and the volume. The model describes the data remarkably well, given the small number of free parameters, which are the following. The baryon chemical potential extracted from the fit is consistent with zero, which is what we expect at such large energy. The volume is 5000 fm^3, which corresponds to a radius of about 10 fm. The extracted temperature of freezeout is 156 MeV, so below the critical temperature, but only slightly. This means the QGP hadronizes and the particle abundances are fixed almost immediately afterwards.

The fact that the fit works so well supports a picture in which we have chemical equilibrium over a large volume. This implies that the system is thermalized. What's surprising then is that such a fit actually works pretty well in elementary collisions, as well. This is shown in Fig. 4 (right), where the multiplicity of various hadron species in data is plotted against the results from a statistical model. While an excellent fit is obtained, crucially, one has to use a canonical ensemble, rather than a grand canonical one, such that quantum numbers are conserved only locally. One also needs to add an additional chemical potential for strange particles, which we'll discuss in the next section. These differences notwithstanding, there is

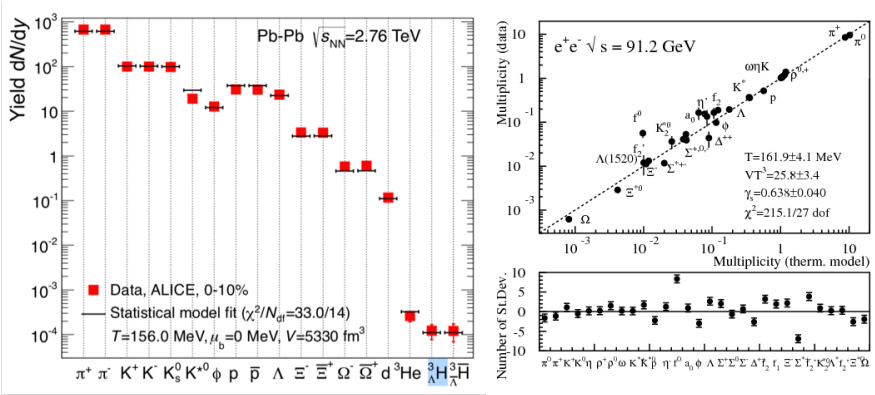

Fig. 4: Left: Data-to-model ratio for the yields of various species of hadrons in central PbPb collisions [16]. Right: Multiplicity of various hadron species in data, compared to a statistical model in e^+e^- collisions. [17]

quite a bit of debate as to why such a description works in elementary collisions, where thermalization is not expected to be reached [18].

7 Strangeness enhancement

The strange quark is indeed rather peculiar with respect to the QGP. For the light quarks, the hadron masses are almost entirely generated dynamically by the confinement mechanism, whereas for the heavy ones, the mass comes from the Higgs mechanism. The strange quark kind of splits the difference, in that bare mass isn't entirely negligible. It is however smaller than the temperature of the QGP, such that we expect thermal production, according to the statistical model. However, if we only have local conservation then the created strange quark pairs have a high probability to annihilate shortly after their creation. On the other hand if we have a situation which may be described as grand canonical ensemble, where we have a large volume of equilibrated matter, the strange quarks can wander off and hadronize with other quarks. That's indeed exactly what we see in the data. If we look at the ratio of strangeness production compared to light quark production, for example by looking at the ratio of kaons to pions, we see that there's an enhancement of a factor 2, which reaches a plateau at large collision energy. This enhancement factor is even larger for particles that contain more than on strange quark, up to about a factor of 15 for the triply strange Ω baryon. This supports a picture of the QGP as an equilibrated state of matter.

8 Collective flow

We now turn to the 2nd part of the lectures, which is more focused on observables that measure properties of the QGP, starting with collectivity. In non-central collisions the overlap zone of a heavy-ion collisions has an ellipsoidal shape. Given rapid thermalization, this leads to larger pressure along the short axis compared to the long one. This *spatial asymmetry* hence gives rise to a *momentum space anisotropy*. This effect is typically quantified via a Fourier decomposition of the particle yield vs. azimuth as

$$E\frac{d^3\sigma}{d\mathbf{p}^3} = \frac{1}{2\pi}\frac{d^2\sigma}{p_{\mathrm{T}}dp_{\mathrm{T}}dy}\left(1 + 2\sum_n v_n \cos\left[n(\phi - \psi_R)\right]\right) \quad , \tag{7}$$

where v_n are the flow coefficients, and ψ_R is the reaction plane angle, defined by the short axis of the ellipsoid. Given the ellipsoidal overlap shape, the 2nd harmonic v_2 is expected to dominate, the so-called elliptic flow term. Naively odd terms are not expected to contribute, due to the symmetry of the overlap shape.

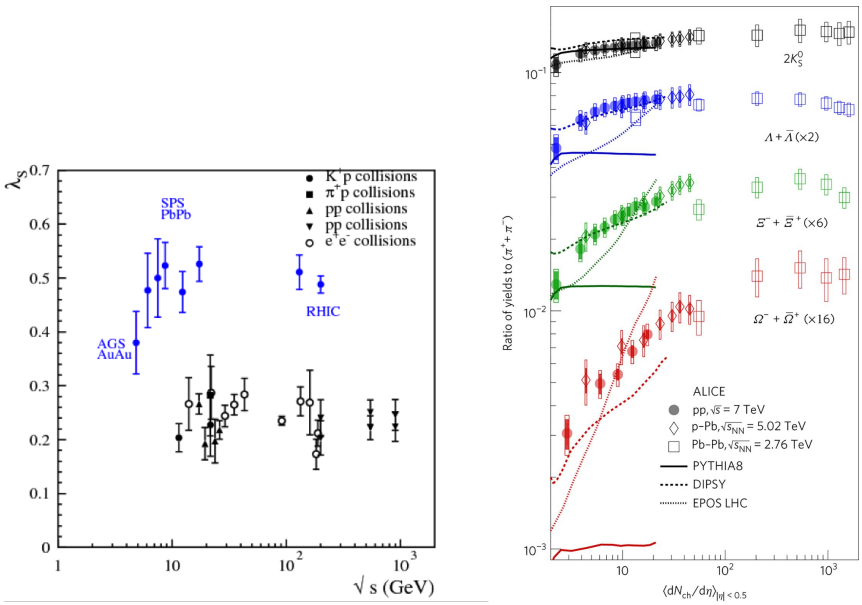

Fig. 5: Left: The ratio of strange quarks to light quarks vs. collision energy for various collision species [19]. Right: The ratio of yields of various strange hadrons to charged pions vs. charged particle multiplicity for pp, pPb and PbPb collisions [20].

Collective flow is modeled using relativistic hydrodynamics. Corrections to ideal hydrodynamics are encoded in the viscous correction terms. The shear viscosity term is particularly relevant as it serves to damp the elliptic flow. Hydrodynamics calculations including viscous terms [21] are compared to the data from STAR [22] in Fig. 6 (left). The elliptic flow term (scaled by the eccentricity of the overlap zone) is plotted as a function of the multiplicity density and compared to hydro calculations for various value of the dimensionless ratio shear viscosity divided by entropy density (η/s). The result comes remarkably close to the value of $1/4\pi$, calculated for conformal field theory, via the AdS/CFT correspondence. Although such a theory does not necessarily represent QCD, as an infinitely strongly coupled theory, it has been conjectured to give a universal lower bound for η/s. A comparison of the fluid imperfection index, η/s scaled by the value of $1/4\pi$, for several low viscosity materials in very different temperature regimes is shown in Fig. 6 (right). Albeit with a sizable systematic uncertainty, the QGP shows the lowest η/s of any known material, leading it to be dubbed the "perfect fluid".

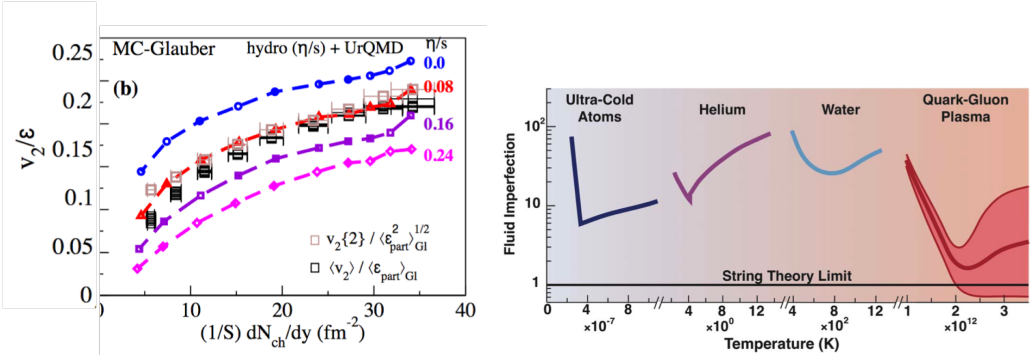

Fig. 6: Left: A comparison of STAR v_2 data [22] to hydrodynamic calculations with different values of η/s [21]. Right: The fluid imperfection index $4\pi\eta/s$ as a function of temperature for various low viscosity material in the vicinity of their respective phase transitions.

Whereas a large v_2 coefficient was expected, higher orders terms were not expected to be large.

In particular the odd terms were expected to be vanishing owing to the symmetry of the overlap zone. This turned out not to be the case, as shown in Fig. 7, where the various v_n coefficients are plotted vs. centrality [23]. The v_3 and v_4 terms were not only found to be non-zero, but have a markedly different centrality dependence than v_2, which increases as collisions become less central. These higher order terms were eventually understood to be due to event-by-event fluctuations of the overlap shape [24], which are not accounted for in the analytic Glauber approach described in Section 5.

These higher order terms turn out to be quite useful, which we can understand by analogy with the cosmic microwave background. Variation in the temperature of the CMB reflects the density fluctuations of the early universe before the inflationary period [25]. In much the same way, the azimuthal anisotropy in long range correlations in central heavy ions collisions reflect fluctuations of the initial state of dense QCD matter before its hydrodynamic expansion. Figure 7 (right) [26] shows the v_n coefficients for a selection of the 0.2% most central PbPb events, where all azimuthal flow should be dominated by initial state fluctuations, as the initial overlap zone is symmetric. The data are compared to hydrodynamic calculations, but with two different initial states: the standard Glauber ones, and MC-KLN, a model that additionally takes into account color field fluctuation at the sub-nucleon scale [27]. While such hydrodynamic calculations provide an accurate description of flow measurements overall, there is quite some sensitivity to how the initial state is modeled, which is demonstrated by the difference between these two calculations. That in turn effects the estimation of the viscosity, and in fact our knowledge of the initial state is at present the limiting factor in how well we can extract this quantity.

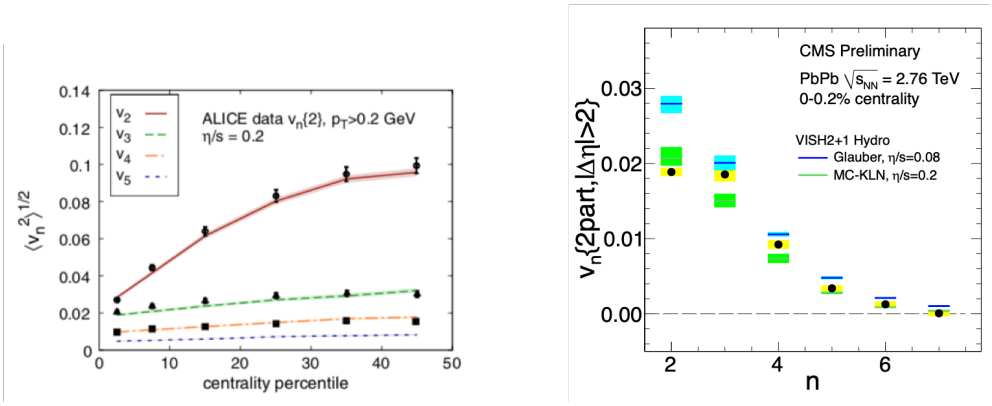

Fig. 7: Left: Azimuthal flow harmonics of different orders vs. centrality for PbPb collisions at 2.76 TeV from ALICE [23]. Right: Flow coefficients in ultra-central PbPb collisions from CMS [26], compared to hydrodynamic calculations with two different models of initial state fluctuations, as described in the text.

9 Jet quenching

The observables we have discussed so far, such as azimuthal flow and strangeness production, are related to bulk particle production. We now turn to hard probes, i.e., processes with large momentum exchange that occur before the QGP is formed. Hard scattering, whether in elementary collisions or in heavy ions, leads to production of jets, a collimated spray of particles resulting from the hadronization of an energetic parton. In heavy ions, however, one expects these jets to be "quenched", an idea first proposed in the early 1980's by Bjorken [28]. In much that same way as a charged particle loses energy as it traverses normal matter, a particle with color charge, i.e., a parton, should lose energy as it passes through the colored medium of the QGP. If one can describe the energy loss process theoretically, then one can in turn extract properties of the medium itself, such as the color charge density.

Jet quenching was first observed at RHIC. Although full reconstruction of jets was not yet possible in the busy heavy-ion environment, the phenomenon was clearly visible as a suppression in the yield of high p_{T} hadrons [29]. Nuclear effects on hard probes are quantified by the nuclear modification factor

$R_{\rm AA}$, which ratio of the yield observed in AA collisions compared to the expectation from pp collisions, taking into account the nuclear thickness from the Glauber model discussed in Section 5. For example, as a function of p_T

$$R_{\rm AA} \equiv \frac{{\rm d}N^{\rm AA}/{\rm d}p_{\rm T}}{\langle T_{\rm AA}\rangle {\rm d}\sigma/{\rm d}p_{\rm T}} = \frac{dN^{\rm AA}/dp_{\rm T}}{\langle N_{\rm coll}\rangle dN/dp_{\rm T}} \quad , \tag{8}$$

where, in the first expression, the denominator contains the inelastic cross section for pp. In the 2nd equivalent form of $R_{\rm AA}$ one sees explicitly that hard processes are expected to scale with $N_{\rm coll}$, such that an $R_{\rm AA}$ of unity would indicate the absence of nuclear effects. Figure 8 (left) shows $R_{\rm AA}$ for the π^0 meson in central AuAu collisions, which shows a suppression of nearly a factor of 5. This suppression is nearly independent of $p_{\rm T}$, which, in the absence of any modification to the fragmentation pattern of jets, would indicate that partons lose a constant fraction of their $p_{\rm T}$. No such suppression is observed in deuteron-gold (dAu) collisions, as would arise if the presence of a nucleus in the initial state was responsible for the effect. In addition to single particle yields, one can look at back-to-back correlations from high $p_{\rm T}$ particles from dijet correlations, as shown in Fig. 8 (right). A high $p_{\rm T}$ "trigger" charged hadron is selected in the range $4 < p_{\rm T} < 6$ GeV. The yield of hadron "partners" with $p_{\rm T}$ lower than the trigger particle, but still above a hard cutoff of 2 GeV, is then studied as function of azimuthal angle between the two particles. One observes a nearly complete disappearance of back-to-back correlations in central AuAu collisions, that once again is not observed in dAu. The standard interpretation of this disappearance is that the selection of a high $p_{\rm T}$ trigger particle preferentially selects jets that come from the surface of the overlap zone, and hence suffer little quenching. As a result, the recoiling jet traverses a large medium path-length, on average, and hence suffers a large quenching.

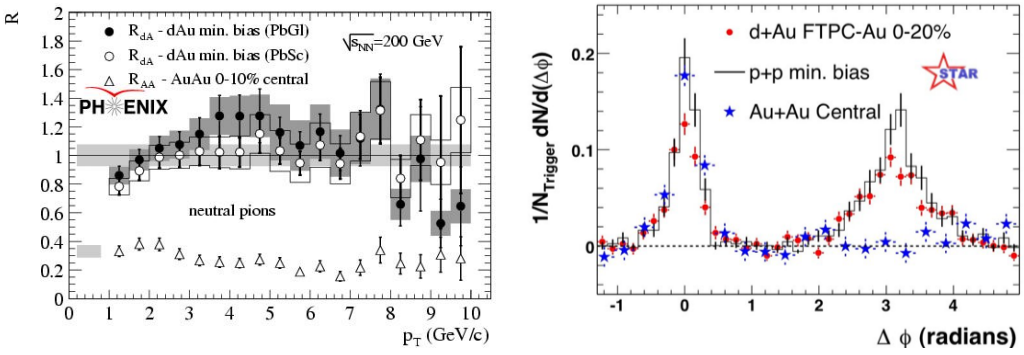

Fig. 8: First observations of jet quenching at RHIC. Left: Nuclear modification factor of π^0 in AuAu and dAu collisions at 200 GeV from PHENIX [30]. Right: Disappearance of back-to-back high $p_{\rm T}$ dihadron azimuthal correlations in AuAu collisions compared to pp and dAu collisions with STAR [31].

Due to the larger cross section for hard processes, as well as the availability of hermetic calorimeters, full reconstruction of jets in heavy ions became possible for the first time at the LHC. Figure 9 (left) shows jet $R_{\rm AA}$ measured over a very wide range of $p_{\rm T}$ out to nearly 1 TeV. A similar approximate flatness of $R_{\rm AA}$ is seen, as was observed for hadrons at RHIC. However, if one looks at the charged hadrons [32], as shown in Fig. 9 (right), one sees that with the much large range of $p_{\rm T}$ that can be measured at the LHC, that the trend is not nearly as flat as was observed at RHIC. In fact, the $R_{\rm AA}$ of hadrons rises rather steeply nearly reaching unity above 100 GeV, albeit with non-negligible uncertainties. How can the measurements of jets and charged hadrons be reconciled? For them to give a different qualitative behavior requires that there be interplay between quenching and the fragmentation pattern of jets.

An advantage of full jet reconstruction is that their fragmentation pattern can then be measured rather directly. Figure 10 (left) shows the jet fragmentation function, defined as the distribution of the ratio (z) of charged hadron $p_{\rm T}$ to the associated jet $p_{\rm T}$, for pp and various centrality selections of PbPb

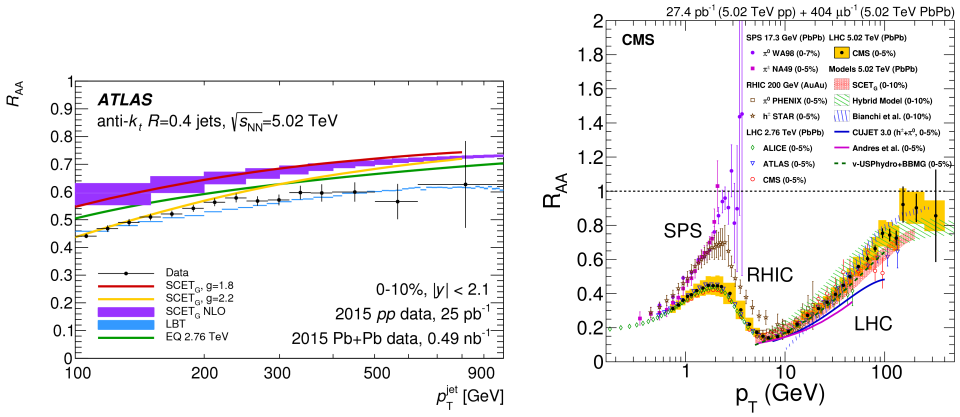

Fig. 9: Left: Jet R_{AA} vs. p_{T} in central PbPb collisions from ATLAS [33]. Right: Charged hadron R_{AA} vs. p_{T} in central PbPb collisions from CMS [32].

collisions [34]. One can study the modification of the fragmentation pattern by forming $R_{D(z)}$, the ratio of the PbPb and pp measurements, as shown in Fig. 10 (right). The fragmentation pattern shows a characteristic modification. The enhancement at low z is generally understood to be the response of the medium to the deposited energy. The depletion and rise at intermediate and high z are still debated. One possibility is that gluons are preferentially suppressed compared to quarks, due to their larger color factor. Because of the softer fragmentation function of gluons, the resulting sample of jets, which initially contains a roughly even mix of quarks and gluons, is biased towards quark jets. Another possibility is that there is a direct dependence of quenching of the fragmentation pattern, via the parton shower. Partons that shower into a larger number of secondary partons may effectively interact with the medium more, compared to lower multiplicity showers. Understanding these details of parton-medium interactions is at the forefront of current investigations, which involve, for example, measurements of jet substructure [35].

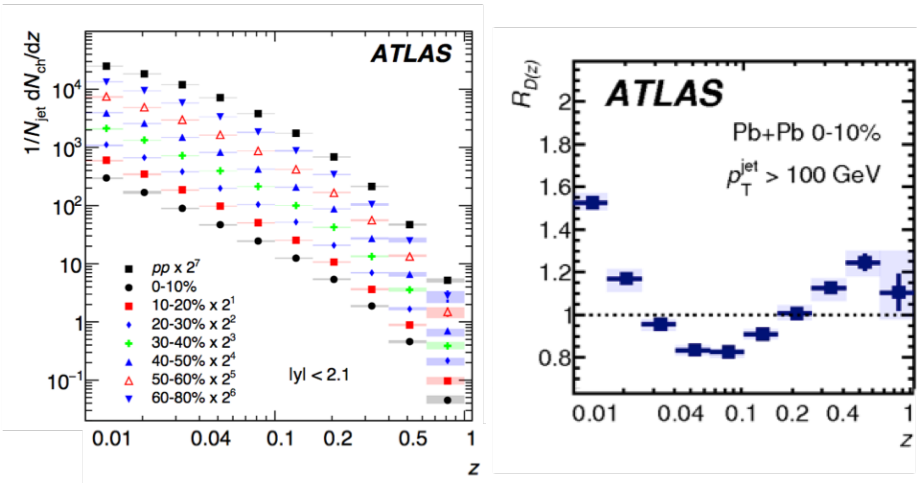

Fig. 10: Left: Jet fragmentation functions measured in pp and in different centrality selections of PbPb collisions [34]. Right: Ratio of jet fragmentation function in central PbPb to pp collisions [34].

10 Quarkonium dissociation

In addition to jets, hard processes can also produce massive states such as quarkonia, bound states of a quark and its anti-quark. The most abundant of these is the J/ψ meson, a bound state of c and $\bar{\text{c}}$. Since the charm mass is above the threshold for thermal production in the QGP, one expects it to be produced only via hard scattering. Just as electric charges are screened by a (QED) plasma, Matsui and Satz predicted that the attractive QCD potential between the quarks could be screened in the QGP [36]. Moreover, different quarkonium states have quite different binding energies, which implies that they should dissociate at different temperatures. For this reason, quarkonia were proposed as a thermometer of the QGP.

In line with this expectation, the J/ψ meson was found to be suppressed in fixed target PbPb collisions at the CERN SPS [37]. As shown in Fig. 11 (left), the data are plotted vs. the pathlength L, another Glauber derived quantity which scales with collision centrality. As a baseline, the J/ψ yield is normalized by that of Drell-Yan, which is not expected to be modified by nuclear effects. The J/ψ shows a rather abrupt suppression in central PbPb events, which is not seen in more peripheral events, nor in lighter collision systems. This was interpreted as one of the earliest observed signatures of QGP formation. The more loosely bound excited state, ψ(2S), was also measured. It shows a suppression starting from even less central events. Naively, one would interpret this as the ψ(2S) dissociating at a lower temperature, a phenomenon often referred to as "sequential melting". As it turns out the more modern interpretation of ψ(2S) data is more nuanced, as it is so fragile that other final state effects likely come into play.

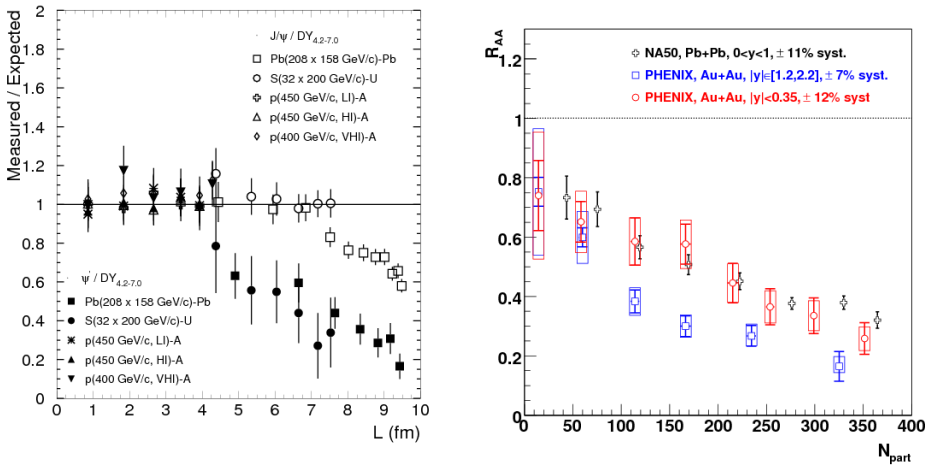

Fig. 11: Left: Suppression of the J/ψ and ψ(2S) in PbPb collisions, and the corresponding lack of suppression in lighter collision systems, by the NA50 experiment [37]. The data are plotted vs. the nuclear pathlength and normalized to the expectation from Drell-Yan events. Right: Nuclear suppression factor for J/ψ vs. the number of participants for a central and a more forward rapidity selection, as measured by PHENIX [38].

Initial results from RHIC showed a similar suppression factor for the J/ψ, despite the larger collision energy, as seen in the mid-rapidity results in Fig. 11 (right), this time evaluated in terms of R_{AA} vs. N_{part}. These results would be most naturally understood as a saturation of the dissociation effect, such that an increase in temperatures would not lead to further suppression. A measurement at forward rapidity showed a larger suppression factor, however, that was not anticipated (also shown in the same figure). Interestingly, when even more forward rapidity measurements were performed at an order of magnitude larger collision energy (2.76 TeV, as compared to 200 GeV), a much reduced suppression was observed, as shown in Fig. 12 (left). This effect is now understood to correspond to an additional mechanism of J/ψ production referred to as statistical regeneration [39]. Although charm quarks are rel-

state	J/ψ	ψ(2S)	ϒ(1S)	ϒ(2S)	ϒ(3S)
mass (GeV)	3.10	3.68	9.46	10.0	10.4
radius (fm)	0.25	0.45	0.14	0.28	0.39

atively rare at the SPS, at RHIC energies they are abundant enough that charm quarks can combine with anti-charm quarks produced in a different nucleon-nucleon scattering. This effect is more pronounced at mid-rapidity, where the charm quark density is larger, explaining the relatively smaller suppression factor, compared to forward rapidity, observed at RHIC. The charm cross section also grows rather rapidly with collision energy, explaining the further reduced suppression observed at the LHC.

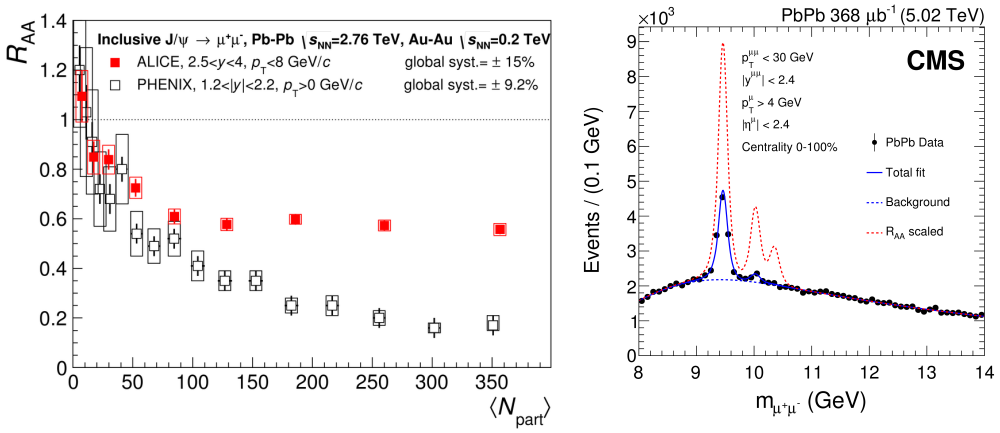

Fig. 12: Left: $R_{\rm AA}$ of J/ψ at vs. $N_{\rm part}$ in 2.76 TeV PbPb collisions at the LHC, compared to 200 GeV collisions at RHIC [40]. Right: CMS measurement of the invariant mass spectrum of dimuon pairs showing the peaks of the three lowest lying ϒ states in PbPb, compared to the expectation from pp collisions [41].

While the statistical regeneration picture provides a satisfactory explanation to the J/ψ puzzle, it certainly complicates a simple interpretation in terms of the afore-mentioned sequential melting picture. In addition to the charmonia states, the large collision energies provided by the LHC enabled the study of the bottomonia states with reasonably high statistics, for the first time. The mass and radii of the five quarkonia states that are readily measured in heavy ions are listed in Table 10. The ϒ family has three states that are readily measurable, compared to two for charmonia. The ground state ϒ(1S) is the most tightly bound of the 5 states, while the ϒ(2S) and ϒ(3S) states are roughly comparable to the J/ψ and ψ(2S), respectively. In contrast to charmonia though, bottom quark production is sufficiently rare that statistical regeneration effects are expected to be small. Figure 12 (right) shows the dimuon invariant mass spectrum in PbPb collisions in the vicinity of the ϒ resonances [41]. The PbPb data are compared to a line-shape based on pp collisions, i.e., in the absence of nuclear effects. A strong suppression is observed for all 3 states, with R values smaller than 0.15 for both excited states (the 3S state is only an upper limit). The 1S state is less suppressed, with an $R_{\rm AA}$ value of around 0.45, as expected from the sequential melting picture. Moreover a sizable fraction of ϒ(1S) come from feed-down of excited states (both S-wave and P-wave), which may account for the majority of its suppression. The ϒ measurements therefore support the scenario of sequential melting of the onium states in the QGP. A quantitative extraction of the medium temperature requires a thorough evaluation of so called cold nuclear matter effects, which will be introduced in the next and final section.

11 Small systems

Thus far I have mainly discussed results in ion-ion collisions, using pp collisions as a baseline from which to benchmark nuclear effects. In Section 9 I mentioned that dAu collisions were studied, to ensure

that jet quenching did not arise from the presence of a heavy nucleus in the initial state (see Fig. 8). Similarly, in Section 10, we have seen that the anomalous suppression of the charmonium states is not observed in lighter collision systems, such as pA with various types of heavy nuclei. This is not to say, however, that nuclear effects are entirely absent from these lighter systems. In this section, I revisit some of the observables that have been discussed, taking a deeper look at their behavior in pA, as well as in pp collisions.

One class of initial state nuclear effects that has been known since the 1970s is modification of the parton distributions in nuclei. The pattern of these modifications, illustrated in Fig. 13 (left), is quantified by looking at the ratio of nuclear PDF to the free nucleon one, as function of Bjorken x. At low x one observes a depletion of partons, known as shadowing, which essentially corresponds to partons obscuring one another in the nucleus. At intermediate values of x, one sees a compensating anti-shadowing effect, where the PDF is enhanced in the nucleus. Finally, at large values of x, one again sees a depletion, referred to as the EMC effect, whose physical explanation remains debated to this day. As for free nucleons, the nuclear PDFs are universal. One may then ask whether data are consistent with various global fits of nuclear PDFs, or whether there are other types of nuclear effects potentially at play. Figure 13 (right) shows the nuclear modification factor in pPb ($R_{\rm pPb}$) for jets, as a function of $p_{\rm T}$ from ATLAS [42] and CMS [43]. The results are compatible with a light enhancement from anti-shadowing, which is dominant in the region of Bjorken x probed by this measurement. By varying the kinematics of dijets, one is able to span a larger range of Bjorken x, including the shadowing, anti-shadowing and EMC regions. The results are also compatible with global fits to the nuclear parton distributions, such as those in Ref. [44].

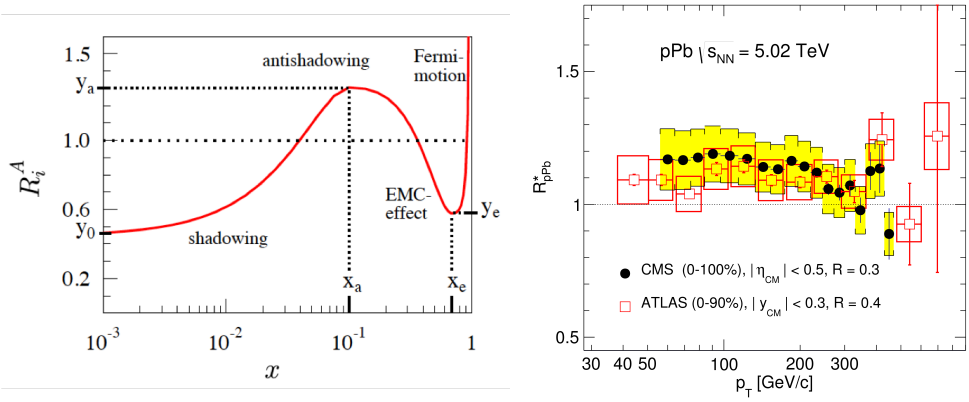

Fig. 13: Left: A schematic of nuclear modifications to the parton distributions, as a function of Bjorken x [45]. Right: Nuclear modification factor of jets in pPb collisions, as a function of $p_{\rm T}$ [43].

In the charmonia sector, one is also able to reproduce the nuclear modification factor based on parameterization of nuclear PDF effects. In Fig. 14, the $R_{\rm pPb}$ of J/ψ is plotted as function of rapidity [46]. While nuclear effects are modest at backward rapidity, the Pb-going direction, they are rather sizable at forward rapidity, which is the proton-side going direction. These effects are consistent with the level of shadowing found in commonly used global fits (the first three theory calculations on the plots) [44,45,47]. Although the suppression due to shadowing is smaller than that observed in PbPb collisions, clearly it needs to be taken into account in the interpretation of quarkonium dissociation measurements.

Whereas both the jet and J/ψ data can be described by nuclear PDF parameterizations, there are effects that cannot be easily described only be these effects. Fig 14 (right) shows the ratio of the yield of the excited state ϒ(2S) to the ground state ϒ(1S), as a function of event multiplicity, for 3 collision systems: pp, pPb and PbPb [49]. One observes that this ratio varies with multiplicity not only in PbPb collisions, where such an effect is expected based on the corresponding variation of collision centrality,

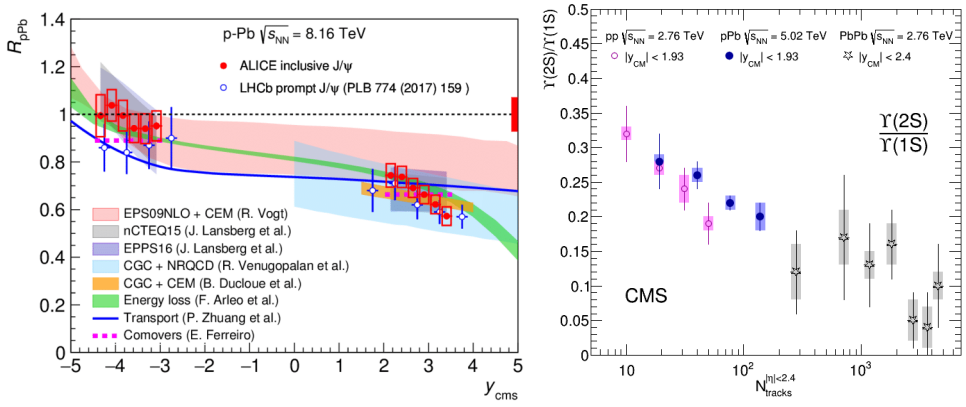

Fig. 14: Left: Nuclear modification factor R_{pPb} of J/ψ, as a function of rapidity [46], compared to calculations using various parameterizations of nuclear PDF effects, as well as other models. Right: Ratio of the yield of the excited state ϒ(2S) to the ground state ϒ(1S), as a function of event multiplicity, for 3 collision systems: pp, pPb and PbPb [48].

but also in pPb, and even in pp. Although the statistical precision of the data is not sufficient to tell whether the data fall along the same trend, clearly there are multiplicity dependent effects present in all three systems. The interpretation of this data remains a subject of debate.

Such multiplicity dependent effects in small systems have also been observed in the flow sector. The elliptic flow can be measured there, by using a technique known as multi-particle cumulants. Figure 15 shows v_2 as a function of event multiplicity in pp, pPb and PbPb. Aside from two-particle correlation results, which are sensitive to a "non-flow" signal from jet correlations, the results for correlations amongst different numbers of particles are very close to one another. This supports collectivity as the origin of the effect. In PbPb collisions, one observes an increasing v_2 with multiplicity, which can be traced back to the centrality dependence of the collisions, as discussed in Section 8. Surprisingly, however, one also observes a non-zero flow appearing in both pp and pPb, which saturates above a multiplicity of about 100 charged particles, and which is of comparable magnitude in the two systems. Just as the higher order flow coefficients (see Fig. 7) reflect event-by-event fluctuations of the overlap zone in PbPb collisions, the non-zero v_2 observed in smaller systems is thought to share a similar origin. The collectivity is thought to arise from fluctuations of color fields on the sub-nucleonic scale. A similar effect is observed in the ratios of strange particle to light hadron yields in Fig. 5 (right). Together these results tend to support the surprising conclusion that there is collectivity and chemical equilibrium over an extended domain, even in light systems. These effects are not captured by standard event generators of proton collisions, and thus could have an impact on measurements thought to be far outside the domain of heavy ion physics, where such ideas have arisen. Detailed studies of the origin of such effects in small systems represents one of the most interesting and active current frontiers in the field.

References

[1] J. C. Collins and M. J. Perry, "Superdense Matter: Neutrons Or Asymptotically Free Quarks?," *Phys. Rev. Lett.* **34** (1975) 1353.

[2] N. Cabibbo and G. Parisi, "Exponential Hadronic Spectrum and Quark Liberation," *Phys. Lett.* **59B** (1975) 67–69.

[3] Y. Aoki *et al.*, "The QCD transition temperature: results with physical masses in the continuum limit II.," *JHEP* **06** (2009) 088, arXiv:0903.4155 [hep-lat].

[4] A. Bazavov *et al.*, "The chiral and deconfinement aspects of the QCD transition," *Phys. Rev.* **D85** (2012) 054503, arXiv:1111.1710 [hep-lat].

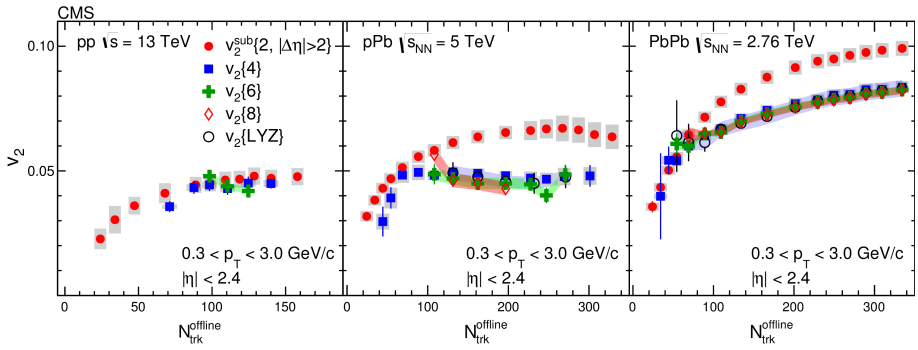

Fig. 15: The elliptic flow coefficient v_2 measured by various methods, as function of event multiplicity in pp (left), pPb (middle) and PbPb (right) collisions [49].

[5] R. Hagedorn, "Statistical thermodynamics of strong interactions at high energies," *Nuovo Cimento, Suppl.* **3** no. CERN-TH-520, (1965) 147–186. http://cds.cern.ch/record/346206.

[6] P. Braun-Munzinger and J. Stachel, "The quest for the quark-gluon plasma," *Nature* **448** (2007) 302–309.

[7] R. C. Hwa and X. N. Wang, eds., *Quark-gluon plasma. Vol. 3*. 2004.

[8] R. J. Glauber, "Cross-sections in deuterium at high-energies," *Phys. Rev.* **100** (1955) 242–248.

[9] M. L. Miller, K. Reygers, S. J. Sanders, and P. Steinberg, "Glauber modeling in high energy nuclear collisions," *Ann. Rev. Nucl. Part. Sci.* **57** (2007) 205–243, arXiv:nucl-ex/0701025 [nucl-ex].

[10] **ALICE** Collaboration, B. Abelev *et al.*, "Centrality determination of Pb-Pb collisions at $\sqrt{s_{\rm NN}}$ = 2.76 TeV with ALICE," *Phys. Rev.* **C88** (2013) 044909, arXiv:1301.4361 [nucl-ex].

[11] **ALICE** Collaboration, J. Adam *et al.*, "Centrality dependence of the charged-particle multiplicity density at midrapidity in Pb-Pb collisions at $\sqrt{s_{\rm NN}}$ = 5.02 TeV," *Phys. Rev. Lett.* **116** (2016) 222302, arXiv:1512.06104 [nucl-ex].

[12] J. D. Bjorken, "Highly Relativistic Nucleus-Nucleus Collisions: The Central Rapidity Region," *Phys. Rev.* **D27** (1983) 140–151.

[13] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjostrand, "Parton Fragmentation and String Dynamics," *Phys. Rept.* **97** (1983) 31–145.

[14] G. Marchesini and B. R. Webber, "Simulation of QCD Jets Including Soft Gluon Interference," *Nucl. Phys.* **B238** (1984) 1–29.

[15] E. Fermi, "High-energy nuclear events," *Prog. Theor. Phys.* **5** (1950) 570–583.

[16] A. Andronic, P. Braun-Munzinger, K. Redlich, and J. Stachel, "Decoding the phase structure of QCD via particle production at high energy," *Nature* **561** (2018) 321–330, arXiv:1710.09425 [nucl-th].

[17] F. Becattini, P. Castorina, J. Manninen, and H. Satz, "The Thermal Production of Strange and Non-Strange Hadrons in e+ e- Collisions," *Eur. Phys. J.* **C56** (2008) 493–510, arXiv:0805.0964 [hep-ph].

[18] R. Stock, "Hadronization revisited: The Dynamics behind hadro-chemical equilibrium," *PoS* **CPOD2006** (2006) 040, arXiv:nucl-th/0703050 [nucl-th].

[19] F. Becattini and J. Manninen, "Strangeness production from SPS to LHC," *J. Phys.* **G35** (2008) 104013, arXiv:0805.0098 [nucl-th].

[20] **ALICE** Collaboration, J. Adam *et al.*, "Enhanced production of multi-strange hadrons in high-multiplicity proton-proton collisions," *Nature Phys.* **13** (2017) 535–539, arXiv:1606.07424 [nucl-ex].

[21] H. Song, S. A. Bass, U. Heinz, T. Hirano, and C. Shen, "200 A GeV Au+Au collisions serve a nearly perfect quark-gluon liquid," *Phys. Rev. Lett.* **106** (2011) 192301, arXiv:1011.2783 [nucl-th]. [Erratum: Phys. Rev. Lett.109,139904(2012)].

[22] **STAR** Collaboration, J. Adams *et al.*, "Azimuthal anisotropy in Au+Au collisions at $\sqrt{s_{\rm NN}}$ = 200-GeV," *Phys. Rev.* **C72** (2005) 014904, arXiv:nucl-ex/0409033 [nucl-ex].

[23] **ALICE** Collaboration, J. Adam *et al.*, "Pseudorapidity dependence of the anisotropic flow of charged particles in Pb-Pb collisions at $\sqrt{s_{\rm NN}} = 2.76$ TeV," *Phys. Lett.* **B762** (2016) 376–388, arXiv:1605.02035 [nucl-ex].

[24] B. Alver and G. Roland, "Collision geometry fluctuations and triangular flow in heavy-ion collisions," *Phys. Rev.* **C81** (2010) 054905, arXiv:1003.0194 [nucl-th]. [Erratum: Phys. Rev.C82,039903(2010)].

[25] **WMAP** Collaboration, G. Hinshaw *et al.*, "Three-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: temperature analysis," *Astrophys. J. Suppl.* **170** (2007) 288, arXiv:astro-ph/0603451 [astro-ph].

[26] **CMS** Collaboration, S. Chatrchyan *et al.*, "Studies of azimuthal dihadron correlations in ultra-central PbPb collisions at $\sqrt{s_{\rm NN}} = 2.76$ TeV," *JHEP* **02** (2014) 088, arXiv:1312.1845 [nucl-ex].

[27] T. Hirano, U. W. Heinz, D. Kharzeev, R. Lacey, and Y. Nara, "Hadronic dissipative effects on elliptic flow in ultrarelativistic heavy-ion collisions," *Phys. Lett.* **B636** (2006) 299–304, arXiv:nucl-th/0511046 [nucl-th].

[28] J. D. Bjorken, "Energy Loss of Energetic Partons in Quark - Gluon Plasma: Possible Extinction of High p(t) Jets in Hadron - Hadron Collisions," 1982. FERMILAB-PUB-82-059-THY, FERMILAB-PUB-82-059-T.

[29] **PHENIX** Collaboration, "Suppression of hadrons with large transverse momentum in central Au+Au collisions at $\sqrt{s_{\rm NN}}$ = 130-GeV," *Phys. Rev. Lett.* **88** (2002) 022301, arXiv:nucl-ex/0109003 [nucl-ex].

[30] **PHENIX** Collaboration, S. S. Adler *et al.*, "Absence of suppression in particle production at large transverse momentum in $\sqrt{s_{\rm NN}}$ = 200-GeV d + Au collisions," *Phys. Rev. Lett.* **91** (2003) 072303, arXiv:nucl-ex/0306021 [nucl-ex].

[31] **STAR** Collaboration, "Disappearance of back-to-back high p_T hadron correlations in central Au+Au collisions at $\sqrt{s_{\rm NN}}$ = 200-GeV," *Phys. Rev. Lett.* **90** (2003) 082302, arXiv:nucl-ex/0210033 [nucl-ex].

[32] **CMS** Collaboration, V. Khachatryan *et al.*, "Charged-particle nuclear modification factors in PbPb and pPb collisions at $\sqrt{s_{\rm NN}} = 5.02$ TeV," *JHEP* **04** (2017) 039, arXiv:1611.01664 [nucl-ex].

[33] **ATLAS** Collaboration, M. Aaboud *et al.*, "Measurement of the nuclear modification factor for inclusive jets in Pb+Pb collisions at $\sqrt{s_{\rm NN}} = 5.02$ TeV with the ATLAS detector," arXiv:1805.05635 [nucl-ex].

[34] **ATLAS** Collaboration, M. Aaboud *et al.*, "Measurement of jet fragmentation in Pb+Pb and pp collisions at $\sqrt{s_{\rm NN}} = 2.76$ TeV with the ATLAS detector at the LHC," *Eur. Phys. J.* **C77** (2017) 379, arXiv:1702.00674 [hep-ex].

[35] **CMS** Collaboration, A. M. Sirunyan *et al.*, "Measurement of the Splitting Function in pp and Pb-Pb Collisions at $\sqrt{s_{\rm NN}} = 5.02$ TeV," *Phys. Rev. Lett.* **120** (2018) 142302, arXiv:1708.09429 [nucl-ex].

[36] T. Matsui and H. Satz, "J/ψ Suppression by Quark-Gluon Plasma Formation," *Phys. Lett.* **B178** (1986) 416–422.

[37] **NA50** Collaboration, B. Alessandro *et al.*, "psi-prime production in Pb-Pb collisions at 158-GeV/nucleon," *Eur. Phys. J.* **C49** (2007) 559–567, arXiv:nucl-ex/0612013 [nucl-ex].

[38] **PHENIX** Collaboration, A. Adare *et al.*, "J/ψ Production vs Centrality, Transverse Momentum, and Rapidity in Au+Au Collisions at $\sqrt{s_{\rm NN}} = 200$ GeV," *Phys. Rev. Lett.* **98** (2007) 232301, arXiv:nucl-ex/0611020 [nucl-ex].

[39] R. L. Thews, M. Schroedter, and J. Rafelski, "Enhanced J/ψ production in deconfined quark matter," *Phys. Rev.* **C63** (2001) 054905, arXiv:hep-ph/0007323 [hep-ph].

[40] **ALICE** Collaboration, J. Adam *et al.*, "Differential studies of inclusive J/ψ and ψ(2S) production at forward rapidity in Pb-Pb collisions at $\sqrt{s_{\rm NN}} = 2.76$ TeV," *JHEP* **05** (2016) 179, arXiv:1506.08804 [nucl-ex].

[41] **CMS** Collaboration, A. M. Sirunyan *et al.*, "Measurement of nuclear modification factors of Υ(1S), Υ(2S), and Υ(3S) mesons in PbPb collisions at $\sqrt{s_{\rm NN}} = 5.02$ TeV," arXiv:1805.09215 [hep-ex].

[42] **ATLAS** Collaboration, G. Aad *et al.*, "Centrality and rapidity dependence of inclusive jet production in $\sqrt{s_{\rm NN}} = 5.02$ TeV proton-lead collisions with the ATLAS detector," *Phys. Lett.* **B748** (2015) 392–413, arXiv:1412.4092 [hep-ex].

[43] **CMS** Collaboration, V. Khachatryan *et al.*, "Measurement of inclusive jet production and nuclear modifications in pPb collisions at $\sqrt{s_{\rm NN}} = 5.02$ TeV," *Eur. Phys. J.* **C76** (2016) 372, arXiv:1601.02001 [nucl-ex].

[44] K. J. Eskola, P. Paakkinen, H. Paukkunen, and C. A. Salgado, "EPPS16: Nuclear parton distributions with LHC data," *Eur. Phys. J.* **C77** (2017) 163, arXiv:1612.05741 [hep-ph].

[45] K. J. Eskola, H. Paukkunen, and C. A. Salgado, "EPS09: A New Generation of NLO and LO Nuclear Parton Distribution Functions," *JHEP* **04** (2009) 065, arXiv:0902.4154 [hep-ph].

[46] **ALICE** Collaboration, S. Acharya *et al.*, "Inclusive J/ψ production at forward and backward rapidity in p-Pb collisions at $\sqrt{s_{\rm NN}}$ = 8.16 TeV," *JHEP* **07** (2018) 160, arXiv:1805.04381 [nucl-ex].

[47] K. Kovarik *et al.*, "nCTEQ15 - Global analysis of nuclear parton distributions with uncertainties in the CTEQ framework," *Phys. Rev.* **D93** (2016) 085037, arXiv:1509.00792 [hep-ph].

[48] **CMS** Collaboration, S. Chatrchyan *et al.*, "Event activity dependence of Y(nS) production in $\sqrt{s_{\rm NN}}$=5.02 TeV pPb and $\sqrt{s}$=2.76 TeV pp collisions," *JHEP* **04** (2014) 103, arXiv:1312.6300 [nucl-ex].

[49] **CMS** Collaboration, V. Khachatryan *et al.*, "Evidence for collectivity in pp collisions at the LHC," *Phys. Lett.* **B765** (2017) 193–220, arXiv:1606.06198 [nucl-ex].